



HCEO WORKING PAPER SERIES

Working Paper



HUMAN CAPITAL AND
ECONOMIC OPPORTUNITY
GLOBAL WORKING GROUP

The University of Chicago
1126 E. 59th Street Box 107
Chicago IL 60637

www.hceconomics.org

Commuting, Migration and Local Employment Elasticities*

Ferdinando Monte[†]
Georgetown University

Stephen J. Redding[‡]
Princeton University

Esteban Rossi-Hansberg[§]
Princeton University

September 25, 2017

Abstract

We provide theory and evidence that the elasticity of local employment to a labor demand shock is heterogeneous depending on the commuting openness of the local labor market. We develop a quantitative general equilibrium model that incorporates spatial linkages in goods markets (trade) and factor markets (commuting and migration). We quantify this model to match the observed gravity equation relationships for trade and commuting. We find that empirically-observed reductions in commuting costs generate welfare gains of around 3.3 percent. We provide separate evidence in support of the model's predictions using decompositions of employment changes, million dollar plants, and trade shocks.

JEL CLASSIFICATION: F12, F14, R13, R23

1 Introduction

Agents spend about 8% of their workday commuting to and from work.¹ They make this significant daily investment, to live and work in different locations, so as to balance their living costs and residential amenities with the wage they can obtain at their place of employment. The ability of firms in a location to attract workers depends, therefore, not only on the ability to attract local residents through migration, but also on the ability to attract commuters from other, nearby, locations. Together, migration and commuting determine the response of local employment to a local labor demand shock, which we term the *local employment elasticity*. This elasticity is of great policy interest since it determines the impact of local policies, such as transport infrastructure investments, local taxation and regional development programs. Estimating its magnitude has been the subject of a large empirical literature on local labor markets, which has considered a variety of sources of local labor demand shocks, including sectoral composition (Bartik shocks), productivity, international trade, natural resource abundance and business cycle fluctuations, as discussed further below.² In this paper we explore the determinants and characteristics of the local employment elasticity (and the corresponding *local resident elasticity*).

*Much of this research was undertaken while Ferdinando Monte was visiting the International Economics Section (IES) at Princeton. We are grateful to the IES and Princeton more generally for research support. We are also grateful to the editor, four anonymous referees, and conference and seminar participants for helpful comments and suggestions.

[†]McDonough School of Business, 37th and O Streets, NW, Washington, DC 20057. ferdinando.monte@georgetown.edu.

[‡]Dept. Economics and WWS, Fisher Hall, Princeton, NJ 08544. 609 258 4016. reddings@princeton.edu.

[§]Dept. Economics and WWS, Fisher Hall, Princeton, NJ 08544. 609 258 4024. erossi@princeton.edu.

¹See for example Redding and Turner (2015).

²For a survey of this empirical literature, see Moretti (2011).

We begin by developing a quantitative general equilibrium model that incorporates spatial linkages between locations in both goods markets (trade) and factor markets (commuting and migration). We show that there is no single local employment elasticity. Instead, the local employment elasticity is an endogenous variable that differs across locations depending on their linkages to one another in goods and factor markets. Quantifying our model to county-level data for the United States, we find that the elasticity of local employment with respect to local productivity shocks varies from around 0.5 to 2.5. Therefore, an average local employment elasticity estimated from cross-section data can be quite misleading when used to predict the impact of a local shock or policy on any individual county. We use our quantitative model to understand the systematic determinants of the local employment elasticity and show that a large part of the variation results from differences in commuting links between a location and its neighbors. We then propose variables that can be included in reduced-form regressions to improve their ability to account for the heterogeneity in local employment responses without imposing the full structure of our model.

Our theoretical framework allows for an arbitrary number of locations that can differ in productivity, amenities and geographical relationship to one another. The spatial distribution of economic activity is driven by a tension between productivity differences and home market effects (forces for the concentration of economic activity) and an inelastic supply of land and commuting costs (dispersion forces). Commuting allows workers to access high productivity employment locations without having to live there and hence alleviates the congestion effect in such high productivity locations. We show that the resulting commuting flows between locations exhibit a gravity equation relationship with a much higher distance elasticity than for goods flows, suggesting that moving people is more costly than moving goods across geographic space. We discipline our quantitative spatial model to match the observed gravity equation relationships for trade in goods and commuting flows as well as the observed cross-section distributions of employment, residents and wages across U.S. counties. Given the observed data on wages, employment by workplace, commuting flows and land area, and a parameterization of trade and commuting costs, we show that our model can be used to recover unique values of the unobserved location fundamentals (productivity and amenities) that exactly rationalize the observed data as an equilibrium of the model. We show how the values of these observed variables in an initial equilibrium can be used to undertake counterfactuals for the impact of local labor demand shocks (captured by productivity shocks in our model) and for the impact of changes in trade or commuting costs.

An advantage of our explicitly modeling the spatial linkages between locations is that our framework can be taken to data on local economic activity at different levels of spatial aggregation. In contrast, existing research that does not explicitly model these spatial linkages is faced with a trade-off when studying local labor markets. On the one hand, larger spatial units have the advantage of reducing the unmodeled spatial linkages between locations. On the other hand, larger spatial units reduce the ability to make inferences about local labor markets. Existing research has typically resolved this trade-off by defining commuting zones (CZs) that are aggregations of counties chosen to minimize commuting flows. However, there exists no choice of boundaries for local labor markets that completely eliminates commuting. For example, should Princeton, NJ be considered part of New York's or Philadelphia's local labor market? Some of its residents commute to New York, but others commute to Philadelphia. Furthermore, often researchers are interested in the effects of policy interventions at spatial scales smaller than CZs. For example, transport

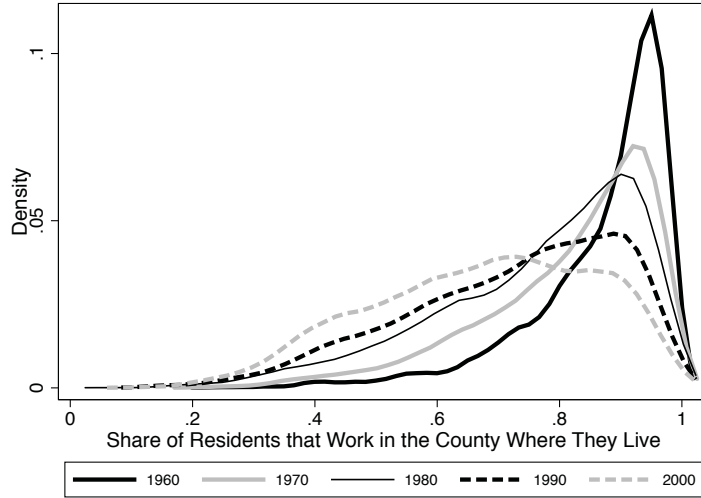


Figure 1: Kernel densities of the share of residents that work in the county where they live

infrastructure improvements can disproportionately affect commuting between counties within CZs. As a preliminary illustration of the size and heterogeneity of these commuting interactions between counties, Figure 1 displays kernel densities of the share of residents that work in the same county where they live (the “residence own commuting share”). In 1960, when the interstate highway system had only recently begun to be constructed, U.S. counties were relatively closed, as shown by the concentration of density at high values, with the median own commuting share equal to 91 percent. Forty years later, the picture is rather different, as shown by the marked shift in density towards lower values, with a median own commuting share equal to 69 percent.³ More generally, the substantial heterogeneity in own commuting shares, evident in Figure 1, already suggests that counties are likely to differ substantially in the extent to which labor market shocks and policies spill over across their boundaries.

We show that our results are robust both theoretically and empirically. From a theoretical perspective, we show that heterogeneous local employment elasticities are not specific to our theoretical model, but rather are a more generic prediction of an entire class of theoretical models consistent with a gravity equation for commuting flows. From an empirical perspective, we show that we continue to find substantial heterogeneity in these local employment elasticities when we incorporate the variable land supply elasticities from Saiz (2010). Introducing this second source of heterogeneity generates more variation in local resident elasticities but does not reduce the variation in local employment elasticities. This pattern of results is intuitive. The housing supply elasticity matters less for employment than for residents, because commuting allows individuals to work in locations with inelastic housing supplies without actually having to live there and pay the resulting high land prices. Therefore, a high productivity location with an inelastic housing supply can increase employment through commuting without requiring substantial changes in the number of residents. An important policy implication is that improvements in commuting technologies provide an alternative to the relaxation of housing supply elasticities in facilitating the allocation of workers to

³In Figure 1, we measure commuting using the share of residents that work in the same county where they live. We show below that this measure is model consistent and it is the only measure available at the county level over the entire 1960-2000 period. The distributions for each decade shown in Figure 1 are statistically significantly different from those in the decade immediately before using a Kolomogorov-Smirnov test. These shifts in the distribution over time are also apparent if we weight counties by residents or use CZs instead of counties (as shown in Section B.6 of the web appendix).

productive locations. While this possibility has been informally discussed in the existing literature on housing supply elasticities (as for example in Hsieh and Moretti 2017), we are the first, as far as we are aware, to provide quantitative empirical evidence on the relevance of commuting for the response of employment to local labor demand shocks.

We provide three separate pieces of empirical evidence independent of our model for the importance of commuting for employment changes using shift-share decompositions, the location of million dollar plants (MDPs following Greenstone, Hornbeck and Moretti 2010, henceforth GHM), and international trade shocks (following Autor, Dorn and Hanson 2013, henceforth ADH). We first decompose the variation in employment across both counties at a given time and within counties over time into the contributions of own residents, own commuting, other residents and other commuting. We show that the commuting terms account for around two thirds of the cross-section variation and around one half of the time-series changes in employment.

We next use quasi-experimental variation from comparisons of winner and runner-up counties in competitions for MDPs from GHM. As runner-up counties are those that have survived a long selection process but narrowly lost the competition, they are likely to provide a better counterfactual for a winner county than other comparison counties. We find no evidence of statistically significant differences in employment growth in winner and runner-up counties prior to the announcement of MDPs, but substantial differences emerge after these announcements. Consistent with the predictions of our model, we find greater increases in employment from the positive labor demand shock from the opening of a MDP in counties with more open commuting markets.⁴

Finally, we use the trade shock provided by China’s emergence into the global economy as another source of variation in local labor demand. We reproduce the central empirical exercise from ADH for the ratio of manufacturing employment to population at the county level, but allow the treatment effect of the shock to vary with the openness of counties to commuting. This central exercise from ADH is harder to interpret in terms of our model, because the local labor demand shock is concentrated in manufacturing, affects all counties simultaneously, and is spatially correlated. Nonetheless, we again find substantial heterogeneity in estimated treatment effects that is captured by our measure of openness to commuting.

Having provided these three separate pieces of evidence in support of commuting as a source of heterogeneity in employment responses to local labor demand shocks, we show that our model provides a platform for evaluating the counterfactual effects of changes in trade and commuting costs. Building on approaches in the international trade literature (e.g. Head and Ries, 2001), we show how observed data on commuting flows over time can be used to back out the empirical distribution of implied changes in commuting costs. We use this empirical distribution to undertake counterfactuals for empirically-plausible changes in commuting costs. For example, reducing commuting costs by the median reduction from 1990-2010 (a reduction of 12 percent), we find an increase in welfare of 3.3 percent. The commuting technology facilitates a separation of workplace and residence, enabling people to work in high productivity locations and live in high amenity locations. Therefore reducing commuting costs increases the concentration of employment in locations that were net importers of commuters in the initial equilibrium (e.g. Manhattan)

⁴These results are consistent with the empirical findings, in another context, of Manning and Petrongolo (2011), which shows that local development policies are fairly ineffective in raising local unemployment outflows, because labor markets overlap, and the associated ripple effects in applications largely dilute the impact of local stimulus across space.

and enhances the clustering of residents in locations that initially were net exporters of commuters (e.g. parts of New Jersey). This logic seems to suggest that commuting might be important only for larger cities in the U.S., but this is in fact not the case. Although the changes in employment as a result of eliminating commuting are well explained by initial commuting intensity, this intensity cannot be easily proxied for using empirical controls such as land area, size or housing supply elasticities.

Our paper is related to several existing literatures. In international trade, our work relates to quantitative models of costly trade in goods following Eaton and Kortum (2002) and extensions. Our research also contributes to the economic geography literature on costly trade in goods and factor mobility, which typically uses variation *across regions or systems of cities*.⁵ Our work also contributes to the urban economics literature on the costly movement of people (commuting), which typically uses variation *within cities*.⁶ In contrast, we develop a framework in which an arbitrary set of regions are connected in both goods markets (through costly trade) and labor markets (through migration and commuting), and which encompasses both within and across-city interactions. Although incorporating costly goods trade and commuting is a natural idea, our first main contribution is to develop a tractable framework that captures these forces and is amenable to both analytic and quantitative analysis. Our second main contribution is to quantify this framework using disaggregated data on trade and commuting for the United States and to show how it provides a platform for evaluating a range of counterfactual interventions. Our third main contribution is to establish theoretically and empirically the importance of spatial interactions between locations (in particular through commuting) in determining the local economic effects of local labor demand shocks.

Our paper is also related to the large empirical literature on local labor markets, which has estimated the effects of local labor demand shocks.⁷ Each of the papers in this literature is concerned with evaluating the local impact of economic shocks using data on finely-detailed spatial units. However, these spatial units are typically treated as independent observations in reduced-form regressions, with little attention paid to the linkages between these spatial units in goods and labor markets, and hence with little consideration of the distinction between employment and residents introduced by endogenous commuting decisions. A key contribution of our paper is to show that understanding these spatial linkages is central to evaluating the local impact of these and other economic shocks.

The remainder of the paper is structured as follows. Section 2 develops our theoretical framework. Section 3 discusses the quantification of the model using U.S. data and reports summary statistics on commuting between counties. Section 4 uses the model to quantify the heterogeneity in local employment elasticities across U.S. counties. Section 5 presents independent evidence in support of the importance

⁵See Krugman (1991), Hanson (1996, 2005), Helpman (1998), Fujita et al. (1999), Rossi-Hansberg (2005), Redding and Sturm (2008), Moretti and Klein (2014), Allen and Arkolakis (2014), Caliendo, et al. (2014), Desmet and Rossi-Hansberg (2014) and Redding (2016).

⁶See Alonso (1964), Mills (1967), Muth (1969), Lucas and Rossi-Hansberg (2002), Desmet and Rossi-Hansberg (2013), Behrens, et al. (2014), Ahlfeldt, et al. (2015), Allen, Arkolakis and Li (2015) and Monte (2016).

⁷Examples include: (a) GHM (2010)'s analysis of million dollar plants; (b) Autor, Dorn and Hanson (2013), which examines the local economic effects from the international trade shock provided by China's emergence into global markets; (c) the many empirical studies that use the Bartik (1991) instrument, which interacts aggregate industry shocks with locations' industry employment shares, including Diamond (2016) and Notowidigdo (2013); (d) research on the geographic incidence of macroeconomic shocks, such as the 2008 Financial Crisis and Great Recession, including Mian and Sufi (2014) and Yagan (2016); and (e) work on the impact of natural resource discoveries on the spatial distribution of economic activity, as in Michaels (2011) and Feyrer, Mansur and Sacerdote (2015). Other related research on local labor demand shocks includes Blanchard and Katz (1992), Bound and Holzer (2000), and Busso, Gregory and Kline (2013).

of commuting for employment changes using shift-share decompositions, million dollar plants (MDPs), and international trade shocks. Section 6 undertakes counterfactuals for changes in commuting costs and Section 7 summarizes our conclusions. A web appendix contains the derivations of theoretical results, the proofs of propositions, additional robustness tests, and a description of the data sources and manipulations.

2 The Model

We develop a spatial general equilibrium model in which locations are linked in goods markets through trade and in factor markets through migration and commuting. The economy consists of a set of locations $n, i \in N$. Each location n is endowed with a supply of land (H_n). Following the new economic geography literature, we begin by interpreting land as geographical land area, which is necessarily in perfectly inelastic supply. We later extend our analysis to interpret land as developed land area, which has a positive supply elasticity that we allow to differ across locations. The economy as a whole is populated by a measure \bar{L} of workers, each of whom is endowed with one unit of labor that is supplied inelastically.

