ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

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

https://www.econometricsociety.org/

Econometrica, Vol. 94, No. 2 (March, 2026), 407–464

SPATIAL ECONOMICS FOR GRANULAR SETTINGS

JONATHAN I. DINGEL
Department of Economics, Columbia University, NBER, and CEPR

FELIX TINTELNOT
Department of Economics, Duke University, NBER, and CEPR

Econometrica, Vol. 94, No. 2 (March, 2026), 407–464

SPATIAL ECONOMICS FOR GRANULAR SETTINGS

JONATHAN I. DINGEL
Department of Economics, Columbia University, NBER, and CEPR

FELIX TINTELNOT
Department of Economics, Duke University, NBER, and CEPR

We examine the application of quantitative spatial models to the growing body of fine spatial data used to study local economic outcomes. In granular settings in which people choose from a large set of potential residence-workplace pairs, observed outcomes in part reflect idiosyncratic choices. Using analytical examples, Monte Carlo simulations, and event studies of neighborhood employment booms, we demonstrate that calibration procedures that equate observed shares and modeled probabilities perform very poorly in these high-dimensional settings. Parsimonious specifications of spatial linkages deliver better counterfactual predictions. To quantify the uncertainty about counterfactual outcomes induced by the idiosyncratic component of individuals' decisions, we introduce a quantitative spatial model with a finite number of individuals. Applying this model to Amazon's proposed second headquarters in New York City reveals that its predicted consequences for most neighborhoods vary substantially across realizations of the individual idiosyncrasies.

KEYWORDS: Commuting, exact hat algebra, overfitting, quantitative spatial economics.

1. INTRODUCTION

ECONOMISTS INCREASINGLY USE quantitative models to evaluate urban policies such as infrastructure investments and land-use planning. The growing availability of economic data observed at fine spatial scales offers tremendous potential for new insights. However, to inform policy, these models need to produce credible counterfactual predictions when applied to granular settings.

Quantitative spatial models feature links between pairs of locations that shape the geography of economic outcomes (Redding and Rossi-Hansberg (2017), Proost and Thisse (2019)). In many of these models, individuals choose a residence and a workplace based on locations' characteristics, the links between them, and an idiosyncratic component. A setting with granular geographic units has many locations and, therefore, a great number of links between pairs of locations. Two concerns arise when the number of links is large

Jonathan I. Dingel: j.dingel@columbia.edu
Felix Tintelnot: felix.tintelnot@duke.edu

We thank four anonymous referees, Rodrigo Adão, Gabriel Ahlfeldt, Treb Allen, Stephane Bonhomme, Kirill Borusyak, Victor Couture, Teresa Fort, Lars Hansen, Elhanan Helpman, Thomas Holmes, John Huizinga, Erik Hurst, Kyle Mangum, Yuhei Miyauchi, Eduardo Morales, Stephen Redding, Jordan Rosenthal-Kay, Esteban Rossi-Hansberg, Chris Severen, Sebastian Sotelo, Daniel Sturm, Miguel Zerecero, Oren Ziv, and numerous conference and seminar participants for helpful feedback. We are grateful to Junbiao Chen, Levi Crews, Gianna Fenaroli, Reigner Kane, Xianglong Kong, Daniil Iurchenko, Leran Qi, John Ruf, Isaac Shon, Ye Sun, Linghui Wu, Shijian Yang, and Mingjie Zhu for excellent research assistance. Dingel thanks the James S. Kemper Foundation Faculty Research Fund at the University of Chicago Booth School of Business. This work was completed in part with resources provided by the University of Chicago Research Computing Center and Columbia University's Shared Research Computing Facility. We gratefully acknowledge support from the National Science Foundation under Grant 2018609.

 https://doi.org/10.3982/ECTA19350

relative to the number of decision makers. First, if one uses a high-dimensional parameterization of spatial linkages, such as one parameter per pair of locations, the modeler runs the risk of fitting the model to the idiosyncratic components of decisions rather than capturing spatial linkages. Second, counterfactual changes in local economic outcomes are uncertain because they depend on individuals' idiosyncratic preferences.

In this article, we examine the application of quantitative spatial models to granular settings. In doing so, we make three main contributions. First, we show that the calibration procedure most often used, which perfectly matches the observed baseline outcomes, predicts changes in outcomes poorly in granular settings. We show how to use simulations that require only the inputs typically used in this procedure to diagnose the problem. We confirm the poor predictive performance using event studies of neighborhood employment booms in New York City. Second, we demonstrate that both a parsimonious parameterization using only a transit-time covariate and parameterizations using matrix approximations to capture unobserved spatial linkages deliver better counterfactual predictions in this high-dimensional setting. Third, we develop a spatial model with a finite number of individuals that we use to quantify the variability of counterfactual outcomes induced by individual idiosyncrasies. Applying this model to Amazon's proposed second headquarters in New York City shows that counterfactual changes for most census tracts vary considerably.

Section 2 reviews the computation of counterfactual outcomes in quantitative spatial models. With a continuum of individuals, counterfactual changes can be characterized by a system of equations containing counterfactual endogenous outcomes relative to baseline endogenous outcomes, counterfactual exogenous parameters relative to baseline exogenous parameters, constant elasticities, and baseline equilibrium shares (a formulation known as "exact hat algebra"). When parameterizing the spatial links between residences and workplaces that determine the baseline equilibrium shares, researchers have used two starkly different approaches. The covariates-based approach makes the cost of commuting a function of pair-specific covariates, such as transit time. The calibrated-shares procedure calibrates a cost parameter for each pair of locations to match the observed commuting shares. The contrast in baseline equilibrium shares produces contrasting predictions about counterfactual outcomes.

Calibrating pair-specific commuting costs to match observed commuting shares offers advantages and disadvantages. The number of commuters between a pair of locations in a finite sample depends on three components: spatial links correlated with included covariates, orthogonal spatial links, and individual-specific idiosyncrasies. The potential advantage of calibrating pair-specific costs to match observed shares is that this can capture spatial links not predicted by covariates like transit times, such as many Columbia University employees living in nearby university-owned residences. The potential disadvantage is that parameterizing the model's spatial links by matching all observed variation means overfitting the model: the calibrated parameters conflate the spatial links common to all individuals and individual-specific idiosyncrasies in the observed data.

Section 3 examines the application of these models to granular empirical settings. The commuting matrices used in urban economics often have so many more pairs of locations than commuters that many of the cells must contain zero commuters. Small flows constitute a substantial share of total commuting: more than 40% of New York City commuters have a residence-workplace pairing populated by five or fewer individuals. Such settings can pose small-sample problems for the calibrated-shares procedure, which equates observed shares and underlying probabilities. This procedure's expected squared error for a baseline share is its finite-sample variance.

Errors in baseline shares generate errors in predicted changes in counterfactual scenarios. We illustrate this in a special case in which the error in the predicted change in commuters is proportional to the error in the corresponding baseline share. More generally, counterfactual changes are a nonlinear function of all baseline shares. We use Monte Carlo simulations to examine the finite-sample behavior of the calibrated-shares procedure in a realistic setting. The data-generating process is the parsimonious covariates-based parameterization of New York City in 2010, which uses the same covariate employed to estimate the commuting elasticity for the calibrated-shares procedure. These simulations show that small-sample bias can seriously impair the calibrated-shares procedure when the model is otherwise correctly specified. While the procedure works well as the number of individuals becomes arbitrarily large, it produces much worse predictions in simulations using the actual number of individuals in New York City.

Next, we use event studies to examine how well the incumbent approaches predict changes in commuting flows, the key spatial linkage between neighborhoods in many models of cities. We study large employment increases in a single census tract that often stem from the arrival or expansion of a large employer. We increase the productivity parameters for these tracts to match their increases in total employment and compare the predicted changes in bilateral commuting flows to those observed in the data. When examining 83 tract-level employment booms in New York City, we find that the covariates-based approach——using transit times as the only bilateral covariate——predicts the observed changes in commuting flows better than the calibrated-shares procedure in 80 of these events. Regressing the observed change in commuters on the covariates-based approach's prediction typically yields a slope near one, whereas the predictions from the calibrated-shares procedure are *negatively* correlated with the observed changes in more than half of the events.

We explore two additional approaches that offer the advantage of a more flexible specification while guarding against the disadvantage of overfitting. One specification augments the covariates-based approach to capture unobserved spatial linkages that have an interactive fixed effects structure (Bai (2009), Chen, Fernández-Val, and Weidner (2021)). This more flexible specification modestly improves predictive performance. An alternative approach computes baseline shares used in the exact hat algebra from a rank-constrained singular value decomposition of the observed commuting matrix. This simple, covariate-free transformation performs as well as the covariates-based approach in our application.[1] Both these approaches allow one to incorporate unobserved components when applying continuum models to granular settings.

The noisy-shares problem is relevant for a number of fields in which economists rely on a "shares inversion" to estimate or calibrate models. Market shares computed from consumer surveys, region-to-region trade shares computed from shipment surveys, and migration shares computed from household surveys that cover a fraction of the population can be noisy.[2] Shares computed using the universe of commuting flows, sectoral exports,

[1]Smoothing the observed shares using rank-restricted matrix approximations before employing them in exact hat algebra is reminiscent of, yet distinct from, practices such as applying the Hodrick and Prescott (1997) filter before calibrating macroeconomic models, grouping firms into classes using k-means clustering before estimating a model of workers and firms (Bonhomme, Lamadon, and Manresa (2019)), and smoothing observations before inverting conditional choice probabilities in dynamic models (e.g., Kalouptsidi (2014), Hsiao (2022)).

[2]The US Commodity Flow Survey, for example, samples about one in seven establishments in the industries it covers and asks them to report at most 40 of the shipments sent during each of 4 weeks within the year. See Appendix D.5 in the Supplemental Appendix (Dingel and Tintelnot (2026)) for the case of state-to-state migration shares.

firm-to-firm sales, or employee–employer matches are noisy when there are few decisions or transactions made among the many possibilities so that the observed shares contain a substantial idiosyncratic component.[3]

In Section 4, we depart from treating the data as a finite sample from the continuum model and introduce a model with a finite number of individuals. We use it to quantify the variability of counterfactual outcomes induced by individual idiosyncrasies. Where the continuum model yields a single equilibrium allocation, this model produces a *distribution* of economic outcomes. Notably, outcomes with strictly positive probabilities may not occur in a particular realization. The key modeling challenge is that individual decisions affect equilibrium wages and rents, but it is computationally infeasible to enumerate all the possible labor allocations and resulting prices. To overcome this combinatorial challenge, we assume that individuals make choices based on wage and rent beliefs that are the equilibrium prices from the continuum model with the same economic primitives.

Our model with finite individuals can be estimated using the same data used to estimate the continuum model. When individuals' price beliefs are the continuum-case equilibrium prices, the estimated parameters of this model coincide with those of the covariates-based continuum model. Given the same parameter values, the mean equilibrium labor allocation equals the allocation of the continuum model. By contrast, equilibrium prices solve a system of nonlinear equations, so their mean values are not necessarily equal to the continuum-case prices. Crucially, the model with a finite number of individuals characterizes the dispersion in quantities and prices that arises from idiosyncrasies.