2.1 Preferences and Endowments

Workers are geographically mobile and have heterogeneous preferences for locations. Each worker chooses a pair of residence and workplace locations to maximize their utility taking as given the choices of other firms and workers.⁸ The preferences of a worker ω who lives and consumes in location n and works in location i are defined over final goods consumption ($C_{n\omega}$), residential land use ($H_{n\omega}$), an idiosyncratic amenities shock ($b_{ni\omega}$) and commuting costs (κ_{ni}), according to the Cobb-Douglas form,⁹

$$U_{ni\omega} = \frac{b_{ni\omega}}{\kappa_{ni}} \left(\frac{C_{n\omega}}{\alpha} \right)^\alpha \left(\frac{H_{n\omega}}{1 - \alpha} \right)^{1 - \alpha}, \quad (1)$$

where $\kappa_{ni} \in [1, \infty)$ is an iceberg commuting cost in terms of utility.¹⁰ The idiosyncratic amenities shock ($b_{ni\omega}$) captures the idea that individual workers can have idiosyncratic reasons for living and working in different locations. We model this heterogeneity in amenities following McFadden (1974) and Eaton and Kortum (2002).¹¹ For each worker ω living in location n and working in location i , idiosyncratic amenities ($b_{ni\omega}$) are drawn from an independent Fréchet distribution,

$$G_{ni}(b) = e^{-B_{ni}b^{-\epsilon}}, \quad B_{ni} > 0, \epsilon > 1, \quad (2)$$

⁸Throughout the following, we use n to denote a worker's location of residence and consumption and i to denote a worker's location of employment and production, unless otherwise indicated.

⁹For empirical evidence using U.S. data in support of the constant housing expenditure share implied by the Cobb-Douglas functional form, see Davis and Ortalo-Magne (2011).

¹⁰Although we model commuting costs in terms of utility, they enter the indirect utility function multiplicatively with the wage, which implies that they are proportional to the opportunity cost of time. Therefore, similar results hold if commuting costs are instead modeled as a reduction in effective units of labor, as discussed in Subsection B.15 of the web appendix.

¹¹A long line of research models location decisions using preference heterogeneity, as in Artuc, Chaudhuri and McClaren (2010), Kennan and Walker (2011), Grogger and Hanson (2011), Moretti (2011) and Busso, Gregory and Kline (2013). Modeling individual heterogeneity in terms of productivity rather than preferences results in a similar specification, as discussed in Section B.14 of the web appendix.

where the scale parameter B_{ni} determines the average amenities from living in location n and working in location i , and the shape parameter $\epsilon > 1$ controls the dispersion of amenities. The idiosyncratic amenities shock ($b_{ni\omega}$) implies that different workers make different choices about their workplace and residence locations when faced with the same prices and wages. All workers ω residing in location n and working in location i receive the same wage and make the same consumption and residential land choices. Hence we suppress the implicit dependence on ω except where important.¹²

To isolate the effects of introducing commuting, we model goods consumption as in the new economic geography literature. The goods consumption index in location n is a constant elasticity of substitution (CES) function of consumption of a continuum of tradable varieties sourced from each location i ,

$$C_n = \left[\sum_{i \in N} \int_0^{M_i} c_{ni}(j)^\rho dj \right]^{\frac{1}{\rho}}, \quad \sigma = \frac{1}{1 - \rho} > 1. \quad (3)$$

Utility maximization implies that equilibrium consumption in location n of each variety sourced from location i is given by $c_{ni}(j) = \alpha X_n P_n^{\sigma-1} p_{ni}(j)^{-\sigma}$, where X_n is aggregate expenditure in location n ; P_n is the price index dual to (3), and $p_{ni}(j)$ is the “cost inclusive of freight” price of a variety j produced in location i and consumed in location n .¹³

Utility maximization also implies that a fraction $(1 - \alpha)$ of worker income is spent on residential land. We assume that this land is owned by immobile landlords, who receive worker expenditure on residential land as income, and consume only goods where they live. This assumption allows us to incorporate general equilibrium effects from changes in the value of land, without introducing a mechanical externality into workers’ location decisions from the local redistribution of land rents.¹⁴ Using this assumption, total expenditure on consumption goods equals the fraction α of the total income of residents plus the entire income of landlords (which equals the fraction $(1 - \alpha)$ of the total income of residents):

$$P_n C_n = \alpha \bar{v}_n R_n + (1 - \alpha) \bar{v}_n R_n = \bar{v}_n R_n \quad (4)$$

where \bar{v}_n is the average labor income of residents across employment locations; and R_n is the measure of residents. Land market clearing determines the land price (Q_n) as a function of the supply of land (H_n):

$$Q_n = (1 - \alpha) \frac{\bar{v}_n R_n}{H_n}. \quad (5)$$

¹²Our baseline specification focuses on a single worker type with a Fréchet distribution of idiosyncratic preferences for tractability, which results in similar choice probabilities to the logit model. In Subsection B.10 of the web appendix, we generalize our analysis to multiple worker types z with different Fréchet scale and shape parameters, which results in similar choice probabilities to the mixed logit model of McFadden and Train (2000).

¹³In Subsection B.12 of the web appendix, we show how this standard specification can be further generalized to introduce non-traded consumption goods.

¹⁴In Subsection B.13 of the web appendix, we show that the model has similar properties if landlords consume both consumption goods and residential land, although expressions are less elegant. In the web appendix, we also report the results of a robustness test, in which we instead assume that land is partially owned locally and partially owned by a national portfolio that redistributes land rents to workers throughout the economy (as in Caliendo et al. 2014).

2.2 Production

Again to isolate the effects of introducing commuting, we model production as in the new economic geography literature. Tradeable varieties are produced using labor under monopolistic competition and increasing returns to scale. To produce a variety, a firm must incur a fixed cost of F and a constant variable cost that depends on a location's productivity A_i .¹⁵ Therefore the total amount of labor ($l_i(j)$) required to produce $x_i(j)$ units of a variety j in location i is $l_i(j) = F + x_i(j)/A_i$.¹⁶

Profit maximization implies that equilibrium prices are a constant mark-up over marginal cost: $p_{ni}(j) = \left(\frac{\sigma}{\sigma-1}\right) \frac{d_{ni}w_i}{A_i}$, where w_i is the wage in location i . Combining profit maximization and zero profits, equilibrium output of each variety is equal to a constant: $x_i(j) = A_i F (\sigma - 1)$. This constant equilibrium output of each variety and labor market clearing together imply that the total measure of produced varieties (M_i) is proportional to the measure of employed workers (L_i), $M_i = \frac{L_i}{\sigma F}$.

2.3 Goods Trade

The model implies a gravity equation for bilateral trade between locations. Using the CES expenditure function, the equilibrium pricing rule, and the measure of firms, the share of location n 's expenditure on goods produced in location i is

$$\pi_{ni} = \frac{M_i p_{ni}^{1-\sigma}}{\sum_{k \in N} M_k p_{nk}^{1-\sigma}} = \frac{L_i (d_{ni} w_i / A_i)^{1-\sigma}}{\sum_{k \in N} L_k (d_{nk} w_k / A_k)^{1-\sigma}}. \quad (6)$$

Therefore trade between locations n and i depends on bilateral trade costs (d_{ni}) in the numerator (“bilateral resistance”) and on trade costs to all possible sources of supply k in the denominator (“multilateral resistance”). Equating revenue and expenditure, and using zero profits, workplace income in each location equals total expenditure on goods produced in that location, namely,¹⁷

$$w_i L_i = \sum_{n \in N} \pi_{ni} \bar{v}_n R_n. \quad (7)$$

Using the equilibrium pricing rule and labor market clearing, the price index dual to the consumption index (3) can be expressed as

$$P_n = \frac{\sigma}{\sigma-1} \left(\frac{1}{\sigma F}\right)^{\frac{1}{1-\sigma}} \left[\sum_{i \in N} L_i (d_{ni} w_i / A_i)^{1-\sigma} \right]^{\frac{1}{1-\sigma}} = \frac{\sigma}{\sigma-1} \left(\frac{L_n}{\sigma F \pi_{nn}}\right)^{\frac{1}{1-\sigma}} \frac{d_{nn} w_n}{A_n}. \quad (8)$$

¹⁵We assume a representative firm within each location. However, it is straightforward to generalize the analysis to introduce firm heterogeneity with an untruncated Pareto productivity distribution following Melitz (2003).

¹⁶In Subsection B.14 of the web appendix, we generalize the production technology to include intermediate inputs (as in Krugman and Venables 1995 and Eaton and Kortum 2002), commercial land use and physical capital. As heterogeneous local employment elasticities are a generic prediction of gravity in commuting, they also hold under this production structure.

¹⁷Although a location's total workplace income equals total expenditure on the goods that it produces, total residential income can differ from total workplace income (because of commuting). Therefore total workplace income need not equal total residential expenditure, which implies that total exports need not equal total imports. When we take the model to the data, we also allow total residential expenditure to differ from total residential income, which provides another reason for trade deficits. Within the model, these two variables can diverge if landlords own land in different locations from where they consume. This is how we interpret trade deficits in the empirical section.

where the second equality uses (6) to write the price index (8).

2.4 Labor Mobility and Commuting

Workers are geographically mobile and choose their pair of residence and workplace locations to maximize their utility. Given our specification of preferences (1), the indirect utility function for a worker ω residing in location n and working in location i is

$$U_{ni\omega} = \frac{b_{ni\omega} w_i}{\kappa_{ni} P_n^\alpha Q_n^{1-\alpha}}. \quad (9)$$

Indirect utility is a monotonic function of idiosyncratic amenities ($b_{ni\omega}$) and these amenities have a Fréchet distribution. Therefore, the indirect utility for a worker living in location n and working in location i also has a Fréchet distribution: $G_{ni}(U) = e^{-\Psi_{ni} U^{-\epsilon}}$, where $\Psi_{ni} = B_{ni} (\kappa_{ni} P_n^\alpha Q_n^{1-\alpha})^{-\epsilon} w_i^\epsilon$. Each worker selects the bilateral commute that offers her the maximum utility, where the maximum of Fréchet distributed random variables is itself Fréchet distributed. Using these distributions of utility, the probability that a worker chooses to live in location n and work in location i is

$$\lambda_{ni} = \frac{B_{ni} (\kappa_{ni} P_n^\alpha Q_n^{1-\alpha})^{-\epsilon} w_i^\epsilon}{\sum_{r \in N} \sum_{s \in N} B_{rs} (\kappa_{rs} P_r^\alpha Q_r^{1-\alpha})^{-\epsilon} w_s^\epsilon} \equiv \frac{\Phi_{ni}}{\Phi}. \quad (10)$$

Therefore the idiosyncratic shock to preferences $b_{ni\omega}$ implies that individual workers choose different bilateral commutes when faced with the same prices (P_n, Q_n, w_i), commuting costs (κ_{ni}) and location characteristics (B_{ni}). Other things equal, workers are more likely to live in location n and work in location i , the lower the consumption goods price index (P_n) and land prices (Q_n) in n , the higher the wages (w_i) in i , the more attractive average amenities (B_{ni}), and the lower the commuting costs (κ_{ni}).

Summing these probabilities across workplaces i for a given residence n , we obtain the overall probability that a worker resides in location n (λ_n^R). Similarly, summing across residences n for a given workplace i , we obtain the overall probability that a worker works in location i (λ_i^L). So,

$$\lambda_n^R = \frac{R_n}{\bar{L}} = \sum_{i \in N} \lambda_{ni} = \sum_{i \in N} \frac{\Phi_{ni}}{\Phi}, \quad \text{and} \quad \lambda_i^L = \frac{L_n}{\bar{L}} = \sum_{n \in N} \lambda_{ni} = \sum_{n \in N} \frac{\Phi_{ni}}{\Phi}, \quad (11)$$

where national labor market clearing corresponds to $\sum_n \lambda_n^R = \sum_i \lambda_i^L = 1$.

The average income of a worker living in n depends on the wages in all the nearby employment locations. To construct this average income of residents, note first that the probability that a worker commutes to location i conditional on living in location n is

$$\lambda_{ni|n}^R \equiv \frac{\lambda_{ni}}{\lambda_n^R} = \frac{B_{ni} (w_i/\kappa_{ni})^\epsilon}{\sum_{s \in N} B_{ns} (w_s/\kappa_{ns})^\epsilon}. \quad (12)$$

Equation (12) implies a commuting gravity equation, with an elasticity of commuting flows with respect to commuting costs (κ_{ni}) of $-\epsilon$. Therefore, the probability of commuting to location i conditional on living in location n depends on the wage (w_i), amenities (B_{ni}) and commuting costs (κ_{ni}) for workplace i in the numerator (“bilateral resistance”), as well as the wage (w_s), amenities (B_{ns}) and commuting costs (κ_{ns})

for all other possible workplaces s in the denominator (“multilateral resistance”). This gravity equation prediction is consistent with the existing empirical literature on commuting and migration, including McFadden (1974), Grogger and Hanson (2011) and Kennan and Walker (2011). In Subsection B.9 of the web appendix, we show that heterogeneous local employment elasticities are a generic prediction of the whole class of models consistent with a gravity equation for commuting flows.

Using these conditional commuting probabilities, we obtain the following labor market clearing condition that equates the measure of workers employed in location i (L_i) with the measure of workers choosing to commute to that location, namely,

$$L_i = \sum_{n \in N} \lambda_{ni|n}^R R_n, \quad (13)$$

Expected worker income conditional on living in location n is then equal to the wages in all possible workplaces weighted by the probabilities of commuting to those workplaces conditional on living in n , or

$$\bar{v}_n = \sum_{i \in N} \lambda_{ni|n}^R w_i. \quad (14)$$

Hence expected worker income (\bar{v}_n) is high in locations that have low commuting costs (low κ_{ni}) to high-wage employment locations.¹⁸

Finally, population mobility implies that expected utility is the same for all pairs of residence and workplace and equal to expected utility for the economy as a whole. That is,

$$\bar{U} = \mathbb{E}[U_{niw}] = \Gamma\left(\frac{\epsilon-1}{\epsilon}\right) \left[\sum_{r \in N} \sum_{s \in N} B_{rs} (\kappa_{rs} P_r^\alpha Q_r^{1-\alpha})^{-\epsilon} w_s^\epsilon \right]^{\frac{1}{\epsilon}} \quad \text{all } n, i \in N, \quad (15)$$

where \mathbb{E} is the expectations operator and the expectation is taken over the distribution for the idiosyncratic component of utility and $\Gamma(\cdot)$ is the Gamma function.

Although expected utility is equalized across all pairs of residence and workplace, real wages differ as a result of preference heterogeneity. Workplaces and residences face upward-sloping supply functions for workers and residents respectively (the choice probabilities (11)). Each workplace must pay higher wages to increase commuters’ real income and attract additional workers with lower idiosyncratic amenities for that workplace. Similarly, each residential location must offer a lower cost of living to increase commuters’ real income and attract additional residents with lower idiosyncratic amenities for that residence. Bilateral commutes with attractive characteristics (high workplace wages and low residence cost of living) attract additional commuters with lower idiosyncratic amenities, until expected utility (taking into account idiosyncratic amenities) is the same across all bilateral commutes.

2.5 General Equilibrium

The general equilibrium of the model can be referenced by the following vector of six variables $\{w_n, \bar{v}_n, Q_n, L_n, R_n, P_n\}_{n=1}^N$ and a scalar \bar{U} . Given this equilibrium vector and scalar, all other endogenous variables of the model can be determined. This equilibrium vector solves the following six sets of equations:

¹⁸We treat agents and workers as synonymous, which abstracts from a labor force participation decision, and enables us to isolate the implications of introducing commuting into the standard new economic geography model.

income equals expenditure (7), average residential income (14), land market clearing (5), workplace choice probabilities (11 for L_n), residence choice probabilities (11 for R_n), and price indices (8). The last condition needed to determine the scalar \bar{U} is the labor market clearing condition, $\bar{L} = \sum_{n \in N} R_n = \sum_{n \in N} L_n$.

In Section B.3 of the web appendix, we provide conditions for the existence and uniqueness of the general equilibrium. We show that the system of equations for general equilibrium in our model can be written in the form required to apply the existence and uniqueness results for gravity equation models from Allen, Arkolakis and Li (2016). In developing our model above, we have followed the new economic geography literature in modeling agglomeration forces through love of variety and increasing returns to scale. In Section B.4 of the web appendix, we show that our new economic geography model is isomorphic to a version of Eaton and Kortum (2002) and Redding (2016) with commuting and external economies of scale or a version of Armington (1969) with commuting and external economies of scale (as in Allen and Arkolakis 2014 and Allen, Arkolakis and Li 2015).

2.6 Computing Counterfactuals

We use our quantitative framework to solve for the counterfactual effects of changes in the exogenous variables of the model (productivity A_n , amenities B_{ni} , commuting costs κ_{ni} , and trade costs d_{ni}) without having to necessarily determine the unobserved values of these exogenous variables. Instead, in the web appendix, we show that the system of equations for the counterfactual changes in the endogenous variables of the model can be written solely in terms of the observed values of variables in an initial equilibrium (employment L_i , residents R_i , workplace wages w_n , average residential income \bar{v}_n , trade shares π_{ni} , and commuting probabilities λ_{ni}). This approach uses observed bilateral commuting probabilities to capture unobserved bilateral commuting costs and amenities. Similarly, if bilateral trade shares between locations are available, they can be used to capture unobserved bilateral trade frictions (as in Dekle, Eaton and Kortum 2007). However, since bilateral trade data are only available at a higher level of aggregation than the counties we consider in our data, we make some additional parametric assumptions to solve for implied bilateral trade shares between counties, as discussed below. Throughout this theoretical section, we assume for simplicity that residential income equals expenditure. However, when taking the model to the data, we allow for intertemporal trade deficits that are treated as exogenous in our counterfactuals, as in Dekle, Eaton and Kortum (2007) and Caliendo and Parro (2015), as discussed further below.

3 Data and Measurement

Our empirical analysis combines data from a number of different sources for the United States.¹⁹ From the Commodity Flow Survey (CFS), we use data on bilateral trade and distances shipped for 123 CFS regions. Data on bilateral commuting between counties come from the American Community Survey (ACS) 2006-10 (for our main calibration) and U.S. Census 1960-2000 (for various other statistics). From the Bureau of Economic Analysis (BEA), we use data on employment and wages by workplace. We combine these data sources with a variety of other Geographical Information Systems (GIS) data. We use our data on employment and commuting to calculate the implied number of residents and their average income by

¹⁹See Section D of the web appendix for further discussion of the data sources and definitions.

county. First, from commuter market clearing (13), we obtain the number of residents (R_n) using data on the number of workers (L_n) and commuting probabilities conditional on living in each location ($\lambda_{ni|n}^R$). Second, we use these conditional commuting probabilities, together with county wages, to obtain average residential income (\bar{v}_n) as defined in equation (14).

3.1 Goods Trade

In the Commodity Flow Survey (CFS) data, we observe bilateral trade flows and distances shipped between 123 CFS regions and trade deficits for each these CFS regions.²⁰ To quantify the model at the county level, we allocate the deficit for each CFS region across the counties within that region according to their shares of CFS residential income (as measured by $\bar{v}_i R_i$). Using the resulting trade deficits for each county (D_i), we solve the equality between income and expenditure (7) for unobserved county productivities (A_i):

$$w_i L_i - \sum_{n \in N} \frac{L_i (d_{ni} w_i / A_i)^{1-\sigma}}{\sum_{k \in N} L_k (d_{nk} w_k / A_k)^{1-\sigma}} [\bar{v}_n R_n + D_n] = 0, \quad (16)$$

where we observe (or have solved for) wages (w_i), employment (L_i), average residential income (\bar{v}_i), residents (R_i) and trade deficits (D_i).

Given the elasticity of substitution (σ), our measures for ($w_i, L_i, \bar{v}_i, R_i, D_i$) and a parameterization of trade costs ($d_{ni}^{1-\sigma}$), equation (16) provides a system of N equations that can be solved for a unique vector of N unobserved productivities (A_i), as shown formally in Proposition B.1 in Section B.5 of the web appendix. We parameterize bilateral trade costs as a constant elasticity function of distance and a stochastic error ($d_{ni}^{1-\sigma} = dist_{ni}^{-\psi(\sigma-1)} \tilde{e}_{ni}$) for all pairs with positive trade, while the model implies prohibitive trade costs for pairs with zero trade. We use the model's gravity equation and observed bilateral trade between CFS regions to estimate the composite parameter $-\psi(\sigma-1) = -1.29$, as discussed further in Section B.5 of the web appendix. We assume a central value for the elasticity of substitution between varieties from the existing empirical literature of $\sigma = 4$, which is in line with the estimates of this parameter using price and expenditure data in Broda and Weinstein (2006), and implies $\psi = 0.43$. We use these estimated parameters and equation (16) to solve for the unique unobserved productivities (A_i) and generate the model's predictions for bilateral trade flows between counties (X_{ni}) from equation (6).²¹ As a check on this specification, we aggregate the model's predictions for trade between counties within pairs of CFS regions, and compare these predictions to the data on CFS bilateral trade. As shown in Section B.5 of the web appendix, we find a strong and approximately log linear relationship between the model's predictions and the data, which is tighter for the larger trade values that account for most of aggregate trade.

²⁰Other recent studies using the CFS data include Caliendo et. al (2014), Duranton, Morrow and Turner (2014) and Dingel (2015). The CFS is a random sample of plant shipments within the United States (foreign trade shipments are not included). CFS regions are the smallest geographical units for which this random sample is representative, which precludes constructing bilateral trade flows between smaller geographical units using the sampled shipments.

²¹We find that measured productivity (A_i) is correlated with observable proxies for productivity, such as access to natural water. Regressing $\log A_i$ on a dummy indicating if a county is in the 10% of counties closest to the ocean or a navigable river we find a positive and statistically significant estimated coefficient (standard error) of 0.21 (0.02) for the ocean and 0.04 (0.02) for a navigable river. The data on distances are from Rappaport and Sachs (2003).

3.2 Commuting Flows

We start by providing evidence on the quantitative relevance of commuting as a source of spatial linkages between counties and CZ's. Our main data source is the American Community Survey (ACS), which reports county-to-county worker flows for 2006-2010. To abstract from business trips that are not between a worker's usual place of residence and workplace, we define commuting flows as those of less than 120 kilometers in each direction (a round trip of 240 kilometers).²² We supplement these data with information on bilateral commuting between counties from the 1990-2000 population censuses and data on the share of the residents of each county that work where they live from the 1960-1980 population censuses.

In Table 1, we report some descriptive statistics on commuting flows from 2006-10. We find that commuting beyond county boundaries is substantial and varies in importance across locations. For the median county, around 27 percent of its residents work outside the county and around 20 percent of its workers live outside the county. For the county at the 95th percentile, these two figures rise to 59 and 43 percent respectively. In Table 1, we report unweighted statistics to capture the heterogeneity in commuting patterns across geographical locations (counties). As we show in Subsection B.6 of the web appendix, commuting linkages remain substantial and heterogeneous if we instead weight counties by residents to capture heterogeneity across residents.

The share of *residents* who work where they live (the “residence own commuting share,” $\lambda_{ii|i}^R$) and the share of *workers* who live where they work (the “workplace own commuting share,” $\lambda_{ii|i}^L$) can be related to one another and the ratio of employment to residents (L_i/R_i) using the commuter market clearing condition (see Section B.7 of the web appendix). Both commuting measures are intuitive and they are inversely related to the flow of either residents or workers from other locations. We find that these two measures are positively correlated with one another, as reflected in a statistically significant correlation of 0.60 from 2006-10.²³ We choose the residence own commuting share ($\lambda_{ii|i}^R$) as our baseline measure since it is both model consistent and available in the population census back to 1960. But we show that our results are robust to using either measure, or the average or minimum of these two measures, for the years for which both measures are available.

These differences in openness to commuting imply substantial variation across counties in the ratio of employment to residents (L_i/R_i), which is greater than one for counties that import commuters and less than one for those that export commuters. This ratio ranges from 0.60 at the 5th percentile to 1.18 at the 95th percentile (third row, columns two and eight respectively).²⁴

One might think that using commuting zones (CZs) circumvents the need to incorporate commuting into the analysis, since the boundaries of these areas are drawn to minimize commuting flows. Nevertheless, we find that CZ's provide only an imperfect measure of local labor markets, with substantial commuting

²²The majority of commutes are less than 45 minutes in each direction (Duranton and Turner 2011). In our analysis, we measure distance between counties' centroids. We choose the 120 kilometers threshold based on a change in slope of the relationship between log commuters and log distance at this distance threshold. See Appendix D.1.2 for further discussion.

²³Although Figure 1 above shows that counties have become increasingly open to commuting over time, we also find that relative openness to commuting is strongly persistent across counties over time. Correlating the share of residents that work outside the county where they live (λ^R) in 2000 with the corresponding values for each previous decade back to 1960, we find correlations of 0.96, 0.91, 0.84 and 0.73 respectively.

²⁴In Subsection B.6 of the web appendix, we show that the ratio of employment to residents is not only heterogeneous across counties, but is also hard to explain with standard empirical controls, such as land area, size or supply elasticities for developed land.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Min	p5	p10	p25	p50	p75	p90	p95	Max	Mean
Commuters from Residence County	0.00	0.03	0.06	0.14	0.27	0.42	0.53	0.59	0.82	0.29
Commuters to Workplace County	0.00	0.03	0.07	0.14	0.20	0.28	0.37	0.43	0.81	0.22
County employment/Residents	0.26	0.60	0.67	0.79	0.92	1.02	1.11	1.18	3.88	0.91
Commuters from Residence CZ	0.00	0.00	0.01	0.03	0.07	0.12	0.18	0.22	0.49	0.08
Commuters to Workplace CZ	0.00	0.00	0.01	0.03	0.07	0.10	0.13	0.15	0.25	0.07
CZ Employment/Residents	0.63	0.87	0.91	0.97	1.00	1.01	1.03	1.04	1.12	0.98

Tabulations on 3,111 counties and 709 commuting zones. The first row shows the fraction of residents that work outside the county. The second row shows the fraction of workers who live outside the county. The third row shows the ratio of county employment to county residents. The fourth row shows the fraction of a CZ's residents that work outside the CZ. The fifth row shows the fraction of a CZ's workers that live outside the CZ. The sixth row shows the ratio of CZ employment to CZ residents across all 709 CZs. p5, p10 etc refer to the 5th, 10th etc percentiles of the distribution.

Table 1: Commuting Across Counties and Commuting Zones (Unweighted)

beyond CZ boundaries that again varies in importance across locations. For the median CZ, around 7 percent of its residents work outside the CZ and around 7 percent of its workers live outside the CZ. For the CZ at the 95th percentile, these two figures rise to 22 percent and 15 percent respectively. Although the ratio of employment to residents (L/R) by construction varies less across CZs than across counties, we still find substantial variation from a minimum of 0.63 to a maximum of 1.12 (final row), which we show below is sufficient to generate substantial heterogeneity in local employment elasticities. To provide a point of comparison to Figure 1 above for counties, Figure B.4 in Subsection B.6 of the web appendix shows kernel densities of the share of residents that work in the same CZ where they live for 1990 and 2000. We find the same pattern of an increase in commuting openness over time. In Subsection B.6 of the web appendix, we also show that we continue to find substantial variation in commuting linkages across CZs if we weight them by residents to capture heterogeneity across residents.

We now use our model's prediction of a gravity equation for bilateral commuting probabilities, which using land market clearing (5) and the price index (8) can be written as

$$\lambda_{ni} = \frac{\mathcal{B}_{ni} \left(\frac{L_n}{\pi_{nn}} \right)^{-\frac{\alpha\epsilon}{\sigma-1}} A_n^{\alpha\epsilon} w_n^{-\alpha\epsilon} \bar{v}_n^{-\epsilon(1-\alpha)} \left(\frac{R_n}{H_n} \right)^{-\epsilon(1-\alpha)} w_i^\epsilon}{\sum_{r \in N} \sum_{s \in N} \mathcal{B}_{rs} \left(\frac{L_r}{\pi_{rr}} \right)^{-\frac{\alpha\epsilon}{\sigma-1}} A_r^{\alpha\epsilon} w_r^{-\alpha\epsilon} \bar{v}_r^{-\epsilon(1-\alpha)} \left(\frac{R_r}{H_r} \right)^{-\epsilon(1-\alpha)} w_s^\epsilon}, \quad (17)$$

where $\mathcal{B}_{ni} \equiv B_{ni} \kappa_{ni}^{-\epsilon}$ is a composite parameter that captures the ease of commuting.

Given the preference heterogeneity parameter (ϵ) and our measures for $(L_n, \pi_{nn}, w_n, v_n, R_n, H_n)$, equation (17) provides a system of $N \times N$ system of equations that can be solved for a unique matrix of $N \times N$ values of the ease of commuting (\mathcal{B}_{ni}), as shown in Proposition B.2 in Section B.6 of the web appendix. We model the ease of commuting (\mathcal{B}_{ni}) for all pairs with positive commuting flows as depending on (i) a residence component (\mathbb{B}_n), (ii) a workplace component (\mathbb{B}_i), (iii) a component that is related to distance ($\text{dist}_{ni}^{-\phi}$), and (iv) an orthogonal component (\mathbb{B}_{ni}), such that $\mathcal{B}_{ni} = \mathbb{B}_n \mathbb{B}_i \text{dist}_{ni}^{-\phi} \mathbb{B}_{ni}$, where the model implies prohibitive commuting costs for pairs with zero commuting flows. We estimate the parameter $\phi = 4.43$ in a first step using the model's gravity equation predictions for bilateral commuting flows and

including workplace and residence fixed effects. We estimate the heterogeneity in location preferences ($\varepsilon = 3.30$) in a second step by using the structure of the model to replace the workplace fixed effects with wages raised to this power. In this second step, we instrument wages using productivity to address the fact that wages themselves depend on commuting flows, as discussed further in Section B.6 of the web appendix. We find that the gravity equation provides a good approximation to observed bilateral commuting flows, and the higher estimated distance coefficient for commuting than goods trade (-4.43 versus -1.29) is in line with the idea that transporting people is more costly than transporting goods.

For the one remaining model parameter, the share of housing in consumer expenditure, we assume a central value from Bureau of Economic Analysis of $1 - \alpha = 0.40$ percent.²⁵ Using our assumption of Cobb-Douglas utility and our interpretation of land as geographical land area, in Subsection C.2 of the web appendix, we show that the model’s predictions for land prices are strongly positively correlated with observed county median housing values. In the next section, we also relax these assumptions to introduce a positive supply elasticity for developed land.

4 Local Employment Elasticities

To provide evidence on local employment elasticities, we compute 3,111 counterfactual exercises where we shock each county with a 5 percent productivity shock (holding productivity in all other counties and holding all other exogenous variables constant).²⁶ Figure 2 shows the estimated kernel density for the distribution of the general equilibrium elasticity of employment with respect to the productivity shock across these treated counties (solid blue line). We also show the 95 percent confidence intervals around this estimated kernel density (gray shading). The mean estimated local employment elasticity of around 1.52 is greater than one because of home market effects and commuting. Around this mean, we find substantial heterogeneity in the predicted effects of the productivity shock, which vary from close to 0.5 to almost 2.5. This variation is surprisingly large. It implies that taking a local employment elasticity estimated for one group of counties and applying that elasticity to another group of counties can lead to substantial discrepancies between the true and predicted impacts of a productivity shock.

To provide a point of comparison, Figure 2 also includes the general equilibrium elasticity of residents in a county with respect to the same 5 percent productivity shock in that county (again holding other parameters constant). Again we show the estimated kernel density across the 3,111 treated counties (dashed red line) and the 95 percent confidence intervals (gray shading). We find substantial differences between the employment and residents elasticities, with the residents elasticity having less dispersion and ranging from around 0.2 to 1.2. Since employment and residents can only differ through commuting, this by itself suggests that the heterogeneity in the local employment elasticity is largely driven by commuting links between counties. In Section C.7 of the web appendix, we provide further evidence that this is indeed the case by simulating productivity shocks in a counterfactual world without commuting between

²⁵Using these assumed parameter values, we correlate our measures of residential amenities with observable proxies for this variable. We regress the solutions for the bilateral ease of commuting (\mathcal{B}_{ni}) from equation (17) on residence and workplace fixed effects and bilateral distance. We use the residence fixed effect as our measure of residential amenities. Regressing this measure on violent crimes per resident, we find a negative and statistically significant coefficient (standard error) of -0.48 (0.10). Crime data is from the U.S. Department of Justice (2007).

²⁶We have experimented with shocks of 1% and 10% as well, with essentially unchanged results.