In Section 5, we contrast methods for evaluating the local economic effects of Amazon's proposed—but later abandoned—HQ2 in New York City. The tract-level changes in residents and land rents across New York City predicted by the covariates-based approach differ considerably from the predictions of the calibrated-shares procedure. We find considerable uncertainty generated by individual idiosyncrasies for these outcomes at the neighborhood level and meaningful uncertainty even for groups of tracts. The uncertainty about aggregate outcomes is negligible. The variation in local counterfactual outcomes due to individual idiosyncrasies in the model with a finite number of individuals is larger than that due to parameter uncertainty in the covariates-based continuum model. This sizable uncertainty regarding the local effects of a headline-grabbing potential employment boom suggests that detailed counterfactual predictions would be highly uncertain in many granular settings.

Our paper introduces a quantitative spatial model suitable for application to small spatial units, such as census tracts. Social scientists study neighborhoods to understand housing markets, intergenerational mobility, racial segregation, and many other phenomena. Spatially precise satellite imagery (Donaldson and Storeygard (2016)) and phone movement data (CDGH+ (2022), Kreindler and Miyauchi (2023)) allow even finer investigations where there are few decision makers, as anticipated by Holmes and Sieg (2015, p. 106). Granular data are valuable in part because a growing body of evidence shows highly localized agglomeration economies (Rosenthal and Strange (2020)).[4] We provide tools for modeling these granular settings.

[3]For example, Gaubert and Itskhoki (2021) report that sectoral trade flows dominated by a small number of firms exhibit greater mean reversion.

[4]Arzaghi and Henderson (2008) estimate that productivity gains from interactions among advertising firms occur primarily within 500 meters. Ahlfeldt, Redding, Sturm, and Wolf (2015) estimate production and residential externalities that decay by more than half within two minutes of travel time. Rossi-Hansberg, Sarte, and Owens (2010) estimate that externalities from nonmarket interactions decline by half approximately every 1000 feet.

Our event studies of neighborhood employment booms contribute to a small literature assessing the predictive power of quantitative spatial models. These studies compare model-predicted changes to observed outcomes using plausibly exogenous shocks (Ahlfeldt, Redding, Sturm, and Wolf (2015), Monte, Redding, and Rossi-Hansberg (2018), Adão, Arkolakis, and Esposito (2022)). More broadly, economists have evaluated the predictions of quantitative trade models about changes in trade policy (Kehoe (2005), Kehoe, Pujolàs, and Rossbach (2017), Adão, Costinot, and Donaldson (2024)). We show that predictions about the "ripple effects" of a local economic shock on other locations perform poorly when economic linkages are calibrated to match noisy shares. In granular settings, this noise stems from the idiosyncratic components of individual choices.

Finally, our work relates to prior studies of how idiosyncrasies affect economic outcomes and parameter estimation. A growing literature examines how firm-specific shocks influence aggregate fluctuations in macroeconomics (Gabaix (2011), Carvalho and Grassi (2019)) and trade (di Giovanni, Levchenko, and Mejean (2014), Gaubert and Itskhoki (2021)). Our paper joins research examining how individual idiosyncrasies affect economists' inferences. Ellison and Glaeser (1997) and Schoefer and Ziv (2022) address the role of individual manufacturing plants when computing the geographic concentration of industries and geographic variation in productivity, respectively. Mogstad, Romano, Shaikh, and Wilhelm (2024) raise concerns about inferring ranks of small geographic units. The finite numbers of firms and shipments have been studied as one explanation for zeros in international trade (Eaton, Kortum, and Sotelo (2013), Armenter and Koren (2014)), and the finite number of consumers is one explanation for zeros in product-level sales data (Quan and Williams (2018), HNPS+ (2023), Gandhi, Lu, and Shi (2023)). Similar to our framework, Panigrahi (2022) models a finite set of firms that form links based on beliefs from a limiting economy with a continuum of firms. Our paper makes two contributions to these various strands of the literature. First, we demonstrate the importance of avoiding overfitting when using continuum models to predict counterfactual outcomes. Second, we introduce a spatial model with a finite number of individuals that quantifies the uncertainty about counterfactual outcomes induced by individual idiosyncrasies.

The Supplemental Appendix (Dingel and Tintelnot (2026)) contains additional proofs, details, and results.

2. COMPUTING COUNTERFACTUAL OUTCOMES IN CONTINUUM MODELS

This section describes the computation of counterfactual outcomes in quantitative spatial models. We use a baseline model with a continuum of individuals making residential and workplace location choices, as in Ahlfeldt, Redding, Sturm, and Wolf (2015) and Redding and Rossi-Hansberg (2017).

2.1. *Primitives*

We consider a closed economy populated by a measure of L individuals who each supply one unit of labor. Discrete locations are indexed by k or n. Each location has a fixed quantity of land, T_k, owned by immobile landlords who consume only goods.[5] Each location has the technology to produce a differentiated good (i.e., the Armington assumption). These goods are freely traded.

[5]This simplifying assumption follows Monte, Redding, and Rossi-Hansberg (2018).

Individuals have Cobb–Douglas preferences over goods and land, devoting α of their expenditure to the latter. They have constant elasticity of substitution (CES) preferences over the differentiated goods, with elasticity of substitution $\sigma > 1$. Individuals (indexed by i) have idiosyncratic preferences for pairs of residential and workplace locations, such that i's indirect utility from residing in k and working in n is

$$U_{kn}^i = \underbrace{\epsilon \ln\left(\frac{w_n}{r_k^\alpha P^{1-\alpha}\delta_{kn}}\right)}_{\equiv U_{kn}} + \nu_{kn}^i, \tag{1}$$

where w_n is the wage in location n, r_k is the land rent in k, P is the common CES price index for goods, δ_{kn} is the commuting cost between k and n, U_{kn} is the mean utility of choice kn, and ν_{kn}^i is the idiosyncratic preference of individual i for residing in k and working in n. The commuting elasticity ϵ governs the importance of mean utility relative to the idiosyncratic preference, which is drawn from the standard Gumbel distribution.[6]

Commuting costs have two components: time ($\bar{\delta}_{kn}$) and disutility (λ_{kn}), so $\delta_{kn} = \bar{\delta}_{kn} \times \lambda_{kn}$. Time spent commuting is not spent working, so individuals residing in k and working in n earn only $w_n/\bar{\delta}_{kn}$ because they only spend $1/\bar{\delta}_{kn}$ of their time working.

Production of each location's differentiated good is linear in labor. Output in n is $q_n = A_n L_n$, where A_n is that location's productivity and L_n is the quantity of labor supplied by workers working in n. The goods market is competitive, so the price of location n's output is w_n/A_n. Thus, the CES price index is $P = [\sum_n (w_n/A_n)^{1-\sigma}]^{1/(1-\sigma)}$.

2.2. *Equilibrium*

Let ℓ_{kn} denote the measure of people residing in k and working in n. Using the Gumbel distribution assumption, one can integrate over the idiosyncratic preferences to obtain the fraction of people residing in k and working in n:

$$\frac{\ell_{kn}}{L} = \frac{w_n^\epsilon \left(r_k^\alpha \delta_{kn}\right)^{-\epsilon}}{\sum_{k',n'} w_{n'}^\epsilon \left(r_{k'}^\alpha \delta_{k'n'}\right)^{-\epsilon}}. \tag{2}$$

Goods market clearing equates each location's output to the quantity demanded. Labor supplied in location n is $L_n = \sum_k \ell_{kn}/\bar{\delta}_{kn}$, and thus output there is $A_n \sum_k \ell_{kn}/\bar{\delta}_{kn}$. Each individual devotes $1-\alpha$ of their expenditure to differentiated goods and α of their expenditure to land, while immobile landlords spend all of their income on differentiated goods, such that total expenditure on differentiated goods equals aggregate income. Aggregate income is $Y \equiv \sum_{k,n} y_{kn}$, where $y_{kn} \equiv w_n \ell_{kn}/\bar{\delta}_{kn}$ is labor income earned in workplace n by workers residing in k. The CES demand for each differentiated good means that equating quantity supplied and quantity demanded requires

$$A_n \sum_k \frac{\ell_{kn}}{\bar{\delta}_{kn}} = \frac{(w_n/A_n)^{-\sigma}}{P^{1-\sigma}} Y \quad \forall n. \tag{3}$$

Note that goods market clearing implies labor market clearing.

[6]The cumulative distribution function of the Gumbel (type-1 extreme value) distribution with a location parameter of zero and scale parameter of one is $F(\nu_{kn}^i) = \exp(-\exp(-\nu_{kn}^i))$.

Similarly, land market clearing equates the fixed land endowment T_k to the quantity demanded by individuals, who devote a constant fraction α of their expenditure to land:

$$T_k = \frac{\alpha}{r_k} \sum_n y_{kn} \quad \forall k. \tag{4}$$

Equilibrium is a set of wages $\{w_n\}$, rents $\{r_k\}$, and labor allocation $\{\ell_{kn}\}$ such that equations (2), (3), and (4) hold. Appendix C.1 shows that if $(\frac{1+\epsilon}{\sigma+\epsilon})(\frac{\alpha\epsilon}{1+\alpha\epsilon}) \leq \frac{1}{2}$, this equilibrium exists and is unique (by Theorem 1 of Allen, Arkolakis, and Li (2023)).

2.3. *Procedures for Counterfactual Predictions*

In this model, counterfactual changes can be characterized by a system of equations containing counterfactual endogenous outcomes relative to baseline endogenous outcomes, counterfactual exogenous parameters relative to baseline exogenous parameters, constant elasticities, and baseline equilibrium shares. This is known as "exact hat algebra" in the trade literature (Dekle, Eaton, and Kortum (2008), Costinot and Rodríguez-Clare (2014)). Denote the counterfactual value of a variable x by x' and denote the counterfactual-baseline ratio of a variable $x > 0$ by $\hat{x} \equiv \frac{x'}{x}$ (and impose $\hat{x} = 0$ if $x = 0$). As shown in Appendix C.2, tedious manipulation of equations (2), (3), and (4) yields the following system of equations characterizing the counterfactual equilibrium:

$$\hat{w}_n = \hat{A}_n \left(\sum_k \hat{y}_{kn} \frac{y_{kn}}{\sum_{k'} y_{k'n}} \right)^{\frac{1}{1-\sigma}} \hat{P} \hat{Y}^{\frac{1}{\sigma-1}}, \tag{5}$$

$$\hat{r}_k = \hat{T}_k^{-1} \sum_n \hat{y}_{kn} \frac{y_{kn}}{\sum_{n'} y_{kn'}}, \tag{6}$$

$$\hat{\ell}_{kn} = \frac{\hat{w}_n^{\epsilon} \left(\hat{r}_k^{\alpha} \hat{\bar{\delta}}_{kn} \hat{\lambda}_{kn} \right)^{-\epsilon}}{\sum_{k',n'} \hat{w}_{n'}^{\epsilon} \left(\hat{r}_{k'}^{\alpha} \hat{\bar{\delta}}_{k'n'} \hat{\lambda}_{k'n'} \right)^{-\epsilon} \frac{\ell_{k'n'}}{L}} \quad \text{if } \ell_{kn} > 0. \tag{7}$$

This system defines the relative endogenous outcomes $\hat{w}_n$, $\hat{r}_k$, and $\hat{\ell}_{kn}$ in terms of relative exogenous parameters $\hat{A}_n$, $\hat{T}_k$, $\hat{\bar{\delta}}_{kn}$, and $\hat{\lambda}_{kn}$, elasticities σ, α, and ϵ, and baseline shares $\frac{\ell_{kn}}{L}$ and $\frac{y_{kn}}{Y}$.