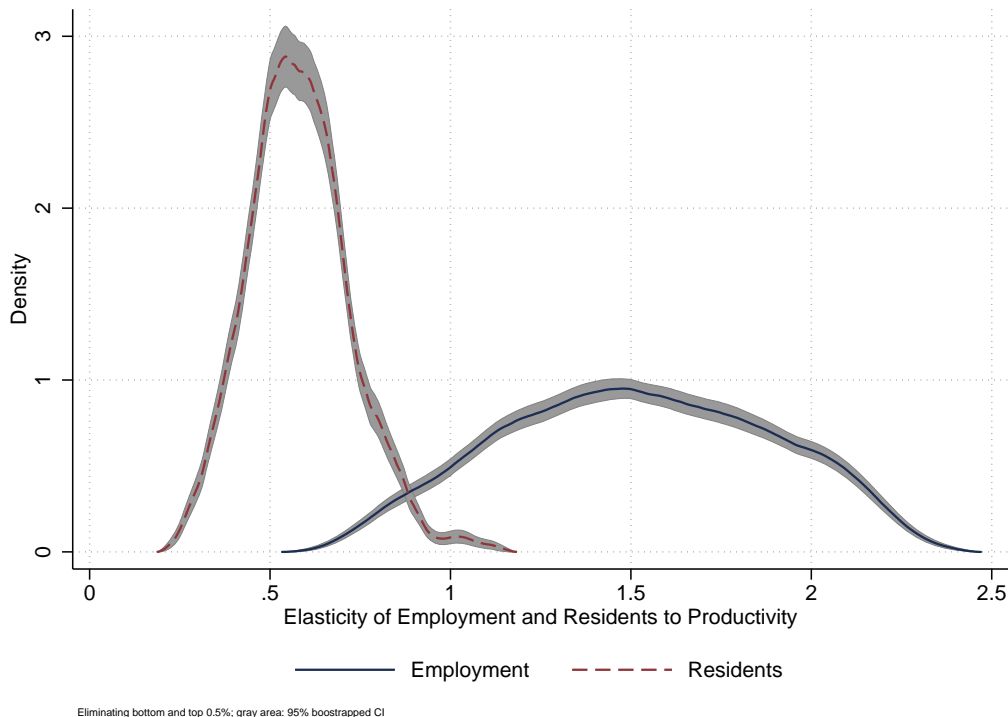


Figure 2: Kernel density for the distribution of employment and residents elasticities in response to a productivity shock across counties

counties. Even in such a counterfactual world, we expect local employment elasticities to be heterogeneous, because counties differ substantially in terms of their initial shares of U.S. employment. However, we find substantially less heterogeneity in local employment elasticities in this counterfactual world than in the actual world with commuting. In fact, the resulting distribution of *employment* (and resident) elasticities is similar to the one for *resident* elasticities in Figure 2.

This heterogeneity in local employment elasticities remains if we shock counties with patterns of spatially correlated shocks reproducing the industrial composition of the U.S. economy (see Subsection C.5 of the web appendix). We also find a similar pattern of results if we replicate our entire quantitative analysis for CZs rather than counties (see Subsection C.11 of the web appendix). Both sets of results are consistent with the fact that heterogeneous local employment elasticities are a generic prediction of a gravity equation for commuting (as shown in Subsection B.9 of the web appendix).

4.1 Explaining the Heterogeneity in Local Employment Elasticities

When we undertake our counterfactual exercises, we solve for the full general equilibrium effect of the productivity shock to each county. To provide intuition for the determinants of these local employment elasticities, in Table 2, we examine the relationship between these general equilibrium elasticities and a range of observed variables. In Column (1), we regress our general equilibrium elasticities on a constant, which captures the mean employment elasticity across the 3,111 treated counties. In Columns (2) through (4), we attempt to explain the heterogeneity in local employment elasticities using standard county controls. In Column (2) we include log county employment as a control for the size of economic activity in a county. In Column (3) we also include log county wages and log county land area. In Column (4) we also include

the average wage and total employment in neighboring counties. Although these controls are all typically statistically significant, we find that they are not particularly successful in explaining the variation in employment elasticities. Adding a constant and all these controls yields an R-squared of only about 0.5 in Column (4). Clearly, there is substantial variation not captured by these controls.

In the remaining columns of the table we attempt to explain the heterogeneity in local employment elasticities using our commuting measures derived from the model. In Column (5), we include the residence own commuting share ($\lambda_{ii|i}^R$) as our baseline measure of commuting linkages.²⁷ The more open the local labor market to commuting, the lower the value of $\lambda_{ii|i}^R$, and the higher the local employment elasticity. This is exactly what we find in Column (5). Furthermore, this variable alone yields a R-squared of 0.89, substantially higher than including all of the standard econometric controls.²⁸ We find a similar pattern of results using the workplace own commuting share ($\lambda_{ii|i}^L$) and the average or minimum of the residence and workplace own commuting shares. Therefore, although our model incorporates several forms of spatial linkages (including trade and migration), we find that the heterogeneity in local employment elasticities is mainly explained by commuting linkages, which is consistent with our gravity equation estimates, where commuting is substantially more local than goods trade.

Even more of the heterogeneity in local employment elasticities can be explained using the partial equilibrium elasticities derived from the model in Section B.8 of the web appendix. In Column (6) we relate the variation in local employment elasticities to the measure of commuting linkages implied by these partial equilibrium elasticities, $\sum_{n \in N} (1 - \lambda_{ni|n}^R) \vartheta_{ni}$, where $\vartheta_{ni} \equiv \lambda_{ni|n}^R R_n / L_i$ is the share of commuters from residence n in workplace i 's employment. We also add the corresponding implied measures of migration and trade linkages, $(\lambda_{ii}/\lambda_{Ri} - \lambda_{Li})$ and $\frac{\partial w_i}{\partial A_i} \frac{A_i}{w_i}$. Including these partial equilibrium measures further increases the R-squared to around 93 percent of the variation in the general equilibrium elasticity. Counties that account for a small share of commuters (small $\lambda_{ni|n}^R$) from their main suppliers of commuters (high ϑ_{ni}) have higher employment elasticities. In Column (7), we use the product of $\frac{\partial w_i}{\partial A_i} \frac{A_i}{w_i}$ and the first two terms rather than each term separately. This restriction yields similar results and confirms the importance of commuting linkages and, to a lesser extent, the interaction between migration and goods linkages. Finally, in the last two columns we combine these partial equilibrium elasticities with the standard controls we used in the first four columns. Clearly, although all variables are significant, these standard controls add little once we control for the partial equilibrium elasticities.²⁹

In sum, Table 2 shows that the heterogeneity in partial equilibrium elasticities is not well explained by standard county controls. In contrast, adding our baseline measure of the openness of the labor market to commuting, or the partial equilibrium elasticities derived from the model, can go a long way towards explaining this heterogeneity.

²⁷See Section B.8 in the web appendix for the derivation our baseline measure and the partial equilibrium elasticities we use below.

²⁸To provide further evidence on the magnitude of these effects, Table C.1 in Section C.3 of the web appendix reports the same regressions as in Table 2 but using standardized coefficients. We find that a one standard deviation change in the own commuting share ($\lambda_{ii|i}^R$) leads to around a one standard deviation change in the local employment elasticity.

²⁹In Subsection B.9 of the web appendix, we report kernel density estimates for the distribution of the partial equilibrium measure of commuting linkages $\sum_{n \in N} (1 - \lambda_{ni|n}^R) \vartheta_{ni}$ that is a generic prediction of any commuting gravity equation. We show a similar distribution of heterogeneous local employment elasticities to that in Figure 2 above, again confirming that this heterogeneity is driven by commuting linkages.

Dependent Variable:	1	2	3	4	5	6	7	8	9
	Elasticity of Employment								
$\log L_i$		-0.003 (0.014)	0.009 (0.012)	-0.054** (0.006)				0.037** (0.004)	0.033** (0.004)
$\log w_i$			-0.201** (0.059)	-0.158** (0.039)				-0.257** (0.016)	-0.263** (0.016)
$\log H_i$			-0.288** (0.021)	-0.172** (0.015)				0.003 (0.009)	0.009 (0.009)
$\log L_{-i}$				0.118** (0.017)				-0.027** (0.009)	-0.027** (0.009)
$\log \bar{w}_{-i}$				0.204* (0.083)				0.163** (0.037)	0.207** (0.038)
$\lambda_{ii i}^R$					-2.047** (0.042)				
$\sum_{n \in N} (1 - \lambda_{Rni}) \vartheta_{ni}$						2.784** (0.192)		2.559** (0.178)	
$\vartheta_{ii} \left(\frac{\lambda_{ii}}{\lambda_{Ri}} - \lambda_{Li} \right)$						0.915** (0.210)		0.605** (0.175)	
$\frac{\partial w_i}{\partial A_i} \frac{A_i}{w_i}$						-1.009** (0.123)		-0.825** (0.150)	
$\frac{\partial w_i}{\partial A_i} \frac{A_i}{w_i} \cdot \sum_{r \in N} (1 - \lambda_{rn r}) \vartheta_{rn}$							1.038** (0.090)		1.100** (0.091)
$\frac{\partial w_i}{\partial A_i} \frac{A_i}{w_i} \cdot \vartheta_{ii} \left(\frac{\lambda_{ii}}{\lambda_{Ri}} - \lambda_{Li} \right)$							-0.818** (0.098)		-0.849** (0.092)
Constant	1.515** (0.034)	1.545** (0.158)	5.683** (0.632)	1.245 (0.797)	2.976** (0.022)	0.840** (0.201)	1.553** (0.087)	1.861** (0.404)	2.064** (0.352)
R^2	0.00	0.00	0.40	0.51	0.89	0.93	0.93	0.95	0.95
N	3,111	3,111	3,111	3,081	3,111	3,111	3,111	3,081	3,081

Note: $L_{-i} \equiv \sum_{n: d_{ni} \leq 120, n \neq i} L_n$ is the total employment in i neighbors whose centroid is no more than 120km away; $\bar{w}_{-i} \equiv \sum_{n: d_{ni} \leq 120, n \neq i} \frac{L_n}{L_{-i}} w_n$ is the weighted average of their workplace wage. Standard errors are clustered by state. * denotes significance at the 5 percent level; ** denotes significance at the 1 percent level.

Table 2: Explaining the general equilibrium local employment elasticities to a 5 percent productivity shock

4.2 Positive Developed Land Supply Elasticities

In the baseline version of the model, we interpret the non-traded good as geographical land, which is necessarily in perfectly inelastic supply. In this section, we develop an extension of the model, in which we interpret the non-traded good as “developed” land and allow for a positive developed land supply elasticity that can differ across locations. The rest of the model remains identical.

We introduce a positive developed land supply elasticity by following Saiz (2010) in assuming that the supply of land (H_n) for each residence n depends on the endogenous price of land (Q_n) as well as on the exogenous characteristics of locations (\bar{H}_n):

$$H_n = \bar{H}_n Q_n^{\eta_n}, \quad (18)$$

where $\eta_n \geq 0$ is the developed land supply elasticity, which we allow to vary across locations; $\eta_n = 0$ is our baseline specification of a perfectly inelastic land supply; and $\eta_n \rightarrow \infty$ is the special case of a perfectly

elastic land supply.

Introducing a positive and heterogeneous developed land supply elasticity only affects one of the conditions for general equilibrium in the model, namely, the land market clearing condition. Using the supply function for land (18) in the land market clearing condition (5), we obtain the following generalization of our earlier expression for the equilibrium price of land (Q_n):

$$Q_n = \left((1 - \alpha) \frac{\bar{v}_n R_n}{\bar{H}_n} \right)^{\frac{1}{1+\eta_n}}. \quad (19)$$

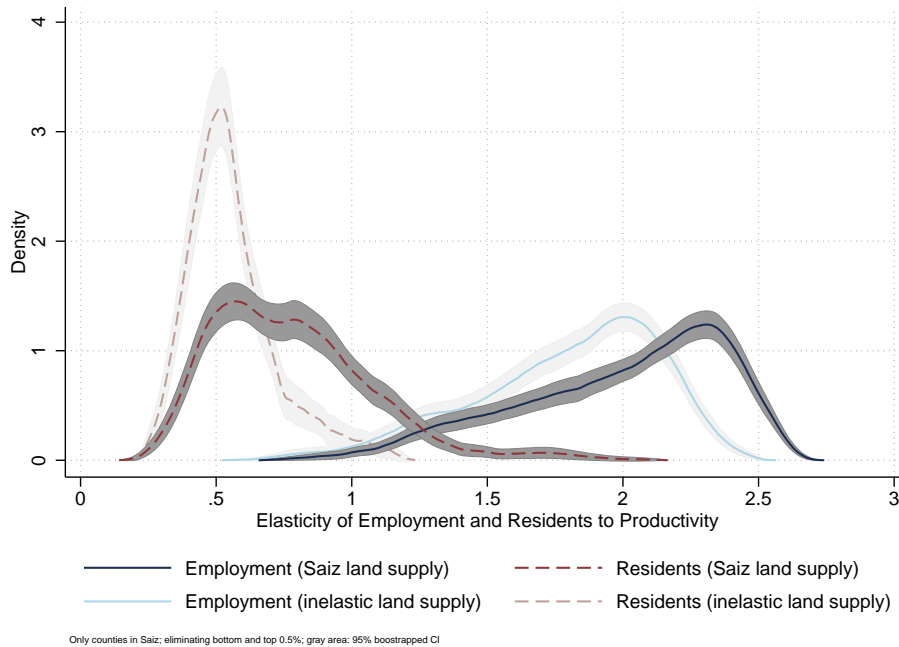
We use the empirical estimates of developed land supply elasticities from Saiz (2010), which are based on physical and regulatory constraints to the geographical expansion of developed land area. Physical constraints are measured using Geographical Information Systems (GIS) data on the location of bodies of water (oceans and lakes) and wetlands and the elevation of terrain (the fraction of surrounding land that has a slope above 15 percent). Regulatory constraints are measured using the Wharton Residential Urban Land Regulation Index, which captures the stringency of residential growth controls. Using these data, Saiz (2010) estimates developed land supply elasticities for 95 Metropolitan Statistical Areas (MSAs) in the United States. The population-weighted average of these land supply elasticities is 1.75, and they range from 0.76 for the 10th ranked MSA (San Jose, CA) to 3.09 for the 85th ranked MSA (Charlotte-Gastonia-Rock Hill, NC-SC).

To calibrate the model to the initial distribution of economic activity, we need to assume a housing supply elasticity for each county. Therefore, we need to decide what to assume about counties that are not part of MSAs and how to treat counties within multi-county MSAs. The Saiz estimates are based on the expansion of the geographical frontier of developed land around the boundaries of the MSA. Counties that are not part of MSAs can expand this geographical frontier, and hence we assume the maximum housing supply elasticity across MSAs for these counties. In MSAs that consist of multiple counties, central counties that are surrounded by other already-developed counties cannot expand this geographical frontier. Hence, we assume a housing supply elasticity of zero for these central counties. Single-county MSAs and outlying counties in multi-county MSAs have the estimated Saiz elasticity.³⁰

When we undertake counterfactuals for productivity shocks, we focus on the subsample of counties within MSAs for which Saiz housing supply elasticities are available and no imputation is required. We shock each county in this subsample with a 5 percent productivity shock (holding productivity in all other counties and holding all other exogenous variables constant). Figure 3 displays the results across this subsample of counties. The dark solid blue line shows the estimated kernel density for the general equilibrium employment elasticity using the Saiz housing supply elasticities. The light solid blue line shows the analogous kernel density for the employment elasticity for this same subsample of counties assuming an inelastic housing supply. The dark dashed red line shows the estimated kernel density for the general equilibrium resident elasticity using the Saiz housing supply elasticities. The light dashed red line shows the analogous kernel density for the resident elasticity for this same subsample of counties assuming an inelastic housing supply. We also show the 95 percent confidence intervals around these estimated kernel

³⁰In robustness checks, we considered variations in these assumptions, such as assuming the minimum housing supply elasticity across MSAs for central counties. Again, we find the same pattern of heterogeneous local employment elasticities.

densities (gray shading).



Note: Counterfactuals for the subsample of counties in the 95 MSAs for which Saiz (2010) reports a housing supply elasticity.

Figure 3: Kernel density for the distribution of employment and residents elasticities in response to a productivity shock across counties (Positive housing supply elasticity)

Comparing these two sets of results, we find that introducing differences in housing supply elasticities across locations shifts both distributions to the right. The reason is that the productivity shock now induces an increase in the supply of housing, which implies a smaller increase in land prices and wages, and hence a larger increase in both residents and employment in response to the productivity shock. We also find that introducing differences in housing supply elasticities increases the heterogeneity in the elasticity of residents with respect to labor demand shocks, but has relatively little impact on the heterogeneity in the elasticity of employment. This pattern of results is also intuitive. The housing supply elasticity matters less for employment than for residents, because commuting allows individuals to work in locations with inelastic housing supplies without actually having to live there and pay the resulting high land prices. Therefore, a productivity shock in a location with an inelastic housing supply can increase employment through commuting without requiring substantial changes in the number of residents. An important implication is that improvements in commuting technologies provide an alternative approach to relaxing housing supply elasticities in enabling individuals to access high productivity locations.

To provide further evidence on the respective roles of commuting and housing supply elasticities, we regress the general equilibrium employment elasticity for this subsample of counties on both determinants and the other controls from Table 2 above (see Section C.6 of the web appendix). Consistent with the pattern in Figure 3, we find that the magnitude and statistical significance of the estimated coefficients on our commuting measures are unaffected by the inclusion of the Saiz housing supply elasticity.

4.3 Measuring the Incidence of Local Labor Demand Shocks

We now relate our model’s prediction of heterogeneous local employment elasticities to the large empirical literature on the incidence of local labor demand shocks. In this section, we illustrate these implications using the counterfactuals from our theoretical model. In the next section, we provide separate evidence, independent of the model, in support of these predictions.

In particular, we compare the general equilibrium elasticities of employment with respect to the productivity shock in the model to reduced-form “differences-in-differences” estimates of the local average treatment effects of the productivity shock. We construct a regression sample including both treated and untreated counties from our 3,111 counterfactuals in which we shock each county in turn with a 5 percent positive productivity shock ($3,111^2 = 9,678,321$ observations). For each of these separate exercises, we estimate a “differences-in-differences” specification given by:

$$\Delta \ln Y_{it} = a_0 + a_1 \mathbb{I}_{it} + a_2 X_{it} + a_3 (\mathbb{I}_{it} \times X_{it}) + u_{it}, \quad (20)$$

where i denotes the 3,111 counties and t indexes the 3,111 counterfactuals; $\Delta \ln Y_{it}$ is the change in log employment between the counterfactual and actual equilibria; \mathbb{I}_{it} is a (0,1) indicator for whether a county is treated with a productivity shock; and X_{it} are controls. We again consider two sets of controls (X_{it}): the model-suggested measures of linkages in goods and factor markets and more standard econometric controls (log employment, log wages and land area). We include both the main effects of these controls (captured by a_2) and their interactions with the treatment indicator to capture heterogeneity in the treatment effects (captured by a_3).

In a specification without the controls ($a_2 = a_3 = 0$), the average effect of the productivity shock on the untreated counties is captured in the regression constant (a_0), and the local average treatment effect (a_1) corresponds to the difference in means between the treated and untreated counties. In specifications with controls, a_2 captures the main effect of these controls, and a_3 allows the response of employment to the productivity shock to depend on these controls. We compare estimating this regression specification including (i) a random untreated county in the control group, (ii) only the nearest untreated county in the control group, (iii) only neighboring counties within 120 kilometers of the treated county in the control group, (iv) only non-neighboring counties located from 120-240 kilometers from the treated county, and (v) all untreated counties in the control group.

In Section C.4 of the web appendix, we show that the model-suggested controls are more successful in explaining the heterogeneity in treatment effects than standard empirical controls from the local labor markets literature. For each of the different control groups (i)-(v), we compute the deviation between the general equilibrium elasticity from the model and the predicted employment elasticity from the “differences-in-differences” regression. None of the “differences-in-differences” specifications completely captures the general equilibrium employment elasticity. However, taking into account commuting linkages with the model-suggested controls substantially increases the predictive power of the “differences-in-differences” specification.

In general, we find similar patterns of results across the different control groups. The one exception is the specification using the nearest county as a control, because employment in the nearest untreated

county is typically negatively affected by the increase in productivity in the treated county. While the use of contiguous locations to take differences is often motivated based on similar unobservables (as in regression discontinuity designs), this pattern of results highlights that contiguous locations are also likely to be the most severely affected by spatial equilibrium linkages in goods and factor markets.

In Subsection C.5 of the web appendix, we show that we find a similar pattern of results if we use spatially correlated shocks reproducing the industrial composition of the U.S. economy, and in Subsection C.11 of the web appendix, we show that we obtain the same pattern of findings if we replicate our entire quantitative analysis using CZs rather than counties. Therefore, while capturing the full general equilibrium effects of the productivity shocks requires solving the model-based counterfactuals, we find that augmenting “difference-in-difference” regressions with measures of commuting linkages captures in a reduced-form way the heterogeneity in the estimated treatment effects.

5 Independent Empirical Evidence

We now provide three separate pieces of empirical evidence independent of our model for the importance of commuting for employment changes using shift-share decompositions, the location of million dollar plants as in GHM, and international trade shocks as in ADH. Appendix D.2 contains details on the data used.

5.1 Shift-Share Decompositions

We first use shift-share decompositions that impose minimal structure on the data. We undertake these decompositions for both cross-section and time-series variation in employment. In particular, we use the accounting identity of the commuter market clearing condition, which requires that employment in each county equals the sum of commuting flows from itself and from other counties:

$$L_{it} = \lambda_{ii|it}^R R_{it} + \sum_{n \neq i} \lambda_{ni|nt}^R R_{nt}. \quad (21)$$

Taking differences between each county i and the county m with the median level of employment, we can decompose cross-section variation in employment between counties into the contributions of differences in the following four terms: (i) the number of residents holding own commuting shares constant; (ii) own commuting shares holding own residents constant; (iii) other residents holding other commuting shares constant; and (iv) other commuting shares holding other residents constant:

$$\Delta^I L_{it} = \underbrace{\lambda_{ii|it}^R \Delta^I R_{it}}_{\text{(i) own residents}} + \underbrace{R_{mt} \Delta^I \lambda_{ii|it}^R}_{\text{(ii) own commuting shares}} + \underbrace{\left(\sum_{n \neq i} \lambda_{ni|nt}^R R_{nt} - \sum_{n \neq m} \lambda_{ni|nt}^R R_{nt} \right)}_{\text{(iii) other residents}} + \underbrace{\sum_{n \neq m} \left(\lambda_{ni|nt}^R - \lambda_{nm|nt}^R \right) R_{nt}}_{\text{(iv) other commuting shares}}, \quad (22)$$

where Δ^I is the cross-section difference operator between an individual county i and the county m with median level of employment (such that $\Delta^I L_{it} = L_{it} - L_{mt}$).³¹

³¹In the other residents term (iii), the sets of other counties $n \neq i$ and $n \neq m$ differ from one another, while the other commuting shares are held fixed at $\lambda_{ni|nt}^R$. In contrast, in the other commuting term (iv), the set of other counties is held

In the special case of no commuting between counties, the final three terms are all necessarily equal to zero, because $\lambda_{ii|it}^R = 1$ and $\lambda_{ni|nt}^R = \lambda_{nm|nt}^R = 0$ for $n \neq i$ and $n \neq m$ in this special case. Therefore the sum of the final three terms relative to the first term reveals the overall importance of commuting for cross-section differences in employment across counties. The values of the final three terms relative to one another inform about the different ways in which commuting affects employment differences between counties, either through variation in own commuting shares, the set of residents in other counties, or variation in commuting shares with other counties.

In Table 3, we report the results of this cross-section decomposition (22) using our bilateral commuting data for 2006-10. As all four terms are equal to zero for the county with median employment, we report the distribution of results across all other counties. Each individual term can be either positive or negative. Therefore, we assess the relative importance of each term by expressing its absolute value as a percentage of the sum of the absolute values of all four terms. We report the means of these percentage contributions, which necessarily sum to one hundred, because the mean is a linear operator, and the four percentage contributions sum to one hundred for each individual county. We also report percentiles of the distribution of each contribution, which need not sum to one hundred, because the county at a given percentile for one contribution may be different from the county at the same percentile for another contribution.

As apparent from the table, we find that the three commuting contributions are large relative to the own residents contribution. On average, 29.3 percent of the differences in employment size are accounted for by differences in the number of residents (bottom row, second column), with the remaining 69.7 percent attributed to the three commuting terms. Of this remainder, differences in commuting shares with other counties contribute 28.7 percent (bottom row, fifth column), differences in the sets of residents in other counties make up 31.3 percent (bottom row, fourth column) and differences in own commuting shares constitute 10.6 percent (bottom row, second column). We also find substantial heterogeneity across counties in the relative importance of the three commuting contributions, which range from 91.9 percent (100 minus the own residents contribution of 8.1 in the second row, second column) to 41.4 percent (100 minus the own residents contribution of 58.6 in the sixth row, second column).

2006-10	(i) Changes Own Residents, Constant Own Commuting	(ii) Changes Own Commuting, Constant Own Residents	(iii) Changes Other Residents, Constant Other Commuting	(iv) Changes Other Commuting, Constant Other Residents	Sum (i)-(iv)
10th percentile	8.1	0.7	13.3	6.1	-
25th percentile	15.5	2.6	26.7	13.9	-
50th percentile	25.5	8.3	33.4	35.0	-
75th percentile	43.7	17.8	38.7	40.4	-
90th percentile	58.6	24.4	43.1	44.1	-
Mean	29.3	10.6	31.3	28.7	100

Mean and percentiles of the distribution of the percentage contributions to cross-section differences in employment between each county and the median county for 2006-10. The four terms are differences in (i) the number of residents holding own commuting shares constant; (ii) own commuting shares holding own residents constant; (iii) other residents holding other commuting shares constant; and (iv) other commuting shares holding other residents constant.

Table 3: Cross-section Decomposition of Employment Differences across Counties for 2006-10

fixed for $n \neq m$, and the other commuting shares vary for $\lambda_{ni|nt}^R \neq \lambda_{nm|nt}^R$. See Section C.1 of the web appendix for a more detailed derivation. This decomposition also can be undertaken using other counties (instead of the median county) as the base. We find the same pattern across alternative choices for the base.

Taking differences over time within an individual county i in equation (21), employment changes over time can be decomposed into four analogous terms, where we replace the cross-section difference operator (Δ^I) in equation (22) by the time-series difference operator ($\Delta^T L_{it} = L_{it} - L_{it-1}$):

$$\Delta^T L_{it} = \underbrace{\lambda_{ii|it}^R \Delta^T R_{it}}_{\text{(i) own residents}} + \underbrace{R_{it-1} \Delta^T \lambda_{ii|it}^R}_{\text{(ii) own commuting shares}} + \underbrace{\sum_{n \neq i} \lambda_{ni|nt}^R \Delta^T R_{nt}}_{\text{(iii) other residents}} + \underbrace{\sum_{n \neq i} R_{nt-1} \Delta^T \lambda_{ni|nt}^R}_{\text{(iv) other commuting shares}}. \quad (23)$$

In the special case of no commuting between counties, the first term for changes in own residents ($\lambda_{ii|it}^R \Delta^T R_{it}$) is the only source of employment changes, because $\lambda_{ii|it}^R = 1$, $\lambda_{ni|nt}^R = 0$ for $n \neq i$, and $\Delta^T \lambda_{ni|nt}^R = 0$ for all n, i in this special case, which implies that the final three terms are all equal to zero. Therefore, comparing the sum of the final three terms for own commuting shares, other residents and other commuting shares to the first term for own residents again reveals the relative importance of commuting for employment variation.

In Table 4, we report the results of this time-series decomposition using the change in the bilateral commuting probabilities between 1990 and 2006-2010. Again the individual contributions in the decomposition (22) can be either positive or negative. Therefore we evaluate the relative importance of each contribution by expressing its absolute value as a percentage of the sum of the absolute values of all four contributions. These percentage contributions have the same properties as for our cross-section decomposition above.

1990 to 2006-10	(i) Changes Own Residents, Constant Own Commuting	(ii) Changes Own Commuting, Constant Own Residents	(iii) Changes Other Residents, Constant Other Commuting	(iv) Changes Other Commuting, Constant Other Residents	Sum (i)-(iv)
10th percentile	17.2	2.2	3.1	1.8	-
25th percentile	33.7	6.3	7.3	5.9	-
50th percentile	50.0	16.0	12.3	13.1	-
75th percentile	66.0	30.4	18.6	22.8	-
90th percentile	80.2	46.3	26.8	34.1	-
Mean	49.7	20.4	14.0	15.8	100

Mean and percentiles of the distribution of the contributions to time-series changes in employment between 1990 and 2006-10 from (i) the number of residents holding own commuting shares constant; (ii) own commuting shares holding own residents constant; (iii) other residents holding other commuting shares constant; and (iv) other commuting shares holding other residents constant.

Table 4: Time-series Decomposition of County Employment Changes between 1990 and 2006-10

As shown in the table, we again find that commuting makes a large and heterogeneous contribution towards the observed employment changes. On average, 49.7 percent of the differences in employment size are accounted for by differences in the number of residents (bottom row, second column), with the remaining 50.3 percent attributed to the three commuting terms. Of this remainder, differences in commuting shares with other counties contribute 15.8 percent (bottom row, fifth column), differences in the number of residents in surrounding counties make up 14.0 percent (bottom row, fourth column) and differences in own commuting shares constitute 20.4 percent (bottom row, second column). The relative importance of the three commuting contributions ranges from 82.8 percent (100 minus the own residents contribution of 17.2 from the second row, second column) to 19.8 percent (100 minus the own residents contribution of 80.2 from the sixth row, second column).

In both the cross-section and over time, variation in county employment is ultimately driven by variation in productivity and other county characteristics. Therefore the findings of these cross-section and time-series decompositions are consistent with the predictions of our model that the response of employment to these underlying shocks is shaped by large and heterogeneous responses of commuting flows.

5.2 Million Dollar Plants Natural Experiment

In this subsection, we follow GHM in using the location decisions of million dollar plants (MDP) as a source of local labor demand shocks.³² A key challenge in evaluating the effects of such local labor demand shocks is constructing a counterfactual for what would have happened in the absence of the shock. To address this challenge, GHM use the revealed rankings of profit-maximizing firms. These rankings come from the corporate real estate journal *Site Selection*, which includes a regular feature titled the “Million Dollar Plants” that describes how a large plant decided where to locate. The “Million Dollar Plants” articles report the county that the plant ultimately chose (i.e., the winner), as well as the one or two runner-up counties (i.e., the losers). The losers are counties that have survived a long selection process, but narrowly lost the competition. Therefore, the identifying assumption is that the losers form a valid counterfactual for the winners, conditional on the controls.

In an initial study using county-level data, Greenstone and Moretti (2004) considered a list of 82 MDPs from *Site Selection*. In a subsequent study of the impact of MDPs on other plants, GHM focused on a subset of 47 plants that could be found in confidential Census micro data. In our analysis of the local employment response to labor demand shocks, we follow Greenstone and Moretti (2004) in using county-level data and the full list of 82 MDPs. For each of these 82 groups of winner and runner-up counties (referred to as “cases”), we construct county-level data on employment and other characteristics from 10 years before the MDP announcement to 10 years afterwards. The resulting sample includes 166 winner and runner-up counties from 39 states from 1972 to 2003. To the extent that some of these announced plants need not have actually opened, this will attenuate the estimates, making the true effects even larger. Although the size of the labor demand shock can differ across cases, because the MDPs are not all exactly the same size, the very fact that these plants appeared in *Site Selection* as MDPs necessarily implies that they are all large plants.

In Table C.3 in the web appendix, we compare the observed characteristics of winner and runner-up counties before a MDP announcement for the full set of 82 MDPs. We find that winner counties have somewhat lower prior values of levels of employment, wages, population and population density than runner-up counties. We also find that they have somewhat more open local labor markets in terms of workplace and residence own commuting shares. Despite these differences in individual observed characteristics, the fact that the firms selected these counties as winners and runners-up suggests that they have similar implied profitability for plant location. As a check on this identifying assumption that the losers form a valid counterfactual for the winners, we now use an event-study specification following GHM and Greenstone and Moretti (2004), which allows us to trace in a flexible way the evolution of relative employment levels in winning counties relative to runner-up counties in the years leading up to and following the MDP

³²See Section D.2 of the web appendix for further discussion of the data sources for this section.

announcement.³³ We estimate

$$\ln L_{it} = \kappa \mathbb{I}_{j\tau} + \sum_{\tau=-10}^{\tau=10} \theta_{\tau} (\mathbb{T}_{\tau} \times \mathbb{W}_i) + \alpha_i + \eta_j + \mu_t + \epsilon_{it}, \quad (24)$$

where we use notation as close as possible to GHM; i indexes counties; j denotes cases (groups of winner-runner-up counties); t corresponds to calendar year; τ is a treatment year index, which is equal to the calendar year minus the MDP announcement year (the treatment year); L_{it} is county employment; $\mathbb{I}_{j\tau}$ is an indicator variable for the treatment, which equals one for a case j from the treatment year onwards and zero otherwise; \mathbb{T}_{τ} is an indicator variable for years relative to the treatment year, which equals one for year τ and zero otherwise; \mathbb{W}_i is an indicator variable for the winner county, which equals one for a winner county and zero otherwise; α_i are county fixed effects; η_j are case fixed effects; μ_t are calendar year fixed effects; and ϵ_{it} is the error term.³⁴ We report standard errors clustered by state, which allows for serial correlation in the error term within counties over time and spatial correlation in the error term across counties within states.