There are many ways to estimate or calibrate the model's parameters to deliver the baseline shares. At one end of the spectrum, one can parsimoniously parameterize the time and disutility components of commuting costs as functions of only observed covariates (e.g., Ahlfeldt, Redding, Sturm, and Wolf (2015)).[7] To compute counterfactual equilibria, one would plug the *fitted* model's values of $\frac{\ell_{kn}}{L}$ and $\frac{y_{kn}}{Y}$ into equations (5)–(7).[8] At

[7] In this approach, equation (2) is generically not satisfied by the observed values of ℓ_{kn} and the parameterization of δ_{kn}. One justification for the model not fitting the observed data would be that the latter is a sample from the population. Section 3.2 describes estimation of model parameters in this covariates-based approach.

[8] Given a parameterized model, one could also solve the system of equations (2)–(4) at counterfactual parameter values. However, the system of equations (5)–(7) is instructive because it indicates that any parameter

the other end of the spectrum, one can calibrate millions of commuting-cost parameters so that the model's baseline equilibrium matches the observed shares. In this case, one would plug the *observed* $\frac{\ell_{kn}}{L}$ and $\frac{y_{kn}}{Y}$ into equations (5)–(7) to compute counterfactual outcomes. The latter calibrated-shares procedure has been widely used in spatial economics recently, often in granular settings with as many residence-workplace pairs as decision makers.[9] Between these two extremes lie a host of strategies to determine the baseline equilibrium shares, some of which we discuss in detail below.

We emphasize the distinction between using the comparative statics defined by equations (5)–(7) to compute counterfactual outcomes and fitting the model's parameters. Because equations (5)–(7) show that computing counterfactual outcomes only requires knowing the model's parameters up to the point where the model delivers the shares $\frac{\ell_{kn}}{L}$ and $\frac{y_{kn}}{Y}$, others have used the phrase "exact hat algebra" to refer to both rewriting the equations in hats and calibrating combinations of model parameters to rationalize observed shares. In fact, the system of equations defines counterfactual outcomes regardless of how one estimates or calibrates the parameters of the baseline equilibrium. The key question is how to fit the model's parameters to data.

3. COUNTERFACTUAL ANALYSIS IN GRANULAR SETTINGS

We assess the predictive performance of quantitative spatial models in granular empirical settings. In Section 3.1, we introduce data on tract-to-tract commuting in New York City. This empirical setting, like many studied in prior research, has many location pairs relative to the number of decision makers. After Section 3.2 describes how we estimate model parameters, Section 3.3 contrasts the nature of the estimation errors made by the covariates-based and calibrated-shares procedures. The covariates-based specification omits unobserved commuting costs, whereas the calibrated-shares procedure includes finite-sample noise in parameter values. We examine this trade-off when modeling New York City using Monte Carlo simulations in Section 3.4 and event studies of neighborhood employment booms in Section 3.5. There is a stark contrast between the two incumbent approaches in predicting the changes in tract-to-tract commuting flows. While the covariates-based specification using transit times predicts quite well, the calibrated-shares procedure's predicted changes are in fact negatively correlated with the observed changes in the majority of events. We also propose and assess alternative specifications that can capture unobserved components of spatial linkages while guarding against overfitting.

3.1. *Commuting Flows in Granular Settings*

Urban economists often study empirical settings with a modest number of decision makers relative to the number of pairs of locations. It is common for these numbers to

combination yielding the same baseline shares and elasticities would make the same counterfactual predictions. In that sense, these are sufficient statistics for the counterfactual outcomes. Appendix Section C.2.1 provides an example.

[9]The calibrated-shares procedure has decades-deep roots in computable general equilibrium models of international trade (Rutherford (1995)). Studies in spatial economics employing this technique include Allen, Arkolakis, and Li (2016), Perez-Cervantes (2016), Waddell and Sarte (2016), Monte, Redding, and Rossi-Hansberg (2018), Heblich, Redding, and Sturm (2020), Krebs and Pflüger (2023), Severen (2023); and Owens, Rossi-Hansberg, and Sarte (2020).

be roughly equal.[10] At even finer resolution, Ahlfeldt, Redding, Sturm, and Wolf (2015) model about 3 million Berliners choosing among about 254 million pairs of city blocks. Here, we examine tract-to-tract commuting flows in New York City. Appendix D.1 reports comparable statistics for Detroit and Minneapolis-St. Paul.

We use data on commuting between census tracts in New York City taken from the Longitudinal Employer-Household Dynamics, Origin-Destination Employment Statistics (LODES).[11] Tracts are defined by the US Census Bureau to be relatively uniform in population size, such that the typical tract has 4000 residents. One implication of granularity is that some data providers perturb observations to protect confidentiality. For example, the published LODES tract-level workplace employment counts are infused with noise. Moreover, the published LODES commuting matrices report synthetically generated residence-workplace pairs.[12] These confidentiality-protecting procedures are both a symptom of the granular setting and an additional cause for caution when interpreting the value of any individual observation.

Small flows constitute a large share of tract-to-tract commuting in New York City. The city has about 2.5 million resident-employees and 4.6 million tract pairs. Thus, its commuting matrix is almost necessarily sparse, as the average cell contains about one-half of a commuter. In fact, 85% of tract pairs have zero commuters between them, so the average positive cell has about three people. As shown in Figure 1, more than half of the tract pairs with a positive number of commuters have only one. Crucially, Figure 1 shows that small flows are a substantial share of the total: 41.1% of New York City commuters have five or fewer commuters in their cell of the commuting matrix. Furthermore, the commuting flows are impersistent over time and exhibit asymmetric zeros (see Appendices D.3 and D.4). Introducing multiple worker types would exacerbate these patterns by reducing the number of decision makers per choice for each type.[13]

Even when studying larger geographic units, there may be few decision makers per pair of locations, particularly when the data describe only a sample of respondents. For example, while some US counties have millions of residents, Appendix D.2 documents that more than half of the positive commuting counts between pairs of counties represent the behavior of five or fewer survey respondents. More broadly, small-sample bias may complicate the analysis of other spatial phenomena, such as migration. Appendix D.5 shows that state-to-state migration flows reported in American Community Survey (ACS) data also represent a small number of respondents, are frequently zero, and are impersistent across time (see Foschi, House, Proebsting, and Tesar (2023) for related findings).

[10] Owens, Rossi-Hansberg, and Sarte (2020) study 1.3 million people commuting between 1.3 million pairs of tracts in the Detroit urban area. Severen (2023) studies the Los Angeles metropolitan area, which has 6.7 million commuters and more than 6 million pairs of tracts. Tsivanidis (2023) models Bogotá's 8 million residents commuting between almost 7 million pairs of tracts. Zárate (2023) examines Mexico City, which has 9 million people and 13 million pairs of tracts.

[11] Davis, Dingel, Monras, and Morales (2019) use the 2010 wave to compute the joint distribution of residences and workplaces for New York City. Owens, Rossi-Hansberg, and Sarte (2020) use the 2014 wave of these data for the Detroit urban area to estimate a gravity model of commuting.

[12] Graham, Kutzbach, and McKenzie (2014): "For each job in a workplace cell, LODES draws from a Dirichlet multinomial posterior distribution of possible residential locations. … The prior adds uncertainty, so that even commutes with few or no observed flows may appear to have a job. Conversely, even when there are commuters from an origin in the likelihood, that residence may not be drawn and thus would not appear in LODES."

[13] For example, individuals in each of the four educational attainment categories reported in the LODES data live in all but one or two residential tracts and are employed in more than 99% of workplace tracts, so each type would have a choice set with more than 4.5 million elements.

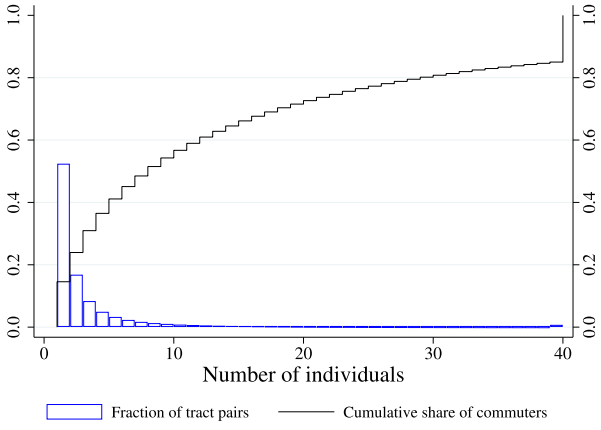

FIGURE 1.—Number of commuters between pairs of tracts in New York City. NOTES: This figure describes tract pairs in New York City by the number of individuals who reside in the origin tract and work in the destination tract in 2010 LODES data. The sample is restricted to the 15% of tract pairs that have a strictly positive number of commuters. The blue histogram depicts the fraction of tract pairs reporting each number of commuters, for each value from 1 to 39 and values of 40 or greater. The black increasing step function depicts the cumulative share of commuters by the number of commuters between the pair of tracts. New York City has 2160 residential census tracts.

3.2. *Estimation*

We estimate the continuum model using the tract-to-tract commuting data for New York City in 2010. Following the literature, we set $\alpha = 0.24$ and $\sigma = 4$ prior to estimating the remaining parameters.[14]

We assume that the time component of commuting costs, $\bar{\delta}_{kn}$, is observed, while the disutility component, λ_{kn}, is unobserved. We construct the observed part of the commuting costs by assuming that each worker has H hours that are spent either working or commuting. We compute $\bar{\delta}_{kn} = \frac{H}{H - t_{kn} - t_{nk}}$, where $H = 9$ hours and t_{kn} is the transit time from k to n according to Google Maps.[15] Assumptions about the unobserved disutility component, λ_{kn}, vary widely across methods, as described below.

One way to estimate the continuum model using the covariates-based approach is to interpret observed data on residence-workplace choices as a finite sample from the continuum model. The covariates-based specification assumes there is no unobserved disutility component of commuting costs, $\lambda_{kn} = 1 \; \forall k, n$. Given observed values of commuting costs $\delta_{kn} = \bar{\delta}_{kn}$, the remaining parameters can be estimated by maximum likelihood. Interpreting the right side of equation (2) as a probability and denoting the number of individuals

[14] Davis and Ortalo-Magne (2011) suggest $\alpha = 0.24$. Monte, Redding, and Rossi-Hansberg (2018) use $\sigma = 4$ when examining goods trade between US counties, citing estimates for international trade between countries. Absent estimates of tract-level labor demand elasticities, we also report results for $\sigma = 1.1$ and $\sigma = \infty$. None of our main conclusions are sensitive to the value of σ.

[15] Given H hours, $1/\bar{\delta}_{kn}$ is the share of that time spent working if the individual resides in k and works in n. The semielasticity of commuting flows with respect to transit time $\mathrm{t}_{kn} = t_{kn} + t_{nk}$ is $-\epsilon/(H - \mathrm{t}_{kn})$. Using $H = 8$ or $H = 10$ yields very little change in the model fit relative to the $H = 9$ results reported in Table I. For New York City, we use Google Maps public-transit times collected by Davis, Dingel, Monras, and Morales (2019). We impute missing observations for fewer than 4.1% of tract pairs in New York City by predicting transit times using physical distance.

who chose the kn pair by ℓ_{kn}, the log likelihood function is

$$\mathcal{L} \equiv \sum_{k,n} \ell_{kn} \ln\left[\mathbb{P}\left(U^i_{kn} > U^i_{k'n'} \ \forall k'n' \neq kn\right)\right] = \sum_{k,n} \ell_{kn} \ln\left[\frac{w_n^{\epsilon}\left(r_k^{\alpha}\bar{\delta}_{kn}\right)^{-\epsilon}}{\sum_{k',n'} w_{n'}^{\epsilon}\left(r_{k'}^{\alpha}\bar{\delta}_{k'n'}\right)^{-\epsilon}}\right]. \tag{8}$$

This is the canonical conditional-logit likelihood of McFadden (1974) applied to location choices, as in McFadden (1978). The k- and n-specific terms are captured by residence and workplace fixed effects, respectively. Maximizing this likelihood function is numerically equivalent to a Poisson pseudo maximum likelihood (PPML) estimator that is available for a variety of software packages (Guimarães, Figueiredo, and Woodward (2003)).[16]

The calibrated-shares procedure, which calibrates λ_{kn}, also requires an estimate of the commuting elasticity ϵ. One might make two distinct assumptions about the unobserved component λ_{kn}. If $\mathbb{E}(\lambda_{kn}^{-\epsilon}|r_k, w_n, \bar{\delta}_{kn}) = 1$, the constant elasticity function (2) can be estimated using the PPML estimator (Silva and Tenreyro (2006)). This yields the same estimate of the commuting elasticity ϵ as maximizing equation (8). If one takes the logarithm of each side of equation (2), assumes $\mathbb{E}(\ln \lambda_{kn}|r_k, w_n, \bar{\delta}_{kn}) = 0$, and restricts the estimation sample to observations for which ℓ_{kn} is strictly positive, then the commuting elasticity can be estimated by ordinary least squares (OLS). As stressed by Silva and Tenreyro (2006), these two estimators can yield very different parameter estimates.

We apply both the PPML and OLS estimators to New York City in 2010. Our estimate of the commuting cost elasticity is presented in column 1 of Table I. The estimate of $\epsilon \approx 8$, which is comparable to the value of 6.8 estimated by Ahlfeldt, Redding, Sturm, and Wolf (2015) for commuting within Berlin, implies that idiosyncratic preferences are modestly dispersed. Estimating the commuting elasticity using OLS yields a much lower elasticity, largely because of the well-understood selection bias associated with omitting four-fifths of observations from the estimation sample.[17] In what follows, we use the ϵ estimate reported in column 1 for both the covariates-based approach and the calibrated-shares procedure.

TABLE I

COMMUTING ELASTICITY ESTIMATES.

	PPML/MLE	OLS
Commuting cost	−7.986 (0.307)	−2.307 (0.00768)
Model fit (R^2 or pseudo-R^2)	0.662	0.561
Location pairs	4,628,878	690,673
Commuters	2,488,905	2,488,905

Note: All specifications include residence fixed effects and workplace fixed effects. The "PPML/MLE" column presents the results from maximum likelihood estimation of equation (8). The "OLS" column presents the results of estimating the log version of equation (2) by ordinary least squares, omitting observations in which $\ell_{kn} = 0$. The model-fit statistic is the pseudo-R^2 for MLE and R^2 for OLS. We report the PPML standard errors (clustered by k and by n), which are larger than the logit MLE standard errors associated with maximizing equation (8).

[16]See Sotelo (2019) for a discussion of the relationship between the multinomial and PPML estimators in the context of gravity models. In practice, we use the Stata package `ppmlhdfe` of Correia, Guimarães, and Zylkin (2020).

[17]Applying the PPML estimator to the OLS sample of strictly positive flows yields an estimate of $\epsilon = 4.4$. Thus, about two-thirds of the discrepancy between the OLS and MLE estimates reflects the omission of the zero flows. This selection bias also affects the fixed effects (see Appendix D.6).

Given the elasticities σ, α, and ϵ, one only needs the model-implied baseline shares $\frac{\ell_{kn}}{L}$ and $\frac{y_{kn}}{Y}$ to compute counterfactual outcomes using equations (5)–(7). The calibrated-shares procedure sets the parameters $\{A_n\}$, $\{T_k\}$, and $\{\lambda_{kn}\}$ such that the model-implied baseline shares $\frac{\ell_{kn}}{L}$ and $\frac{y_{kn}}{Y}$ equal their observed values. For the covariates-based approach, the model parameters $\{A_n\}$ and $\{T_k\}$ can be computed from the estimated fixed effects and the market-clearing conditions. In particular, the residence and workplace fixed effects are proportional to $r_k^{-\alpha\epsilon}$ and w_n^{ϵ}, respectively. Given α, σ, ϵ, $\{\delta_{kn}\}$, $\{r_k\}$, and $\{w_n\}$, equations (2), (3), and (4) can be solved to obtain $\{T_k\}$ and $\{A_n\}$, and one can compute the model-implied baseline shares $\frac{\ell_{kn}}{L}$ and $\frac{y_{kn}}{Y}$.

3.3. *Errors in Estimated Baseline Shares and Prediction Errors*

In practice, each estimation method will imperfectly recover the data-generating process. Estimation error is important because errors in baseline shares translate into errors in counterfactual predictions through their roles in equations (5)–(7). For example, consider an economy in which $\sigma = \infty$ and $\alpha = 0$ with exogenous residential amenities that has a productivity change at one workplace. In this special case, the error in the predicted change in commuters $\Delta\ell_{kn}$ is simply proportional to the error in the baseline share ℓ_{kn}/L (see Appendix C.4).

The two approaches' estimation errors stem from different sources, so researchers face a trade-off. If there are unobserved links ($\lambda_{kn} \neq 1$), the covariates-based model is misspecified. If observed shares ℓ_{kn}/L are noisy, the calibrated-shares procedure will yield noisy estimates of λ_{kn}.[18] Thus, the trade-off depends on the magnitude of unobserved commuting costs and the number of decision makers per pair of locations. To justify applying the calibrated-shares procedure to a granular setting, the researcher must believe unobserved spatial links play a large role.

The potential advantage of the calibrated-shares procedure is its ability to capture links that covariates miss. For example, the largest commuting flow in the 2010 LODES data for New York City is 827 commuters who reside between 110th and 114th Streets in Morningside Heights and work at adjacent Columbia University. The specification that uses transit times predicts only 70 of the 827 observed commuters, failing to capture the effect of the university's dual role as employer and landlord.

The disadvantage of this approach is the sheer number of parameters to be estimated. If the economy has N locations, the calibrated-shares procedure estimates N^2 commuting-cost parameters, whereas the covariates-based specification only estimates $2N$ fixed effects (and the coefficients on the covariates).[19] In New York City, $2N \ll I \ll N^2$ (because I is 2.5 million and N^2 is 4.6 million), so the calibrated-shares procedure will produce estimates incorporating substantial finite-sample noise. The parsimonious covariates-based specification exhibits much less sampling error.[20]

[18] We focus on the finite-sample noise stemming from a granular setting. Sanders (2024) considers how other forms of measurement error affect counterfactual outcomes in quantitative spatial models.

[19] While we focus on finite-sample properties, this difference in dimensionality has implications for asymptotic properties like consistency as both the number of individuals and the number of locations grow. Suppose the number of locations N is an power function of the number of individuals I: $N \propto I^a$. The ratio of individuals to parameters grows at rate $1 - a$ for the covariates-based approach and $1 - 2a$ for the calibrated-shares procedure. For $a \in (0.5, 1)$, $\frac{I}{N}$ goes to infinity and $\frac{I}{N^2}$ goes to zero. In the United States, the Census Bureau aims for Census tracts to have constant size of about 4000 residents, so a is about one.

[20] Appendix C.6 presents a simple example for the NYC setting in which sampling errors differ by orders of magnitude.

We examine the predictions made by these two approaches using two metrics. One way to assess predictions is to regress the changes in outcomes on the predicted changes. An unbiased forecast would yield a slope coefficient of one. In an otherwise correctly specified model, finite-sample noise will typically attenuate the slope coefficient on the calibrated-shares procedure's predicted changes towards zero.[21] The covariates-based specification's omission of the unobserved commuting cost does not necessarily attenuate its slope coefficient because the omitted residual is orthogonal to the covariates by construction. Appendix C.5 provides analytical expressions for the slope coefficient in the special case mentioned above in which predicted changes are proportional to baseline shares.

A second way to assess predictions is to compute their mean squared errors. In the special case, the expected squared error for a predicted change in commuters is proportional to the expected squared error for the corresponding baseline share. The calibrated-shares procedure's expected squared error for a baseline share is its finite-sample variance, which vanishes as I goes to infinity. By imposing more structure, a covariates-based approach typically has lower finite-sample variance but omits the unobserved commuting cost that contributes to its expected squared error for a baseline share. Their relative performance depends on how many commuters are observed per residence-workplace pair and the magnitudes of the unobserved commuting costs. Appendix C.6 derives this result in a simple example.

Errors in baseline shares generate errors in predicted changes in counterfactual scenarios. The magnitude of the resulting errors hinges on two key questions: how large are the errors in the baseline shares and how sensitive is the counterfactual outcome of interest to individual linkages? Appendix C.4 uses the special case to illustrate that residential responses to a workplace productivity shock are more sensitive to sampling errors in baseline shares ℓ_{kn}/L than employment responses. This is because the former depends on the baseline share of a single residence-workplace pair, whereas the latter depends only on the workplace's share of total employment. Beyond this special case, the predicted changes in these models are generally a nonlinear function of the economy's baseline equilibrium quantities. To explore a more typical case, we turn to Monte Carlo simulations.

3.4. *Monte Carlo: The Calibrated-Shares Procedure in a Granular Setting*

This section uses Monte Carlo simulations to assess how well the calibrated-shares procedure for computing counterfactual outcomes performs in a granular setting. In anticipation of examining New York City neighborhood employment booms in the next section, the data-generating process is the estimated covariates-based model of New York City in 2010. We imagine the researcher observes equilibrium wages and a finite-sample labor allocation drawn from a multinomial distribution defined by probabilities given by the right side of equation (2). The researcher then uses the calibrated-shares procedure to predict the consequences of a "counterfactual" productivity increase in one workplace tract. The "true" consequences of this productivity change are given by the counterfactual labor allocation and prices in the covariates-based continuum model. If the observed sample

[21] To gain intuition, consider a regression of the true probabilities on finite-sample shares. Because the frequency estimator is an unbiased estimator of the true probability, the finite-sample share is equal to the true probability plus finite-sample error. With infinite observations, the slope coefficient would be one, and in a typical finite sample, the slope coefficient would suffer attenuation bias, akin to classical measurement error.

had an infinite number of individuals, the realized labor allocation would equal the true shares (because the frequency estimator is consistent), and the calibrated-shares procedure would perfectly predict the changes caused by the productivity increase. In short, our simulations examine the finite-sample behavior of the calibrated-shares procedure using the actual number of individuals in New York City ($I \approx 2.5$ million).