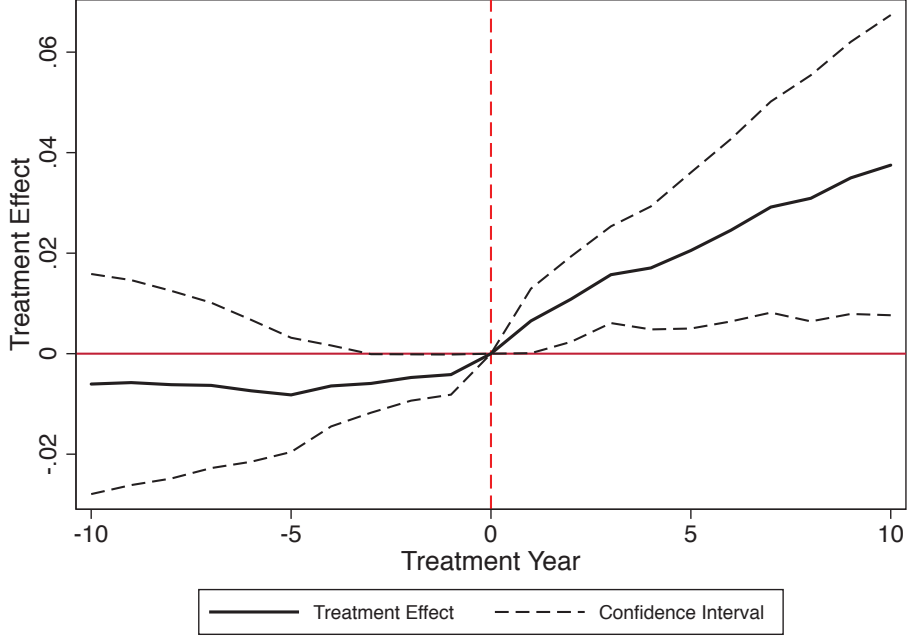
In this event-study specification, the county fixed effects control for unobserved heterogeneity across counties; the case fixed effects control for unobserved determinants of employment that affect both winning and runner-up counties during the ten-year period around each MDP announcement; and the calendar year fixed effects control for secular trends in employment over time. The coefficient κ captures the average change in employment in both winner and runner-up counties following the announcement of a MDP. This coefficient is separately identified from the calendar year fixed effects, because the treatment year occurs in different calendar years for different cases. The key coefficients of interest are θ_{τ} , which capture a “difference-in-difference”: the difference in log employment for a winner county between the treatment year ($\tau = 0$) and another year compared to the same difference for a runner-up county.

In Figure 4, we display the estimates of θ_{τ} from equation (24) and their 95 percent confidence intervals. Prior to the MDP announcement, we find no evidence of statistically significant differences in employment between winner and runner-up counties. In the treatment year in which a MDP is announced ($\tau = 0$), we find an increase in employment in treatment counties relative to control counties, which becomes statistically significant at the 5 percent level by the first year after the announcement ($\tau = 1$), and continues to increase in magnitude from there onwards. By ten years after the MDP announcement, employment in winner counties is on average close to 4 percent larger than in runner-up counties. This pattern of results is consistent with the MDP announcement gradually leading to an expansion in employment in winner counties relative to runner-up counties, through the direct effects of the construction, operation and expansion of the MDP, and through the indirect effects of the construction, operation and expansion of suppliers and ancillary services. Therefore, the timing of the positive estimated treatment effect in this event study provides support for the MDP identifying assumptions.

Having validated the MDP experiment, we now use it to provide evidence on the key prediction of our

³³This event-study specification for county employment is analogous to that for incumbent plant productivity reported in the lower panel of Figure 1 in GHM.

³⁴While GHM includes industry-year effects in its plant-level specifications, we include year fixed effects in our baseline county-level specifications, because unlike plants counties do not have a straightforward allocation to industries. We report robustness tests below using industry-year fixed effects based on assigning an industry for each case, census-region-year fixed effects and state-year fixed effects.



Note: Event-study for 82 MDP announcements. Specification from equation (24) includes county, case and year fixed effects, a post-MDP announcement dummy, and interaction terms between the dummy for winner county and treatment year. Standard errors are clustered by state.

Figure 4: Treatment Effect for MDP Announcement

model that the treatment effect of the MDP should be heterogeneous across cases depending on openness to commuting. In particular, we generalize the baseline “differences-in-differences” specification in GHM³⁵ to include an interaction term between the treatment effect for a MDP and the openness of the local labor market to commuting:

$$\ln L_{it} = \kappa \mathbb{I}_{j\tau} + \theta (\mathbb{I}_{j\tau} \times \mathbb{W}_i) + \beta (\mathbb{I}_{j\tau} \times \lambda_{ii|i}^R) + \gamma (\mathbb{I}_{j\tau} \times \mathbb{W}_i \times \lambda_{ii|i}^R) + \alpha_i + \eta_j + \mu_t + \epsilon_{it}, \quad (25)$$

where $\lambda_{ii|i}^R$ is the residence own commuting share in 1990; the other variables are defined as in equation (24); the main effect of $\lambda_{ii|i}^R$ is captured in the county fixed effects (α_i); β allows the average change in employment in both winner and runner-up counties following the announcement of a MDP to depend on the residence own commuting share; θ captures the mean increase in employment in the winner county relative to the runner-up counties following a MDP announcement for a completely open local labor market ($\lambda_{ii|i}^R = 0$); and γ allows the mean increase in employment in the winner county relative to the runner-up counties following a MDP announcement to depend on the residence own commuting share. If we impose $\beta = \gamma = 0$ in equation (25), we obtain the baseline “differences-in-differences specification” in GHM.

Table 5 reports the estimation results, where we weight county observations by population at the beginning of the sample. In Column (1), we impose $\beta = \gamma = 0$, and estimate that a MDP increases county employment by 5.7 percent, which is statistically significant at the 1 percent level. The magnitude of this estimated treatment is broadly in line with the county-level results reported in Table 9 of GHM. Our estimate for employment is somewhat larger than their estimates for wages (2.7 percent) but somewhat

³⁵Equation (7) in GHM with $\psi = \Omega = \gamma = \theta_2 = 0$.

smaller than their estimates for the number of manufacturing plants (12.6 percent) and manufacturing output (14.5 percent). Our employment data are for all sectors (including non-manufacturing), whereas most MDPs are in manufacturing, which plausibly explains why our employment estimate is smaller than the estimates using manufacturing sector outcomes in GHM.³⁶

In Column (2), we augment Column (1) with our interaction terms for the residence own commuting share ($\lambda_{ii|i}^R$). Consistent with the predictions of our theoretical model, we find a negative estimated coefficient on the commuting interaction term ($\gamma < 0$), implying greater increases in employment in response to the positive labor demand shock in counties with more open local labor markets (lower $\lambda_{ii|i}^R$). In contrast, we find no effect of residence own commuting shares in runner-up counties (we are unable to reject the null hypothesis that β is equal to zero), which is consistent with runner-up counties not experiencing a labor demand shock from the opening of a MDP. In Columns (3)-(5), we show that this finding of a greater employment response in more open commuting markets is robust across our different commuting measures, including workplace (Column (3)), the average of workplace and residence (Column (4)), and the minimum of workplace and residence (Column (5)). These estimates are not only statistically significant but also economically relevant. Using our baseline estimates from Column (2), the standardized coefficient for our interaction term (γ) with the residence own commuting share is -0.050, which is close in absolute value to the standardized coefficient for the main effect (β) of 0.067.³⁷

In Columns (6)-(7), we demonstrate the robustness of our results across alternative possible specifications. In Column (6), we address potential concerns about heterogeneous industry trends by replacing the year fixed effects with industry-year fixed effects.³⁸ In Columns (6) and (7), we examine potential concerns of heterogeneous trends across different geographical regions by replacing the year fixed effects with census-region-year fixed effects and state-year fixed effects respectively. Across all three specifications, we continue to find larger employment responses to the labor demand shock in more open commuting markets, with estimated coefficients around the same magnitude as in our baseline specification in Column (2). As a final check for pre-trends for the commuting interaction, Section C.8 of the web appendix uses a generalization of the event-study specification in equation (24) to estimate a coefficient on this commuting interaction for each year. As shown in Figure C.9, we find that this estimated coefficient is flat and not statistically significantly different from zero before the MDP announcement, but turns sharply negative and becomes statistically significant following the MDP announcement. Therefore, this time pattern of the estimated coefficients for the commuting interaction term provides further support for our interpretation that commuting openness shapes the employment response to the local labor demand shock from the MDP.

Although our model points to commuting linkages as the source of heterogeneity in the treatment effect of a MDP, another potential concern is that the treatment effect could vary with other location characteristics. To address this concern, we now consider a non-parametric specification that estimates a

³⁶We find a smaller estimated treatment effect using population rather than employment, with an estimated coefficient (standard error) of 0.050 (0.018).

³⁷We compute these standardized coefficients as $\gamma \times \text{s.d.}(\mathbb{I}_{j\tau} \times \mathbb{W}_i \times \lambda_{ii|i}^R) / \text{s.d.}(\ln L_{it})$ and $\theta \times \text{s.d.}(\mathbb{I}_{j\tau} \times \mathbb{W}_i) / \text{s.d.}(\ln L_{it})$, respectively, where s.d. denotes the standard deviation.

³⁸We use the reported industry for each case from Appendix Table 2 in Greenstone and Moretti (2004).

Variable	Coefficient	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\mathbb{I}_{j\tau} \times \mathbb{W}_i$	θ	0.057** (0.018)	0.250*** (0.078)	0.191*** (0.065)	0.244*** (0.068)	0.260*** (0.078)	0.223*** (0.078)	0.160** (0.060)	0.159** (0.066)
$\mathbb{I}_{j\tau} \times \mathbb{W}_i \times \lambda_{ii i}^R$	γ		-0.242** (0.096)				-0.219** (0.096)	-0.190** (0.077)	-0.195** (0.066)
$\mathbb{I}_{j\tau} \times \mathbb{W}_i \times \lambda_{ii i}^L$	γ			-0.177** (0.087)					
$\mathbb{I}_{j\tau} \times \mathbb{W}_i \times \lambda_{ii i}^{ARL}$	γ				-0.241*** (0.088)				
$\mathbb{I}_{j\tau} \times \mathbb{W}_i \times \lambda_{ii i}^{MRL}$	γ					-0.281** (0.110)			
$\mathbb{I}_{j\tau} \times \lambda_{ii i}^R$	β		0.012 (0.135)				-0.048 (0.108)	-0.203*** (0.075)	-0.213** (0.082)
$\mathbb{I}_{j\tau} \times \lambda_{ii i}^L$	β			0.243* (0.129)					
$\mathbb{I}_{j\tau} \times \lambda_{ii i}^{ARL}$	β				0.124 (0.160)				
$\mathbb{I}_{j\tau} \times \lambda_{ii i}^{MRL}$	β					0.133 (0.145)			
$\mathbb{I}_{j\tau}$	κ	-0.015* (0.008)	-0.024 (0.109)	-0.200** (0.096)	-0.113 (0.125)	-0.113 (0.106)	0.021 (0.086)	0.160** (0.060)	0.159** (0.066)
County Fixed Effects		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Case Fixed Effects		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects		Yes	Yes	Yes	Yes	Yes			
Industry-year Fixed Effects							Yes		
Census-region-year Fixed Effects								Yes	
State-year Fixed Effects									Yes
Observations		4,431	4,431	4,431	4,431	4,431	4,431	4,431	4,431
R-squared		0.991	0.0991	0.991	0.991	0.991	0.992	0.994	0.996

Estimation results for equation (25); $\mathbb{I}_{j\tau}$ is an indicator that equals one for a case j from the treatment year ($\tau = 0$) onwards and zero otherwise; \mathbb{W}_i is an indicator that equals one for a winner county i and zero otherwise; $\lambda_{ii|i}^R$ is the residence own commuting share; $\lambda_{ii|i}^L$ is the workplace own commuting share; $\lambda_{ii|i}^{ARL}$ is the average of the residence and workplace own commuting shares; $\lambda_{ii|i}^{MRL}$ is the minimum of the residence and workplace own commuting shares. County observations are weighted by population at the beginning of the sample period. Standard errors are clustered by state. * denotes significance at the 10 percent level; ** denotes significance at the 5 percent level; *** denotes significance at the 1 percent level.

Table 5: Estimated MDP Treatment and Commuting Openness

separate treatment effect for each case

$$\ln L_{it} = \kappa \mathbb{I}_{j\tau} + \sum_{j=1}^J \theta_j (\mathbb{I}_{j\tau} \times \mathbb{W}_i) + \alpha_i + \eta_j + \mu_t + \epsilon_{it}, \quad (26)$$

where all variables are defined in the same way as in equation (25); the key coefficient of interest is θ_j , which captures the increase in employment in the winner county in case j following the MDP announcement relative to the same increase for the runners-up.³⁹

In Figure C.8 in Section C.8 of the web appendix, we display the estimated treatment effects for each case, which range from negative values to just below one. Therefore, although the average estimated treatment effect is positive, there is substantial variation around this average. As a result, we reject the null hypothesis that these estimated treatment effects are all equal to the same common value at conventional

³⁹The treatment effect for each case is estimated relative to an excluded category, which is the runner-up counties for each case. We are therefore able to estimate a treatment effect for each case.

levels of significance (p-value 0.000). To provide further evidence on the determinants of this heterogeneity, we regress these estimated treatment effects for each case on a range of observed characteristics. We measure these observed characteristics for each case as the average of the values for the winner and runner-up counties. We weight the observations on the 82 cases in these regressions by the inverse of the standard errors with which the treatment effects are estimated. We cluster the standard errors in these regressions on the states in which treatment and control counties are located for each case.

In Table 6, we report the results from these regressions. Consistent with our earlier results in Table 5, we again find greater employment increases in response to the positive labor demand shock in local labor markets that are more open to commuting. Across Columns (1)-(4), we find negative and statistically significant coefficients for our four (inverse) measures of openness to commuting. In Columns (5)-(8), we show that we find a similar pattern of results if we also control for log population and log land area. Although in some cases, land area or population are statistically significant, these correlations are not robust across specifications. Therefore, the heterogeneity in estimated treatment effects indeed appears to be related to the openness of the local labor market to commuting rather than size or population density.

Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\lambda_{jj j}^R$	-0.666*** (0.183)				-0.875*** (0.198)			
$\lambda_{jj j}^L$		-0.456* (0.239)				-0.832** (0.358)		
$\lambda_{jj j}^{ARL}$			-0.690*** (0.226)				-1.078*** (0.279)	
$\lambda_{jj j}^{MRL}$				-0.625*** (0.222)				-1.005*** (0.271)
ln Population _j					0.045** (0.022)	0.014 (0.028)	0.028 (0.023)	0.030 (0.022)
ln Land Area _j					0.078** (0.038)	0.092 (0.061)	0.115** (0.047)	0.121** (0.047)
Observations	82	82	82	82	82	82	82	82
R-squared	0.18	0.05	0.13	0.13	0.25	0.08	0.20	0.21

Regressions of the estimated treatment effect ($\hat{\theta}_j$ from equation (26)) for each of the 82 MDP cases on case characteristics; case characteristics measured as the average of those for winner and runner-up counties for that case; $\lambda_{jj|j}^R$ is the residence own commuting share; $\lambda_{jj|j}^L$ is the workplace own commuting share; $\lambda_{jj|j}^{ARL}$ is the average of the residence and workplace own commuting shares; $\lambda_{jj|j}^{MRL}$ is the minimum of the residence and workplace own commuting shares. Observations are weighted by the inverse of the standard error of the estimated treatment effect ($\hat{\theta}_j$) for each case. Standard errors are clustered on the states in which the treatment and control counties are located for each case. * denotes significance at the 10 percent level; ** denotes significance at the 5 percent level; *** denotes significance at the 1 percent level.