We contrast the true and predicted changes in the number of commuters from each residential origin to the workplace with increased productivity. As detailed in Appendix A.1, the modeled productivity increases for both the covariates-based model and the calibrated-shares procedure are engineered so that the total employment increase in the "treated" tract is perfectly predicted. We assess the predictive power by contrasting true and predicted changes in bilateral commuting flows and rents. We do so by computing the mean squared error (MSE) and by regressing the observed changes on the predicted changes. With an unbiased forecast, the regression of observed on predicted changes would yield a slope coefficient of one and an intercept of zero. Of course, the estimated coefficients in a single simulation will also reflect finite-sample noise. Simulating this process 100 times yields a distribution of regression coefficients and MSEs, as shown in Figure 2.

Before assessing the calibrated-shares procedure's predictive performance, we note that the covariates-based model is well behaved in the finite sample. The covariates-based model is correctly specified because $\lambda_{kn} = 1 \; \forall k, n$ in the data-generating process. It delivers accurate predictions when applied to a finite sample: its slope and intercept coefficients are centered on one and zero, respectively. It also accurately predicts the changes in rents (see Appendix Figure A.1).

The calibrated-shares procedure makes worse predictions. While it perfectly fits the commuting shares from the finite sample, it produces less accurate predictions about the changes in commuters and rents caused by the productivity increase. This illustrates an overfitting problem: a more flexible parameterization improves in-sample fit but wors-

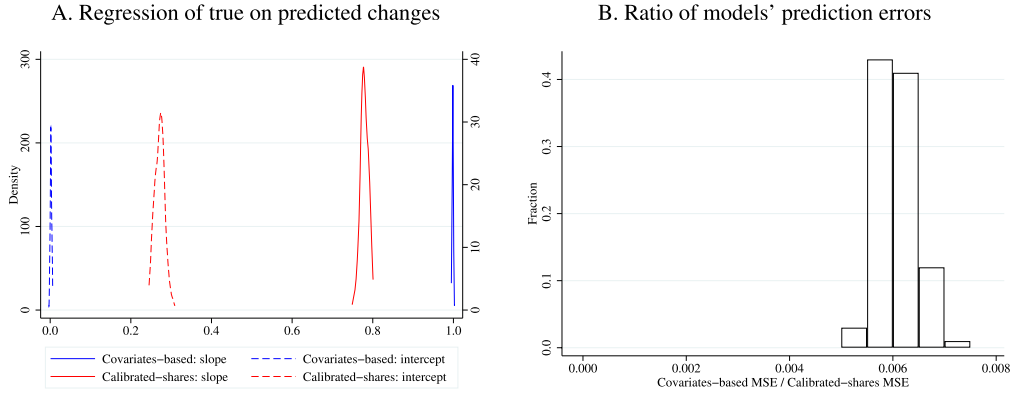

FIGURE 2.—Calibrated-shares procedure overfits in Monte Carlo simulations. NOTES: This figure depicts the regression coefficients and mean squared errors from 100 simulations in which $I = 2{,}488{,}905$, the data-generating process is the estimated covariates-based model of New York City in 2010, and the counterfactual change is a 9% increase in productivity in one workplace tract. See Appendix A.1 for a detailed description of the simulation procedure. "Covariates-based" means estimating the covariates-based parameterization described in Section 3.2 using the simulated data. "Calibrated-shares" means using the observed (simulated) shares in the exact hat algebra defined by equations (5)–(7). In panel A, the densities for the covariates-based model are on the left vertical axis and the densities for the calibrated-shares procedure are on the right vertical axis.

TABLE II

CALIBRATED-SHARES PROCEDURE'S FINITE-SAMPLE PERFORMANCE.

I	2.5	5	12.5	25	50	125	250	2560
A. Regressand is continuum change in commuters								
Calibrated-shares: slope	0.779	0.876	0.948	0.973	0.987	0.995	0.997	1.000
Calibrated-shares: intercept	0.272	0.153	0.064	0.034	0.016	0.006	0.003	0.000
Calibrated-shares: MSE	0.228	0.113	0.045	0.023	0.011	0.005	0.002	0.000
B. Regressand is finite-sample change in commuters								
Calibrated-shares: slope	−0.437	0.228	0.670	0.822	0.917	0.969	0.983	0.998
Calibrated-shares: intercept	1.739	0.960	0.406	0.219	0.102	0.039	0.021	0.003
Calibrated-shares: MSE	17.125	8.518	3.418	1.697	0.846	0.340	0.170	0.017

Note: This table reports the average values of the slope coefficient, intercept coefficient, and mean squared error from 100 simulations as we vary I, the number of individuals in the simulated economy. The column titles denote the number of individuals in millions. The "2.5" million case corresponds to $I = 2{,}488{,}905$, which is the number of individuals who reside and work in New York City in the 2010 LODES data. In panel A, the regressand is the change in commuters with a continuum of individuals. This change does not vary across simulations. In panel B, the regressand is the simulation-specific change in commuters from a realization drawn from the model using preshock parameter values to a realization drawn from the model using post-shock parameter values.

ens out-of-sample performance (Hastie, Tibshirani, and Friedman (2009, p. 221), Belloni, Chernozhukov, and Hansen (2014, p. 30)). Consistent with the attenuation bias mentioned in Section 3.3, regressing the true change in commuters on the change predicted by the calibrated-shares procedure yields a slope of only 0.78 in the median simulation in Figure 2. Its forecast error is larger than that of the covariates-based model, which has a mean squared error that is 1% of the calibrated-shares forecast error in the median simulation. Its predictions about rents are similarly inferior: the regression slopes for rents are centered on 0.19, and the covariates-based model's mean squared error for rents is 0.003% of the calibrated-shares forecast error (see Appendix Figure A.1).

Since the only element of the data-generating process at odds with the assumptions of the calibrated-shares procedure is the finite number of individuals, these simulation results demonstrate that finite-sample bias can severely limit that procedure's predictive power. The procedure suffers from an overfitting problem when the number of individuals is small relative to the number of parameters. As shown in panel A of Table II, the calibrated-shares procedure performs better as we increase the number of individuals. Its mean squared error is inversely proportional to the number of people.[22] Only when this number is one to two orders of magnitude larger do the calibrated-shares procedure's predictions become closely correlated with outcomes. Its rent predictions become closely correlated with outcomes only when the number of individuals is two to three orders of magnitude larger (Appendix Table A.1).

In any empirical application, the researcher does not observe changes in outcomes for a continuum of individuals. The observed changes in outcomes come from finite-sample draws from the pre- and post-shock data-generating processes. We simulate this setting by drawing two labor allocation realizations, the first from the estimated model of New York City in 2010 and the second from that model after the productivity increase in one workplace tract. Because the two realizations are independent draws, commuter counts are mean reverting: absent a shock, the changes are negatively correlated with the initial

[22]The covariates-based model exhibits the same speed of convergence, as can be seen in Appendix Tables A.2 and A.3.

counts. This does not change the slope coefficient for the covariates-based model's predictions (see Appendix Figure A.2), but it severely degrades the predictive performance of the calibrated-shares procedure. For the smallest number of individuals in panel B of Table II, the calibrated-shares procedure's predicted changes are negatively correlated with the simulated changes. As explained in Appendix C.12, this particular symptom of overfitting is a case of Galton's fallacy.[23]

In the simulations described thus far, the covariates-based model is correctly specified. This neglects the calibrated-shares procedure's potential advantage over the covariates-based specification: it can capture spatial links that covariates do not. In Appendix A.3, we conduct simulations in which the covariates-based model is misspecified because $\lambda_{kn} \neq 1$ in the data-generating process. Tables A.2 and A.3 show that, for a given number of individuals, the calibrated-shares procedure's performance improves relative to the covariates-based model as the share of variation in δ_{kn} orthogonal to the included covariates increases. With sufficiently large unobserved commuting costs, the calibrated-shares procedure outperforms the covariates-based model. It achieves this by calibrating λ_{kn} for each pair of locations, which is highly data demanding. For the actual number of individuals in this empirical setting, the calibrated-shares procedure performs better when the unobserved component of commuting costs has a standard deviation at least half as large as that of the observed component (equivalent to a standard deviation of 15 minutes of transit time each way). As the number of individuals grows, the overfitting problem diminishes and the advantage of a more flexible parameterization is realized.

3.5. *Event Studies: Predicting Commuting Responses to a Local Shock*

We now examine how various approaches are able to predict changes in commuting flows using neighborhood employment booms in New York City. As in the Monte Carlo simulations, we investigate changes in commuting patterns to workplace tracts that had large increases in employment. In particular, we conduct "event studies" in which we examine the commuting flows to the 83 workplace tracts in New York City that had a 2-year increase in total employment from 2010 to 2012 of at least 400 employees and at least 12.5% from a 2010 level of at least 400 employees. We focus on local employment booms because these large changes are likely driven by workplace-specific shocks, such as new office openings or an expansion by a large employer, rather than resident-workplace-specific shocks. For example, in 2011, Tiffany & Co. moved its corporate headquarters to 260,000 square feet of office space at 200 Fifth Avenue. In late 2010, Google acquired a building of nearly 3 million square feet at 111 Eighth Avenue. These locations are two of the 83 workplace tracts we examine.[24]

To assess the predictive performance of a model parameterization, we use its baseline shares for New York City in 2010 and equations (5)–(7) to compute the increase in productivity required to match the observed 2010–2012 change in employment for the "treated" tracts.[25] Since different procedures fit the 2010 data differently, these produc-

[23] One can use the model of Section 4 to take finite-sample draws from a pre- and post-shock data-generating process without inducing mean reversion. One does so by fixing the idiosyncratic preferences $\{\nu_{kn}^i\}$ to be the same in the two draws. A Monte Carlo exercise using this assumption shows that the calibrated-shares procedure still suffers from an overfitting problem and predicts outcomes worse than the covariates-based specification (see Appendix A.4).

[24] Figure B.1 depicts the employment changes for these two tracts.

[25] Because the 83 employment booms simultaneously occurred from 2010 to 2012, we compute the vector of 83 productivity increases that matches the observed changes in employment in these 83 workplace tracts.

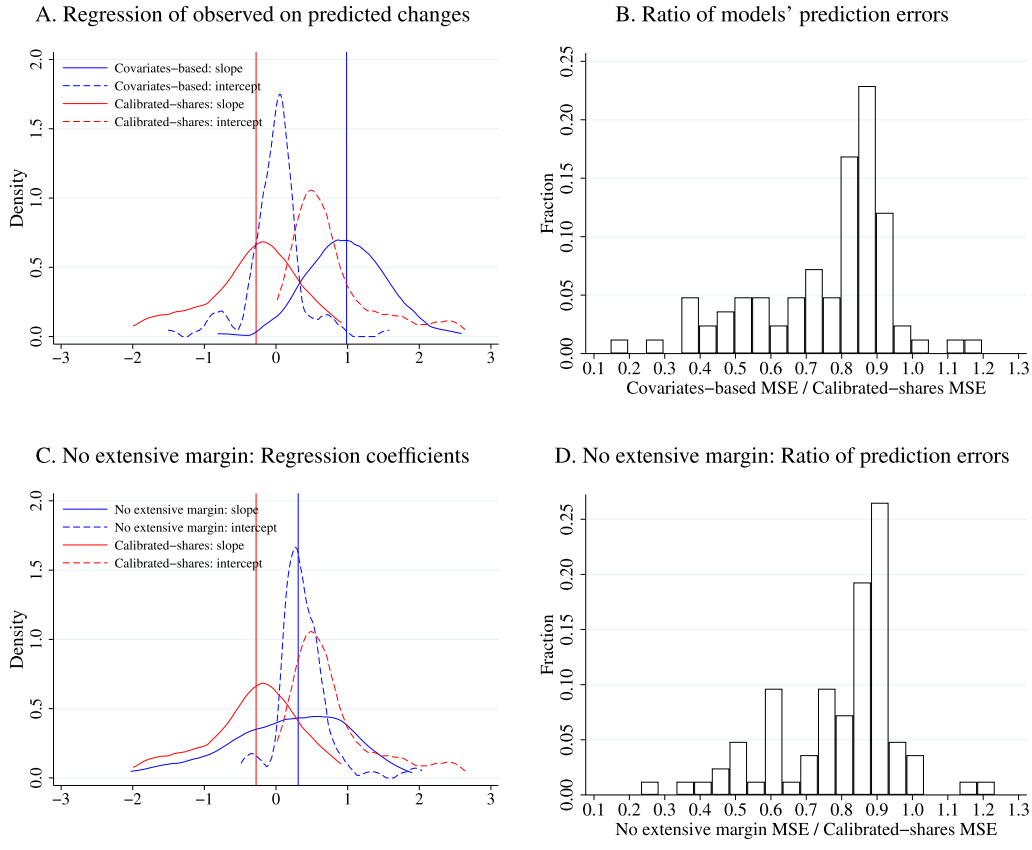

FIGURE 3.—Comparison of models' predictive performance across 83 events. NOTES: "Calibrated-shares" refers to using the observed shares in the exact hat algebra defined by equations (5)–(7). "Covariates-based" in panels A and B refers to the parameterization described in Section 3.2. "No extensive margin" in panels C and D retains only the intensive-margin predictions of the covariates-based model. This case uses the covariates-based model's predictions for residence-workplace pairs with $\ell_{kn} > 0$ in the 2010 data and predicts zero change for pairs with $\ell_{kn} = 0$ in the 2010 data. In panel A, 3 slope and 1 intercept coefficients are not depicted for the calibrated-shares procedure. In panel C, 3 slope and 1 intercept coefficients are not depicted for the calibrated-shares procedure. Vertical lines depict the medians of the slope coefficient distributions.

tivity increases could vary, but in practice they are very similar. Since the productivity increases are defined so that both procedures match the observed increase in employment, we examine their predictions for the change in bilateral commuter counts.

3.5.1. *Covariates-Based Model Outperforms Calibrated-Shares Procedure*

Figure 3 contrasts the changes in commuter flows predicted by the covariates-based model and the calibrated-shares procedure for all 83 employment booms from 2010 to 2012. For each booming tract and each counterfactual procedure, we regress the observed change in the number of residents from each residential tract who work in the treated workplace tract on the predicted change. An unbiased prediction procedure should yield a

The results obtained when computing 83 separate productivity increases, which neglects the simultaneity, are virtually identical to those in Figure 3.

slope of one and an intercept of zero. Panel A of Figure 3 depicts the distribution of these coefficients for both procedures. For our covariates-based model, the slope coefficients are roughly centered on one (median of 0.99), and the intercept coefficients are roughly centered on zero (median of 0.01).

The calibrated-shares procedure does not perform as well. Across the 83 events, the median slope coefficient is −0.28, and the median intercept coefficient is 0.60. That is, the calibrated-shares procedure's predictions are negatively correlated with observed outcomes in more than half of the events. As a result, the covariates-based approach typically has a lower forecast error. Panel B of Figure 3 contrasts the two models' MSEs for each event. The covariates-based approach has a lower MSE than the calibrated-shares procedure in 80 of the 83 events.

Looking at in-sample fit for the 2010 commuting flows would suggest a very different contrast between the two approaches: the χ^2 test statistic for in-sample fit for the calibrated-shares procedure is literally zero, while the estimated covariates-based model would be rejected by a χ^2 test (see Appendix B.6). This warns against evaluating model specifications by in-sample fit. The covariates-based specification, while not perfectly describing the observed outcomes in 2010, is much more informative about how commuting flows change with local employment booms. The calibrated-shares procedure predicts changes in commuting flows very poorly.

The covariate-based approach's superior predictive power is not due solely to the presence of zeros in the commuting matrix. To illustrate this, we produce predictions using a hybrid approach that uses the covariates-based model's predictions for residence-workplace pairs that have nonzero commuters in the 2010 data and predicts zero change for pairs where the baseline flow is zero (as in the calibrated-shares procedure). This hybrid approach, which retains only the intensive-margin predictions of the covariates-based model, still substantially outperforms the calibrated-shares procedure (Figure 3C and Figure 3D).

3.5.2. *Temporal Aggregation and Geographic Aggregation*

One may try to smooth out idiosyncratic elements of the data before fitting the model. For example, one might pool multiple years of data. In many empirical settings, particularly historical contexts, consecutive years of data are not available. In the case of LODES, annual data are available since 2002. We average commuter and wage observations for 2008–2010 before applying each approach.

Pooling 3 years of data yields a modest improvement for the calibrated-shares procedure (see Figure B.2). Its predictions are now positively correlated with observed outcomes for 69% of the events. Nonetheless, the covariates-based approach typically forecasts the changes in commuter counts much better. Its slope coefficients are closer to one (median of 0.98 vs. 0.16), and its intercept coefficients are closer to zero (0.03 vs. 0.35). The covariates-based approach applied to the pooled data has a lower MSE than the pooled calibrated-shares procedure in 74 of the 83 events.

Using larger spatial units is another potential way to address the overfitting problem. We can aggregate the 2160 residential tracts into 195 Neighborhood Tabulation Areas (NTAs) defined by the New York City Department of City Planning. We consider both aggregating the tract-level predictions up to NTA-level predictions and defining locations in the model to be NTAs.

When we aggregate the predicted changes in number of commuters to the booming workplace tracts from residential tracts up to residential NTAs, the covariates-based

model still outperforms the calibrated-shares procedure (see Figure B.2). When aggregating predictions, the covariates-based model has a lower MSE than the calibrated-shares procedure in 55 of the 83 events.

Estimating the model using NTAs reduces the number of locations by an order of magnitude and thus the number of location pairs by two orders of magnitude. 35 NTAs had 2010–2012 employment booms. As described in Appendix B.3, the covariates-based approach and the calibrated-shares procedure perform similarly well when they are applied to NTAs. The covariates-based model has a lower MSE than the calibrated-shares procedure in 18 of the 35 events. In line with our Monte Carlo results, the calibrated-shares procedure is viable when locations are defined as NTAs such that there are more than 50 individuals per pair of locations.

Using larger spatial units avoids the overfitting problem that arises when the calibrated-shares procedure is applied to granular settings. But doing so leaves one unequipped to address the many research questions that concern small spatial units such as neighborhoods. Research does not need to avoid granular settings: the covariates-based approach performs well when applied to small spatial units.

3.5.3. *Exact Hat Algebra Using Shares From an Approximated Matrix*

Temporal and geographic aggregation are simply two of many statistical tools that researchers could use to reduce idiosyncratic noise in granular settings. We now consider a noise-reduction strategy that can be applied to a single year of data from a granular setting (without geographic aggregation). The idea is to use the shares from an approximation of the commuting matrix in the exact hat algebra, which will reduce overfitting to the extent that the approximation reduces idiosyncratic noise.

In particular, we use a rank-restricted singular value decomposition (SVD) of the observed commuting matrix. The truncated SVD is widely used both to remove noise and compress data (Eldén (2007, Chapter 6)). Low-rank models are pervasive in data science and effective because matrices from many different data domains—movie preferences, text documents, medical records—are approximately low rank (Udell and Townsend (2019)). In economics, Conlon, Mortimer, and Sarkis (2023) explore using a low-rank approximation of the substitution matrix in demand estimation.

The commuting matrix $\mathbf{L} = \{\ell_{kn}\}$ of dimension $K \times N$ can be decomposed into three separate matrices $\mathbf{U}$, $\mathbf{\Sigma}$, and $\mathbf{V}$ such that $\mathbf{L} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$. This is an SVD of the commuting matrix where $\mathbf{U}$ is a $K \times K$ orthonormal matrix (the left singular vectors matrix), $\mathbf{\Sigma}$ is a non-negative $K \times N$ diagonal matrix with diagonal elements representing the singular values sorted in descending order, and $\mathbf{V}$ is an $N \times N$ orthonormal matrix (the right singular vectors matrix). The rank-$\mathfrak{r}$-restricted SVD retains only the $\mathfrak{r}$ largest singular values and sets the remaining singular values to zero. We replace negative elements of this rank-restricted matrix by zeros and rescale all values so that the total number of individuals in the resulting approximation is equal to that in the observed commuting matrix.[26]

This approximation requires the researcher to choose the rank $\mathfrak{r}$. The rank-1 approximation would underfit the data because it is merely the product of a k component and an n component, lacking a bilateral kn component. Using a full-rank matrix would mean using the observed commuting matrix. The statistics literature suggests computationally cheap rules of thumb for selecting the rank, such as visually identifying a point in a "scree

[26]This replacement means the approximated matrix may not be rank $\mathfrak{r}$. A nonnegative rank-$\mathfrak{r}$ approximation, which is costlier to compute, delivers very similar counterfactual predictions. See Appendix B.4.3.

plot" in which adjacent singular values differ little or selecting all SVD layers that explain more than a given share of the total variation (Cohen (2021, p. 496)). Alternatively, one could also use a Monte Carlo analysis to pick the optimal rank. We discuss these approaches in Appendix B.4 and use them in our application to New York City. More importantly, it is computationally trivial to vary the rank in this approximation, making it easy to conduct sensitivity analyses.

When studying neighborhood employment booms in New York City, using such an approximation of the commuting matrix yields substantially improved predictive performance over calibrating the model to the observed shares. We use an approximation derived from a rank-16 SVD, selected by a 0.5% threshold for the share of singular values. A Monte Carlo simulation using the covariates-based model as the data-generating process would have selected a rank of 18 (see Appendix B.4.2). Comparing the observed and rank-approximated commuting matrices visualized in Appendix Figure B.4 suggests that these approximations capture much of the relevant observed variation.[27] Plugging this approximation of the shares into equations (5)–(7) yields counterfactual predictions that are much closer to the predictions from the covariates-based model estimated on the full data. As shown in Figure 4, the median regression slope is 0.81, and the mean squared error is similar to that of the covariates-based approach. The regression slope coefficient from the rank-16 SVD is closer to 1.0 in 46 of the 83 events. This performance is not very sensitive to the precise choice of rank used in the approximation. Appendix Table B.1 reports the slope and MSE for a range of ranks. The MSE for the median event is similar for ranks 5 through 20.