Table 6: Heterogeneous Treatment Effects and Commuting Openness

5.3 China Shock

ADH find that commuting zones (CZs) more exposed to the China shock have experienced a range of adverse labor market outcomes. We now show that the labor market response to this shock varies systematically across counties with their openness to commuting.⁴⁰ For brevity, we focus on the central empirical specification from ADH (their Table 3), using the share of manufacturing employment in the working-age

⁴⁰See Section D.2 of the web appendix for further discussion of the data sources for this section.

population, but we also report results using log total employment.

We construct the U.S. China shock and its instrument in exactly the same way and for the same time period as in ADH, but using county rather than CZ data. In particular, we measure the change in Chinese import exposure per worker for each U.S. county (ΔIPW_{it}^{US}) as

$$\Delta IPW_{it}^{US} = \sum_j \frac{L_{ijt}}{L_{jt}} \frac{\Delta M_{cjt}}{L_{it}}, \quad (27)$$

where we index counties by i , sectors by j , China by c and time by t ; ΔM_{cjt} is the change in Chinese imports in sector j for the U.S. as a whole; L_{ijt}/L_{jt} is the start-of-period share of county i in U.S. employment in industry j at time t ; and L_{it} is county i 's start-of-period total employment at time t .

We instrument the U.S. China shock (ΔIPW_{it}^{US}) with an analogous measure constructed using the imports of eight other developed countries and employment lagged by a decade (ΔIPW_{it}^{Oth}).⁴¹ This instrument exploits the idea that the growth in Chinese imports is largely driven by a supply-side shock in China, which should be reflected in rising Chinese import penetration in other developed countries. We construct these two measures for each U.S. county for the time periods 1990-2000 and 2000-7.

Pooling these two cross-sections, we estimate the same ‘‘differences-in-differences’’ empirical specification as in ADH, augmented to allow the impact of the China shock to vary across counties depending on our (inverse) measure of openness to commuting

$$\Delta Y_{it} = \gamma_t + \beta_1 \Delta IPW_{it}^{US} + \beta_2 \lambda_{ii|it}^R + \beta_3 \left(\lambda_{ii|it}^R \times \Delta IPW_{it}^{US} \right) + \mathbf{X}'_{it} \beta_4 + e_{it}, \quad (28)$$

where ΔY_{it} is the percentage point change in manufacturing employment as a share of the working age population; $\lambda_{ii|it}^R$ is the residence own commuting share for the initial period; \mathbf{X}'_{it} is a matrix of controls for the initial period; and e_{it} is a stochastic error. We allow for county fixed effects in the level of the manufacturing share (Y_{it}), which are differenced out over time. The time fixed effect (γ_t) allows for secular trends in the manufacturing share over time. We weight the observations by the start-of-period population and cluster the standard errors by state, which allows for serial correlation in the error term over time within counties and spatial correlation in the error term across counties within states.

In Table 7, we report the estimation results. In Column (1), we include only the main effect of the China shock (ΔIPW_{it}^{US}), and find a negative and statistically significant impact of increased Chinese import competition on the manufacturing share. Our estimated county-level coefficient of -0.612 is broadly in line with the estimated CZ coefficient of -0.746 in Table 3 of ADH. We find that the instruments have power in the first-stage regression, with first-stage F-statistics well above the conventional threshold of ten. In Column (2), we augment this specification with the main effect and interaction term for the residence own commuting share. Both the main effects of the China shock and commuting openness are now positive but not statistically significant. In contrast, the interaction term between these two variables is negative and statistically significant. Therefore, we find that the estimated effects of the China shock are indeed heterogeneous, depending on the openness of counties to commuting. These effects are not only statistically significant but also economically large, with a one standard deviation change in the interaction

⁴¹The eight countries are Australia, Denmark, Finland, Germany, Japan, New Zealand, Spain, and Switzerland.

term reducing the manufacturing share by around 0.5 standard deviations.⁴²

Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)
ΔIPW_{it}^{US}	-0.612*** (0.118)	0.182 (0.154)	0.180 (0.179)	0.107 (0.124)	0.121 (0.113)	0.003 (0.127)	0.040 (0.235)
$(\lambda_{ii it}^R \times \Delta IPW_{it}^{US})$		-1.128*** (0.237)	-1.026*** (0.314)	-0.927*** (0.210)	-0.696*** (0.162)	-0.499*** (0.177)	-0.573* (0.320)
$\lambda_{ii it}^R$		0.115 (0.862)	-2.140** (0.886)	-4.544*** (0.842)	-6.588*** (1.020)	-5.900*** (0.906)	-7.444*** (1.215)
ln Area					0.865*** (0.216)	0.806*** (0.282)	0.883*** (0.310)
Share Man Emp					-16.745*** (2.576)	-20.461*** (2.744)	-21.423*** (2.683)
Share College Educ						-14.855*** (3.065)	-12.882*** (2.701)
Share Foreign Born						-1.145 (2.258)	-0.395 (2.533)
Share Female Employed						0.831 (1.769)	0.015 (1.805)
State Fixed Effects	No	No	Yes	No	No	No	No
Commuting Zone Fixed Effects	No	No	No	Yes	Yes	Yes	Yes
Year 2000 Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	6,214	6,214	6,214	6,214	6,214	6,214	6,214
First-stage F (ΔIPW_{it}^{US})	160.56	66.27	62.26	34.85	27.15	27.14	23.43
First-stage F ($\lambda_{ii it}^R \times \Delta IPW_{it}^{US}$)		73.08	66.39	37.54	29.69	29.74	22.06
First-stage F ($\lambda_{ii it}^R$)							72.82

Regressions of the percentage point change in manufacturing employment as a share of the working-age population; (ΔY_{it}) for 1990-2000 and 2000-7 from equation (28); ΔIPW_{it}^{US} is the county U.S. China shock from equation (27); $\lambda_{ii|it}^R$ is the county residence own commuting share at the beginning of the period; Share Man Emp is the share of manufacturing in county residence employment at the beginning of the period; Share College Educ is the share of county's residents with a college degree in the population at least 25 years old; Share Foreign Born is the share of foreign-born residents in county's population; Share Female Employed is the share of working-age women that are employed; The U.S. China shock (ΔIPW_{it}^{US}) and its interaction with the residence own commuting share ($(\lambda_{ii|it}^R \times \Delta IPW_{it}^{US})$) are instrumented with the China shock using imports for other developed countries and lagged employment (ΔIPW_{it}^{Oth}) and its interaction with the residence own commuting share ($(\lambda_{ii|it}^R \times \Delta IPW_{it}^{Oth})$). Reported standard errors are clustered by state. * denotes significance at the 10 percent level; ** denotes significance at the 5 percent level; *** denotes significance at the 1 percent level.

Table 7: Employment Growth, the China Shock and Commuting Openness

In general, our model can generate either a positive or a negative coefficient on the interaction term with commuting openness when labor demand shocks occur simultaneously and are spatially correlated across locations. In such a setting, the sign of the coefficient on the interaction term with commuting openness depends on whether the neighbors of more open counties have shocks that are either positively or negatively correlated with the shocks to the own county. Clearly, these spatial interactions cannot be fully summarized in this reduced-form regression. Nevertheless, regardless of the sign of the estimated coefficient, our structural model rationalizes the observed changes in employment, and therefore the estimated reduced-form coefficient, by determining a mapping between the ADH shock and the fundamentals of the model.⁴³

⁴²Consistent with ADH, we find smaller estimates using total employment than using the manufacturing share. Re-estimating Column (1) using the log change in total employment, we obtain a negative and statistically significant coefficient (standard error) of -0.015 (0.004). Re-estimating Column (2) using the log change in total employment, the main effect of the China shock again becomes statistically insignificant, with coefficient (standard error) of 0.003 (0.010), while the interaction term is negative and statistically significant at the 10 percent level, with coefficient (standard error) of -0.026 (0.014).

⁴³In particular, using Propositions B.1 and B.2 in the web appendix, we can solve for the unique values of unobserved changes in productivity, the residential component of amenities, and the bilateral ease of commuting that exactly rationalize

In Column (3), we augment this specification with state fixed effects, which control for heterogeneous growth in the manufacturing share across states. In Column (4), we instead include commuting zone fixed effects, which control for all observed and unobserved factors that affect the growth in the manufacturing share over time differentially across commuting zones. In both cases, we continue to find a negative and statistically significant interaction term with the residence own commuting share. Therefore, even when we use only variation across counties within commuting zones, we continue to find heterogeneity in the estimated treatment effect of the China shock depending on the openness of counties to commuting. In Column (5), we augment the specification in Column (4) with log county land area and the initial share of employment in manufacturing, which confirms that these findings are not driven by county geography or solely by variation in exposure to the China shock from the size of the manufacturing sector as a whole.

In Column (6), we include a range of additional demographic controls from ADH, but measured at the county level, including the initial share of the population with a college degree.⁴⁴ Even after controlling for these observed county characteristics, we continue to find heterogeneity in the estimated treatment effects of the China shock across open versus closed counties. In Column (7), we instrument the residence own commuting share and its interaction with the U.S. China shock with the 1960 value of the residence own commuting share and its interaction with the predicted China shock using the imports of other developed countries and lagged employment. Even when we use this persistent component of variation in the openness of counties to commuting from some thirty years before the China shock, we continue to find that the interaction between the own commuting share and the China shock is statistically significant at the 10 percent level. Across all specifications, the first-stage F-statistics remain well above ten, confirming the power of the instruments.

Taken together, the results in Table 7 paint a consistent picture of heterogeneous treatment effects across counties depending on openness to commuting.⁴⁵ More generally, if we regress the change in the manufacturing share on commuting zone fixed effects, interactions between these fixed effects and the China shock, and all the controls from Column (7), we strongly reject the null hypothesis of the same elasticity of employment across counties with respect to the China shock (p-value 0.000). Importantly, we again conclude that some of this heterogeneity can be well explained by our measures of openness to commuting.

6 Changes in Commuting Costs

Having provided independent empirical evidence that commuting linkages matter for the local impact of labor demand shocks, we now use our quantitative model to show that they also matter for the aggregate spatial distribution of economic activity and welfare. Commuting enables workers to access high productivity locations without having to pay the high cost of living in those locations. An increase in the cost of commuting restricts the opportunity set available to firms and workers and hence reduces welfare. Loca-

the observed changes in bilateral commuting probabilities, employment, residents, wages and bilateral trade flows as an equilibrium of the model.

⁴⁴The CZ-level measures of the share of employment in routine occupations and the average offshorability of occupations from ADH are captured in the CZ fixed effects.

⁴⁵We find similar results with our other commuting measures. For example, re-estimating the specification in Column (2) but using the average between the residence and workplace own commuting share we obtain an estimated coefficient (standard error) of -1.390 (0.451). If we use the minimum of both shares we obtain -1.326 (0.327).

tions that were initially net exporters of commuters become less attractive residences, while locations that were initially net importers of commuters become less attractive workplaces. As agents relocate in response to the restricted opportunity set, counties become less specialized as residential or business locations.

We begin by using the observed commuting data to back out implied values of the composite parameter capturing the ease of commuting ($\mathcal{B}_{ni} \equiv B_{ni}\kappa_{ni}^{-\epsilon}$). Following Head and Ries (2001) in the international trade literature, we use the flows of commuters between locations n and i in both directions relative to their own commuting flows. Using the commuting gravity equation (10), and taking the geometric mean of these relative flows in both directions, we obtain the following measure of the average ease of commuting between locations n and i relative to the ease of commuting to themselves:

$$\tilde{\mathcal{B}}_{ni} \equiv \left(\frac{\mathcal{B}_{ni} \mathcal{B}_{in}}{\mathcal{B}_{nn} \mathcal{B}_{ii}} \right)^{1/2} = \left(\frac{L_{ni} L_{in}}{L_{nn} L_{ii}} \right)^{1/2}. \quad (29)$$

We compute this measure for both 1990 and 2010. Between these two years, both miles of paved roads and vehicle kilometers travelled increased substantially.⁴⁶ Consistent with this, we find a substantial increase in the relative ease of commuting from 4 percent ($\tilde{\mathcal{B}}_{ni} = 0.96$) at the 25th percentile, to 12 percent ($\tilde{\mathcal{B}}_{ni} = 0.88$) at the median, and 21 percent ($\tilde{\mathcal{B}}_{ni} = 0.79$) at the 75th percentile.

We use this distribution of implied changes in the relative ease of commuting to undertake counterfactuals for empirically-realistic changes in commuting costs. We assume a common reduction or increase in the costs of commuting for all counties equal to percentiles of this distribution (e.g. we assume that all counties experience a reduction in commuting costs equal to the median value of $\tilde{\mathcal{B}}_{ni}$). Given this assumption, we use the system of equations for general equilibrium in the model to solve for the new counterfactual equilibrium after the reduction in commuting costs, as discussed in Section 2.6 above. Using the commuting probability (10), expected utility (15), the price index (8) and land market clearing (5), the change in the common level of welfare across all locations from the shock to commuting costs can be decomposed as follows:

$$\hat{U} = \left(\frac{1}{\hat{\lambda}_{ii}} \right)^{\frac{1}{\epsilon}} \left(\frac{1}{\hat{\pi}_{ii}} \right)^{\frac{\alpha}{\sigma-1}} \left(\frac{\hat{w}_i}{\hat{v}_i} \right)^{1-\alpha} \frac{\hat{L}_i^{\frac{\alpha}{\sigma-1}}}{\hat{R}_i^{1-\alpha}}, \quad (30)$$

where we have used the fact that $\{\kappa_{ii}, B_{ii}, A_i, d_{ii}\}$ are unchanged; the first term in $\hat{\lambda}_{ii}$ captures the impact through changes in openness to commuting; the second term in $\hat{\pi}_{ii}$ captures the effect through adjustments in openness to goods trade; the remaining terms capture the influence of changes in the spatial distribution of wages (\hat{w}_i), expected residential income (\hat{v}_i), employment (\hat{L}_i) and residents (\hat{R}_i).

As shown in Table 8, we find substantial effects of these empirically-relevant reductions in commuting costs on aggregate welfare. Reducing commuting costs by the median proportional change observed over our time period from 1990 to 2010 is predicted to increase welfare by around 3.3 percent (second column). In contrast, raising commuting costs by the same proportional amount decreases welfare by around 2.3 percent. As we scale up the reduction in commuting costs to the 75th percentile observed over our time period, we amplify the welfare gain to 6.9 percent (first column). As we scale down the reduction in

⁴⁶Between 1990 and 2010, kilometers of paved public roads in the United States increased by over 20 percent (from 3.6 to 4.4 million), and vehicle kilometers travelled increased by more than 38 percent (from 3,451,016 to 4,775,352 million). For further discussion of this expansion in transport use, see for example Duranton and Turner (2011).

	Decrease by p75	Decrease by p50	Decrease by p25	Increase by 1/p50
Implied $\hat{\mathcal{B}}_{ni}$	0.79	0.88	0.96	1.13
Welfare Change	6.89%	3.26%	0.89%	-2.33%

This table shows the percentage change in welfare for different counterfactual changes in commuting costs. Each column reports a different counterfactual exercise; p75, p50 and p25, respectively, are the 75th, 50th and 25th percentiles of the empirical distribution of changes in the ease of commuting $\hat{\mathcal{B}}_{ni}$ from 1990-2010. The first row reports the implied $\hat{\mathcal{B}}_{ni}$ for all counties. The second row reports the percentage change in welfare for each counterfactual.