Using a noise-reducing approximation of the matrix of spatial outcomes in exact hat algebra may be fruitful in other contexts. The approximation guards against overfitting the spatial linkages between pairs of locations even in the absence of a high-quality co-

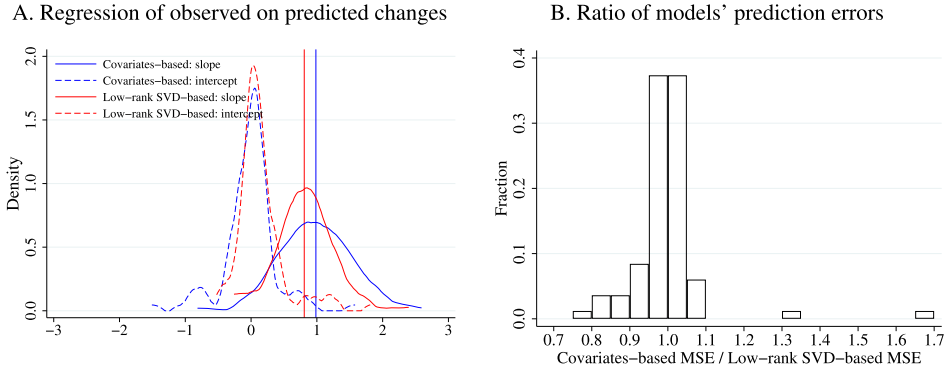

FIGURE 4.—Predictive performance across 83 events using low-rank approximation. NOTES: This figure depicts the predictive performance of using an approximation of the commuting matrix in exact hat algebra. The approximation is derived from a rank-16 singular value decomposition of the observed commuting matrix, as described in Appendix B.4. The covariates-based model is the same as in Figure 3. Panel A depicts the distributions of slope and intercept coefficients, and the vertical lines denote the median slope coefficients.

[27]For example, the rank-16 SVD predicts 815 commuters for the Columbia University outlier, quite close to the observed 827 commuters and much larger than the covariates-based model's prediction of 70. One feature of the observed matrix that these low-rank approximations do not capture is the large counts on the diagonal of the matrix: the proclivity for individuals to reside and work in the same tract. Appendix Figure B.4 shows that both an approximation derived from a rank-200 SVD and the covariates-based model do capture this pattern.

variate that predicts bilateral flows. Moreover, the computational costs of our suggested approximation are very low, as any modern linear algebra library allows one to compute the SVD of million-element matrices very quickly.

3.5.4. *Interactive Fixed Effects*

One can leverage the explanatory power of observed covariates while allowing for more flexibility by extending the covariates-based approach to have an explicitly modeled unobserved component of commuting costs. As detailed in Appendix B.5, we estimate an interactive-fixed-effects (IFE) specification that assumes the disutility component of commuting costs has an R-dimensional factor structure, $\lambda_{kn} = \exp(\psi_k' \gamma_n)$ (Bai (2009), Chen, Fernández-Val, and Weidner (2021)). For intermediate values of R, this specification is more flexible than the covariates-based model ($R = 0$) without going to the extreme of one parameter per pair of locations (the full-rank case). As with the covariates-based model, one can use the equilibrium shares of the estimated IFE model with equations (5)–(7) to compute counterfactual outcomes.

A rank-1 IFE specification offers a modest improvement over the covariates-based model in the event studies. As an encompassing specification, the IFE necessarily outperforms the covariates-based model (CBM) in terms of in-sample fit (see pseudo-R^2 in Appendix Table B.4 and χ^2 test in Appendix B.6). The event studies, summarized in Figure 5, examine its out-of-sample predictive performance. In panel A, the IFE slope coefficient for the median event is 0.85, which is farther from 1.0 than the CBM. While not centered on 1.0, its distribution of slope coefficients is narrower such that the IFE slope coefficient is closer to 1.0 in 42 of the 83 events. The IFE's forecast error is smaller than that of the CBM in the typical event: Panel B shows that its MSE is lower in 47 of the 83 events. Thus, the more flexible interactive-fixed-effects specification offers a modest improvement in predictive performance. Rank-2 and rank-3 IFE specifications perform very similarly (see Appendix Figure B.6).

3.5.5. *Sensitivity Analysis*

One difference between the approaches described above is in their use of estimated or observed wages. The covariates-based model and interactive-fixed-effects specification

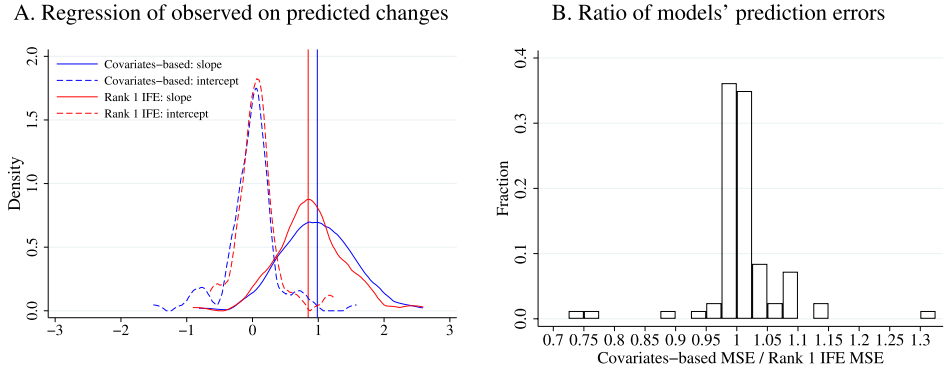

FIGURE 5.—Predictive performance of the rank-1 IFE specification across 83 events. NOTES: This figure summarizes the predictive performance of the covariates-based model and the rank-1 IFE specification. Panel A depicts the distributions of slope coefficients, and the vertical lines denote the median slope coefficients.

estimate wages using workplace fixed effects. By contrast, observed wages are used in the calibrated-shares procedure and in the exact hat algebra with the SVD-approximated shares. We have verified that using alternative wage values—such as equal wages across locations—does little to alter the predictions of the calibrated-shares procedure in the event studies.

We find similar event-study results when employing different assumptions about preferences and production. Thus far, we have assumed idiosyncratic preferences are spatially independent, as in prior research. Making residential locations in the same NTA closer substitutes in a nested-logit demand system has little effect on the predictions produced by either method (see Appendix B.7). Similarly, the changes in commuters predicted by both methods are not sensitive to the elasticity of substitution between goods σ (see Appendix B.8) nor to introducing local increasing returns at workplaces (see Appendix B.9).

3.6. *Counterfactuals in Continuum Models: Takeaways*

The predictive performance of quantitative spatial models applied to granular empirical settings is highly sensitive to the baseline equilibrium shares used in the exact hat algebra. The incumbent methods—the covariates-based approach and calibrated-shares procedure—produce sharply contrasting predictions for bilateral commuting flows because they parameterize spatial linkages very differently. Using shares from a covariates-based model predicts changes in commuting flows remarkably well. Using observed shares to calibrate millions of parameters suffers from an overfitting problem. Aggregating data across a few years does not remedy this problem. Using shares produced from a low-rank approximation of the commuting matrix alleviates the overfitting problem and produces predictions comparable to the covariates-based approach. An interactive-fixed-effects specification that extends the covariates-based approach offers modest improvements in predictive performance at greater computational cost.

When are the counterfactual predictions from the calibrated-shares procedure likely to suffer an overfitting problem? Generally, this depends on the magnitude of errors in the baseline shares and the sensitivity of counterfactual outcomes of interest to these errors. To examine the calibrated-shares procedure's finite-sample performance in a specific context, we recommend using a covariates-based model as the data-generating process in Monte Carlo simulations, as in Section 3.4.[28] The inputs necessary for such simulations are the same as those used in the calibrated-shares procedure, assuming the researcher has used a covariate to estimate the commuting elasticity. Verifying finite-sample performance is relevant beyond spatial economics, as calibrating model parameters to match observed market shares is also common in fields such as international trade and industrial organization. Whether perfectly matching the market shares is a feature or a bug depends on the context. In settings with noisy observed shares, using shares produced by data-smoothing procedures, such as low-rank approximations, or using more parsimonious models is likely to deliver better counterfactual predictions.

4. A SPATIAL MODEL WITH A FINITE NUMBER OF INDIVIDUALS

Thus far, we studied the continuum model assuming that the observed data were a finite sample drawn from that model. The overfitting problem that troubles the calibrated-shares procedure in granular settings stems from substantial idiosyncratic noise embodied

[28] Our code to simulate quantitative urban models is at https://github.com/jdingel/DingelTintelnotSEGS.

in the observed baseline shares. These idiosyncrasies matter because there are few decision makers per spatial link, even when we observe, as researchers often do, the universe of decision makers. In this section, we explore a related implication: when baseline outcomes depend on individual idiosyncrasies, so do counterfactual outcomes. In such settings, the patterns of substitution elicited by changes in economic primitives are shaped by individuals' idiosyncratic preferences. Unfortunately, the continuum model discussed in Sections 2 and 3 does not capture this. Contrary to individual idiosyncrasies affecting outcomes in granular settings, the premise of the continuum model is that they are integrated away.

We introduce a spatial model with a finite number of individuals in which equilibrium outcomes depend on individual idiosyncrasies. Its outcomes differ from the continuum model in two ways. First, because these models are not linear, the equilibrium prices in a continuum model need not equal the average equilibrium prices from a model with finite individuals. Second, counterfactual changes in quantities and prices depend on which individuals change their decisions in response to the exogenous shock, which depends on their vectors of unobserved idiosyncratic preferences. There is therefore uncertainty about realized counterfactual outcomes, as the distribution of idiosyncrasies induces a distribution of counterfactual changes.

In a model with a finite number of people, individuals' decisions affect wages and rents, raising two issues. First, do individuals internalize the effects of their own choices on local labor supplies and land demands? Second, are individuals able to enumerate the prices induced by every possible combination of others' choices? For tractability, we assume that individuals have common point-mass beliefs about wages and rents. Therefore, individuals act as price takers, as in Gabaix (2011). In our application, we assume that, given the model parameters, individuals have beliefs about wages and rents that are the equilibrium prices of the conventional continuum model. We label these beliefs "continuum-case rational expectations" because they would be rational if there were a continuum of individuals.[29] This model with a finite number of individuals is intentionally similar to the conventional continuum model: they coincide as the number of individuals becomes arbitrarily large and can be estimated similarly.

4.1. *Setup*

Our model with a finite number of individuals features the same utility function, technology, and commuting costs as in Section 2, so the set of primitives $\Upsilon \equiv \{L, \{A_n\}, \{T_k\}, \{\bar{\delta}_{kn}\}, \{\lambda_{kn}\}, \alpha, \epsilon, \sigma\}$ is the same.[30] It differs in two important respects. First, the aggregate labor endowment L is embodied in I individuals each supplying $\frac{L}{I}$ units of labor. Second, we introduce a timing assumption: individuals first choose their residence-workplace pairs, and then all markets clear given individuals' locations. This timing would not matter if there were a continuum of individuals.