Table 8: Welfare Impacts for different Changes in Commuting Costs

commuting costs to the 25th percentile, we diminish the welfare gain to 0.89 percent (third column). These proportional changes in welfare are large relative to standard empirical estimates of the welfare gains from opening the closed economy to international trade, which for example range from less than 1 percent for the United States to just over 10 percent for Belgium in Eaton and Kortum (2002). One caveat is that we abstract from non-tradables different from housing or agriculture, which may lead us to overestimate the welfare gain from reductions in commuting costs. Nonetheless, these results clearly highlight that commuting not only shapes the local impact of shocks but can be consequential for aggregate welfare.⁴⁷

These empirically-realistic changes in commuting costs also result in substantial changes in the spatial distribution of employment and residents across locations. In Figure C.10 in the web appendix, we show the counterfactual change in employment in each county from reducing commuting costs by the median proportional change observed over our time period from 1990 to 2010. We find a strong relationship between these counterfactual changes and the initial ratio of employment to residents (L_i/R_i). As discussed in Subsection 3.2 above and shown in Subsection B.6 of the web appendix, this initial employment to residents ratio (L_i/R_i) is itself hard to explain in terms of standard empirical controls, such as land area, size and housing supply elasticities, and hence cannot easily be proxied by these variables.

In Subsection C.11 of the web appendix, we provide further evidence that the importance of commuting is by no means restricted to large cities. We undertake counterfactuals for reductions in commuting costs for CZs rather than counties (replicating our entire quantitative analysis for CZs). We show that the counterfactual changes in CZ employment from reductions in commuting costs are well explained by measures of the extent to which the CZ uses the commuting technology in the initial equilibrium. In contrast, these counterfactual changes in CZ employment are not well explained by initial CZ employment or residents size, confirming the importance of measures of commuting linkages.

Given this importance of commuting links in shaping the distribution of economic activity across locations, it is natural to expect that these links also determine the magnitude of the impact of reductions in trade costs. In Subsection C.10 of the web appendix, we explore this interaction between trade and commuting costs. We compare the counterfactual effects of a 20 percent reduction of trade costs in the actual world with commuting to the effects in a hypothetical world without commuting. In general,

⁴⁷Smaller (larger) values for the Fréchet shape parameter (ϵ) imply more (less) heterogeneity in preferences for residence-workplace pairs, which magnifies (diminishes) the effects of changes in commuting costs on welfare. For example, in a world with a 50 percent lower (higher) value of ϵ , reducing commuting costs by the median proportional change increases welfare by 6.9 (2.1) percent, while increasing commuting costs by the same amount reduces welfare by 4.8 (1.5) percent.

reductions in trade costs lead to a more dispersed spatial distribution of economic activity in the model. But this dispersal is smaller with commuting than without commuting. As trade costs fall, commuting increases the ability of the most productive locations to serve the national market by drawing workers from a suburban hinterland, without bidding up land prices as much as would otherwise occur. These results further underscore the prominence of commuting linkages in shaping the equilibrium spatial distribution of economic activity, and the necessity of incorporating them in models of economic geography.

7 Conclusions

The economic effects of local labor demand shocks have been the subject of an extensive empirical literature, which has considered a wide range of such shocks, including industry composition, international trade, macro and financial crises, and natural resource discoveries, among others. To understand the impact of these types of shocks, we develop a quantitative spatial general equilibrium model that incorporates spatial linkages between locations in both goods markets (trade) and factor markets (commuting and migration). Although we allow for a large number of locations and a rich geography in both goods and factor markets, our framework remains amenable to both analytical and quantitative analysis. We show how the model can be quantified using available data to match the observed gravity equation relationships for trade in goods and commuting, as well as the observed cross-section distributions of employment, residents and wages. Thus, our framework provides a tractable platform for undertaking a range of counterfactuals for realistic changes in trade and commuting costs.

While previous research has often worked at relatively high levels of spatial aggregation (e.g. commuting zones) to reduce the effect of unmodeled commuting links, we explicitly model the spatial interactions between locations in goods and commuting markets, thereby providing a framework for examining the local impact of labor demand shocks at finer spatial scales within commuting zones. Commuting allows workers to separate their workplace and residence, thereby introducing a quantitatively relevant distinction between the effects of local labor demand shocks on employment and residents. We find substantial heterogeneity across both counties and CZs in the elasticity of local employment to a productivity shock, which ranges from around 0.5 to 2.5. Therefore an average local employment elasticity estimated in one context can be quite misleading if applied in another context without controlling for this heterogeneity. We show that this heterogeneity is hard to explain with standard empirical controls, such as area and size, but is well explained by measures of linkages in commuting networks. We use our model to highlight a summary statistic of the share of residents that work locally that can be included in reduced-form regressions to help to control for this heterogeneity.

We provide three separate pieces of empirical evidence independent of our model for the importance of commuting in determining the employment response to changes in the local economic environment. Using shift-share decompositions, we show that commuting accounts for around two thirds of the cross-section variation and around one half of the time-series changes in employment. Using quasi-experimental evidence from million dollar plants (MDPs), we provide direct evidence in support of our model's prediction of larger increases in employment in response to labor demand shocks in counties with more open commuting markets. International trade shocks are harder to interpret, because they affect all locations simultaneously

depending on industry composition, and the neighbors of more open counties can experience shocks that are either positively or negatively correlated with those for the own county. Nonetheless, we again show that the openness of a county to commuting helps to control for the heterogeneity in its response to the local labor market demand shock.

Finally, as well as shaping the effects of local labor demand shocks, we show that commuting also matters in the aggregate for the spatial distribution of economic activity and welfare. We use observed commuting flows between pairs of counties over time to back out the empirical distribution of implied reductions in commuting costs from 1990-2010. Reducing commuting costs for all counties by the median of this empirical distribution, we find an increase in welfare of 3.3 percent, and employment changes across counties that range from increases of 28 percent to reductions of 19 percent. Taken together, our findings are consistent with the view that the openness of labor markets to commuting is central to shaping both the local response to shocks and the aggregate spatial distribution of economic activity across locations.

References

- [1] Ahlfeldt, Gabriel, Stephen Redding, Daniel Sturm and Nikolaus Wolf (2015) “The Economics of Density: Evidence from the Berlin Wall,” *Econometrica*, 83(6), 2127-2189.
- [2] Allen, Treb, and Costas Arkolakis (2014) “Trade and the Topography of the Spatial Economy,” *Quarterly Journal of Economics*, 129(3), 1085-1140.
- [3] Allen, Treb, Costas Arkolakis and Yuta Takahashi (2014) “Universal Gravity,” *NBER Working Paper*, 20787.
- [4] Allen, Treb, Costas Arkolakis and Xiangliang Li (2015) “Optimal City Structure,” Yale University, mimeograph.
- [5] Allen, Treb, Costas Arkolakis and Xiangliang Li (2016) “On the Existence and Uniqueness of Trade Equilibria,” Yale University, mimeograph.
- [6] Alonso, William (1964) *Location and Land Use*, Cambridge MA: Harvard.
- [7] Arkolakis, Costas, Arnaud Costinot and Andres Rodriguez-Clare (2012) “New Trade Models, Same Old Gains,” *American Economic Review*, 102(1), 94-130.
- [8] Armington, Paul S. (1969) “A Theory of Demand for Products Distinguished by Place of Production,” *IMF Staff Papers*, 16(1), 159-178.
- [9] Artuc, E., S. Chaudhuri, and J. McLaren (2010) “Trade Shocks and Labor Adjustment: A Structural Empirical Approach,” *American Economic Review*, 100(3), 1008-1045.
- [10] Autor David, David Dorn and Gordon Hanson (2013) “The China Syndrome: Local Labor Market Effects of Import Competition in the United States,” *American Economic Review*, 103(6), 2121-2168.
- [11] Autor David, David Dorn and Gordon Hanson (2013) “The China Syndrome: Local Labor Market Effects of Import Competition in the United States,” *American Economic Review*, 103(6), 2121-2168.
- [12] Autor David, David Dorn and Gordon Hanson (2016) “The China Shock: Learning from Labor-Market Adjustment to Large Changes in Trade,” *Annual Review of Economics*, 8, 205-40.
- [13] Autor David, David Dorn, Gordon Hanson and Jae Song (2014) “Trade Adjustment: Worker Level Evidence,” *Quarterly Journal of Economics*, 129(4), 1799-860.

- [14] Bartik, Timothy J. (1991) "Who Benefits from State and Local Economic Development Policies?" W.E. Upjohn Institute for Employment Research, Upjohn Press.
- [15] Behrens, K., G. Mion, Y. Murata and J. Südekum, 2014. "Spatial frictions," DICE Discussion Papers 160, Heinrich-Heine-Universität Düsseldorf, Düsseldorf Institute for Competition Economics (DICE).
- [16] Blanchard, Olivier Jean, and Lawrence F. Katz (1992) "Regional Evolutions," *Brookings Papers on Economic Activity*, 23(1), 1-76.
- [17] Bound, John and Holzer, Harry J (2000) "Demand Shifts, Population Adjustments, and Labor Market Outcomes during the 1980s," *Journal of Labor Economics*, 18(1), 20-54.
- [18] Broda, Christian and David E. Weinstein (2006) "Globalization and the Gains from Variety," *Quarterly Journal of Economics*, 121(2), 541-86.
- [19] Busso, M., J. Gregory, and P. Kline (2013) "Assessing the Incidence and Efficiency of a Prominent Place Based Policy," *American Economic Review*, 103(2), 897-947.
- [20] Caliendo, L. and F. Parro (2015) "Estimates of the Trade and Welfare Effects of NAFTA," *Review of Economic Studies*, 82(1), 1-44.
- [21] Caliendo, L., F. Parro, E. Rossi-Hansberg, and P.-D. Sarte (2014) "The Impact of Regional and Sectoral Productivity Changes on the U.S. Economy," *NBER Working Paper*, 20168.
- [22] Dekle, R., J. Eaton, and S. Kortum (2007) "Unbalanced Trade," *American Economic Review*, 97(2), 351-355.
- [23] Desmet, Klaus and Esteban Rossi-Hansberg (2013) "Urban Accounting and Welfare," *American Economic Review*, 103:6, 2296-2327.
- [24] Desmet, Klaus and Esteban Rossi-Hansberg (2014) "Spatial Development," *American Economic Review*, 104(4), 1211-1243.
- [25] Diamond, Rebecca (2016) "The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000," *American Economic Review*, 106(3), 479-524.
- [26] Dingel, Jonathan (2015) "The Determinants of Quality Specialization," Chicago Booth School of Business, mimeograph.
- [27] Duranton, Gilles and Matthew A. Turner (2011) "The Fundamental Law of Road Congestion: Evidence from U.S. Cities," *American Economic Review*, 101(6), 2616-2652.
- [28] Duranton, Gilles, Peter Morrow and Matthew A. Turner (2014) "Roads and Trade: Evidence from the U.S.," *Review of Economic Studies*, 81(2), 681-724.
- [29] Eaton, Jonathan and Kortum, Samuel (2002) "Technology, Geography, and Trade," *Econometrica*, 70(5), 1741-1779.
- [30] Feyrer, James, Erin T. Mansur and Bruce Sacerdote (2015) "Geographic dispersion of economic shocks: Evidence from the fracking revolution" *NBER Working Paper*, 21624.
- [31] Fujita, Masahisa, Paul Krugman and Anthony J. Venables (1999) *The Spatial Economy: Cities, Regions and International Trade*, Cambridge: MIT Press.
- [32] Greenstone, M. and E. Moretti (2004) "Bidding for Industrial Plants: Does Winning a 'Million Dollar Plant' Increase Welfare?" *NBER Working Paper*, 9844.

- [33] Greenstone, M., R. Hornbeck, and E. Moretti (2010) "Identifying Agglomeration Spillovers: Evidence from Winners and Losers of Large Plant Openings," *Journal of Political Economy*, 118(3), 536-598.
- [34] Grogger, J., and G. Hanson (2011) "Income Maximization and the Selection and Sorting of International Migrants," *Journal of Development Economics*, 95(1), 42-57.
- [35] Hanson, G. H. (1996) "Localization Economies, Vertical Organization, and Trade," *American Economic Review*, 86(5), 1266-1278.
- [36] Hanson, Gordon H. (2005) "Market Potential, Increasing Returns, and Geographic Concentration," *Journal of International Economics*, 67(1), 1-24.
- [37] Head, Keith and John Ries (2001) "Increasing Returns versus National Product Differentiation," *American Economic Review*, 91, 858-876.
- [38] Helpman, Elhanan (1998) "The Size of Regions," In *Topics in Public Economics: Theoretical and Applied Analysis*, ed. David Pines, Efraim Sadka, and Itzhak Zilcha, Cambridge: Cambridge University Press.
- [39] Hsieh, Chang-Tai and Enrico Moretti (2017) "Housing Constraints and Spatial Misallocation," University of California, Berkeley, mimeograph.
- [40] Kennan, J., and J. R. Walker (2011) "The Effect of Expected Income on Individual Migration Decisions," *Econometrica*, 79(1), 211-251.
- [41] Krugman, Paul (1991) "Increasing Returns and Economic Geography," *Journal of Political Economy*, 99(3), 483-99.
- [42] Krugman, Paul and Anthony J. Venables (1995) "Globalization and the Inequality of Nations," *Quarterly Journal of Economics*, 110(4), 857-880.
- [43] Lucas, Robert E. and Esteban Rossi-Hansberg (2002) "On the Internal Structure of Cities", *Econometrica*, 70(4), 1445-76.
- [44] Manning, Alan and Barbara Petrongolo (2011) "How Local are Labour Markets? Evidence from a Spatial Search Model", *CEP Discussion Paper*, No. 1101, London School of Economics.
- [45] McFadden, D. (1974) "The Measurement of Urban Travel Demand," *Journal of Public Economics*, 3(4), 303-328.
- [46] McFadden, Daniel and Kenneth Train (2000) "Mixed MNL Models of Discrete Response," *Journal of Applied Econometrics*, 15, 447-470.
- [47] Melitz, Marc J. (2003) "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71(6), 1695-1725.
- [48] Mian, Atif and Amir Sufi (2014) "What Explains the 2007-2009 Drop in Employment?" *Econometrica*, 82(6), 2197-2223.
- [49] Michaels, Guy (2011) "The Long Term Consequences of Resource-Based Specialisation," *Economic Journal*, 121(551), 31-57.
- [50] Mills, Edwin S. (1967) "An Aggregative Model of Resource Allocation in a Metropolitan Centre," *American Economic Review*, 57(2), 197-210.
- [51] Monte, Ferdinando (2016) "The Local Incidence of Trade Shocks," Georgetown University, mimeograph.

- [52] Moretti, Enrico (2011) "Local Labor Markets," in *Handbook of Labor Economics*, (ed.) D. Card, and O. Ashenfelter, Volume 4b, Amsterdam: Elsevier North Holland, 1238-1303.
- [53] Moretti, Enrico and Patrick Klein (2011) "Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority," *Quarterly Journal of Economics*, 129(1), 275-331.
- [54] Muth, Richard (1969) *Cities and Housing*, Chicago: University of Chicago Press.
- [55] Notowidigdo, Matthew (2013) "The Incidence of Local Labor Demand Shocks," Northwestern University, mimeograph.
- [56] Rappaport, Jordan and Jeffrey D. Sachs (2003) "The United States as a Coastal Nation", *Journal of Economic Growth*, 8(1), 5-46.
- [57] Redding, Stephen J. and Daniel Sturm (2008) "The Costs of Remoteness: Evidence from German Division and Reunification," *American Economic Review*, 98(5), 1766-1797.
- [58] Redding, Stephen J. and Matthew Turner (2015) "Transportation Costs and the Spatial Organization of Economic Activity," in (eds) Gilles Duranton, J. Vernon Henderson and William Strange, *Handbook of Urban and Regional Economics*, Chapter 20, 1339-1398.
- [59] Redding, Stephen J. (2016) "Goods Trade, Factor Mobility and Welfare," *Journal of International Economics*, 101, 148-167.
- [60] Rossi-Hansberg, Esteban (2005) "A Spatial Theory of Trade," *American Economic Review*, 95(5), 1464-1491.
- [61] Saiz, Albert (2010) "The Geographic Determinants of Housing Supply," *Quarterly Journal of Economics*, 125(3), 1253-1296.
- [62] Simonovska, Ina and Michael E. Waugh (2014) "The Elasticity of Trade: Estimates and Evidence," *Journal of International Economics*, 92(1), 34-50.
- [63] Tolbert, Charles M. and Molly Sizer (1996) "U.S. Commuting Zones and Labor Market Areas: a 1990 Update," *U.S. Department of Agriculture Staff Paper*, AGES-9614.
- [64] United States Department of Justice, Federal Bureau of Investigation. (September 2008). Crime in the United States, 2007. Retrieved 10/15/2016, from <https://www2.fbi.gov/ucr/cius2007/>.
- [65] Yagan, Danny (2016) "Moving to Opportunity? Migratory Insurance over the Great Recession," University of California, Berkeley.