[29] Similar simplifications to agents' beliefs have been made in industrial organization (e.g., oblivious equilibrium in Weintraub, Benkard, and Van Roy (2008)) and macroeconomics (e.g., Krusell and Smith (1998)). Specifically, in an oblivious equilibrium, a finite number of firms optimize assuming that the industry state always equals its long-run expected value, even though the industry state is not constant given a finite number of firms and idiosyncratic shocks. In the Krusell and Smith (1998) model, heterogeneous agents condition their choices on moments from the distribution of state variables (rather than tracking the entire distribution).

[30] Model extensions that introduce trade costs, residential amenities, local increasing returns, and production employing land are presented in Appendix C.11.

We denote the set of all idiosyncratic preferences ν^i_{kn} by $\boldsymbol{\nu}^I$. As before, these are i.i.d. draws from the standard Gumbel distribution. The superscript I indicates that the set grows with the number of individuals.

We make the following assumptions about information and expectations. All workers know the economic primitives Υ and have (common) expectations about equilibrium prices r_k and w_n. The price expectations are point-mass beliefs, such that each individual assigns 100% probability to a vector of wages $\{\tilde{w}_n\}$ and a vector of land prices $\{\tilde{r}_k\}$. Worker i knows her idiosyncratic preferences $\{\nu^i_{kn}\}$.

Decisions are made and markets clear in the following order. Based on beliefs $\{\tilde{w}_n\}$ and $\{\tilde{r}_k\}$, each worker chooses the residential location and the work location that maximize expected utility

$$\tilde{U}^i_{kn} = \underbrace{\epsilon \ln\left(\tilde{w}_n / \left(\tilde{P}^{1-\alpha} \tilde{r}_k^{\alpha} \delta_{kn}\right)\right)}_{\equiv \tilde{U}_{kn}} + \nu^i_{kn}, \tag{9}$$

where $\tilde{P} = [\sum_n (\tilde{w}_n / A_n)^{1-\sigma}]^{1/(1-\sigma)}$. After these decisions are made, individuals cannot relocate.[31] As detailed below, realized equilibrium land prices r_k and wages w_n are those that clear goods, labor, and land markets given individuals' residential and workplace locations.

Assuming point-mass beliefs about prices considerably simplifies the analysis. Otherwise, individuals would need to compute expectations using (equilibrium) probabilities for an enormous set of feasible allocations. For example, an economy with only 10 individuals and 16 residence-workplace pairs has more than 3 million possible allocations.[32] For empirically relevant magnitudes, this is infeasible. While our approach permits arbitrary point-mass beliefs about wages and rents, in our analysis below we assume that the expected prices are the equilibrium wages and rents of the continuum model with the same economic primitives Υ.

4.2. *Equilibrium*

We distinguish between a *trade equilibrium*, which clears goods and land markets taking individuals' locations as given, and a *commuting equilibrium with finitely many individuals*, in which individuals choose locations based on beliefs about the trade equilibrium that will result.

DEFINITION 4.1: Trade equilibrium. Given the labor allocation $\{\ell_{kn}\}$ and economic primitives Υ, a trade equilibrium is a set of wages $\{w_n\}$ and land prices $\{r_k\}$ satisfying equations (3) and (4).

[31] The assumption that individuals make irreversible decisions is common in static spatial models. For example, in the open-city model of Ahlfeldt, Redding, Sturm, and Wolf (2015), individuals choose to live in Berlin based on expected utility, which is equal to the reservation level of utility in the wider economy. Individuals who choose Berlin and realize utility below the citywide average cannot leave. Redding and Rossi-Hansberg (2017), Heblich, Redding, and Sturm (2020), Owens, Rossi-Hansberg, and Sarte (2020), and Brinkman and Lin (2022) make the same irreversibility assumption. In dynamic location choice models (e.g., Diamond, McQuade, and Qian (2019)), moving costs cause agents to not necessarily change locations in response to shocks.

[32] With I individuals and N^2 residence-workplace pairs, the set of possible allocations (the support of the multinomial distribution) contains $\binom{I+N^2-1}{N^2-1} = \frac{(I+N^2-1)!}{(N^2-1)!I!}$ elements. For $I = 10$ and $N = 4$, this is about 3.27×10^6.

Note that these market-clearing conditions are the same as in the continuum model.[33] Given point-mass beliefs $\{\tilde{w}_n\}$ and $\{\tilde{r}_k\}$ and idiosyncratic preferences $\boldsymbol{\nu}^I$, individuals choose the residence-workplace pair that maximizes their utility in equation (9).

DEFINITION 4.2: Commuting equilibrium with finitely many individuals. Given a number of individuals I, economic primitives Υ, idiosyncratic residence-workplace preferences $\boldsymbol{\nu}^I$, and a set of point-mass beliefs $(\{\tilde{w}_n\}, \{\tilde{r}_k\})$, a commuting equilibrium with finitely many individuals is a labor allocation $\{\ell_{kn}\}$, wages $\{w_n\}$, and land prices $\{r_k\}$ such that

- $\ell_{kn} = \frac{L}{I}\sum_{i=1}^{I} \mathbf{1}\{\tilde{U}_{kn} + \nu^i_{kn} > \tilde{U}_{k'n'} + \nu^i_{k'n'}\ \forall (k', n') \neq (k, n)\}$; and
- wages $\{w_n\}$ and land prices $\{r_k\}$ are a trade equilibrium given the labor allocation $\{\ell_{kn}\}$.

We now define a set of price beliefs that are the equilibrium prices of the continuum model from Section 2 with the same economic primitives.[34]

DEFINITION 4.3: Rational expectations for the continuum case. Given economic primitives Υ, $\{\tilde{w}_n\}$ and $\{\tilde{r}_k\}$ are "continuum-case rational expectations" if $\{w_n\} = \{\tilde{w}_n\}$ and $\{r_k\} = \{\tilde{r}_k\}$ constitute an equilibrium of the continuum model defined in Section 2.2 for the same economic primitives Υ.

Distinguishing between the number of individuals I and aggregate labor L allows us to study locational decisions as $I \to \infty$ without changing aggregate labor supply. As Appendix C.9 shows, the labor allocation in the $I \to \infty$ limit is the same as in the continuum model in equation (2). Note that given idiosyncratic preferences and point-mass beliefs, the commuting equilibrium with finitely many individuals is unique.

4.3. *Estimation of Economic Primitives*

One can estimate a covariates-based specification of this model in the same way we estimated its continuum-model counterpart in Section 3.2. We take α, σ, $\bar{\delta}_{kn}$ as given, impose $\lambda_{kn} = 1\ \forall kn$, and assume that ν^i_{kn} are i.i.d. draws from a standard Gumbel distribution.[35] The resulting log likelihood function is

$$\mathcal{L} \equiv \sum_{k,n} \ell_{kn} \ln\left[\mathbb{P}\left(\tilde{U}^i_{kn} > \tilde{U}^i_{k'n'}\ \forall k'n' \neq kn\right)\right] = \sum_{k,n} \ell_{kn} \ln\left[\frac{\tilde{w}_n^{\epsilon}\left(\tilde{r}_k^{\alpha}\bar{\delta}_{kn}\right)^{-\epsilon}}{\sum_{k',n'} \tilde{w}_{n'}^{\epsilon}\left(\tilde{r}_{k'}^{\alpha}\bar{\delta}_{k'n'}\right)^{-\epsilon}}\right]. \tag{10}$$

This is equation (8) with rent and wage beliefs instead of equilibrium rents and wages. The maximum-likelihood estimate of the commuting elasticity ϵ is identical to that in the covariates-based continuum model (in column 1 of Table I). The remaining economic

[33]The trade equilibrium is unique. Given the labor allocation, there is a unique set of relative wages satisfying equation (3), as shown in Appendix C.8. Equation (4) can be rewritten as $r_k = \frac{\alpha}{T_k}\sum_n \frac{\ell_{kn}}{\bar{\delta}_{kn}} w_n$, so there is a unique set of land prices associated with that wage vector.

[34]The wage and rent beliefs in Definition 4.3 are unique if the parametric condition for uniqueness in the continuum model in Section 2.2 is satisfied.

[35]The assumption that $\lambda_{kn} = 1$ can be relaxed by allowing the unobserved disutility component of commuting costs to have a factor structure, $\lambda_{kn} = \exp(\psi_k' \gamma_n)$, as in Section 3.5.4 and Appendix B.5. If individuals have continuum-case rational expectations, the economic primitives of this specification of the model with a finite number of individuals match those of the interactive fixed-effects specification of the continuum model.

primitives $\{T_k\}$ and $\{A_n\}$ can be obtained from the estimated rent and wage beliefs: the residence and workplace fixed effects are proportional to $\tilde{r}_k^{-\alpha\epsilon}$ and $\tilde{w}_n^{\epsilon}$, respectively. If $\{\tilde{r}_k\}$ and $\{\tilde{w}_n\}$ are continuum-case rational expectations, plugging these estimated beliefs into equation (2) yields the continuum-case labor allocation. Given α, σ, $\{\tilde{r}_k\}$, $\{\tilde{w}_n\}$, $\{\delta_{kn}\}$, and that labor allocation, equations (3) and (4) can be solved to obtain $\{T_k\}$ and $\{A_n\}$. These estimated land and productivity parameters are identical to those of the covariates-based continuum model.

4.4. *Price Dispersion and Ex Post Regret*

The realized equilibrium rents and wages vary with individual idiosyncrasies, so they generically differ from the point-mass beliefs about rents and wages that govern individuals' choices of residences and workplaces. How often would individuals choose a different residence-workplace pair at the realized equilibrium prices if they were able? In our New York City setting, ex post regret is rare: 96% of individual have idiosyncratic preferences such that they would not want to switch. See Appendix E for details of these results. Given the return to making the model tractable and computationally feasible, we judge this magnitude to be modest.

4.5. *Contrasts With the Continuum Model*

With a finite number of individuals, individual idiosyncrasies affect equilibrium outcomes. Given economic primitives Υ, this model produces a distribution of outcomes associated with the distribution of idiosyncratic preference shocks $\boldsymbol{\nu}^I$, whereas the continuum model delivers deterministic outcomes. Such a contrast has been emphasized by studies of how aggregate outcomes depend on firm-level productivity shocks (e.g., Gabaix (2011), di Giovanni, Levchenko, and Mejean (2014), Eaton, Kortum, and Sotelo (2013), Gaubert and Itskhoki (2021)).

Given the same primitives Υ, the expected equilibrium labor allocation of this model equals that of the continuum model. This is because the multinomial distribution of residence-workplace outcomes $\{\ell_{kn}\}$ defined by the probabilities in equation (10) has expected values equal to the quantities in equation (2). By contrast, the model's equilibrium prices solve a system of nonlinear equations, so their expected values need not equal the continuum-case prices. For our baseline parameter values, these differences are small (Appendix E.1).

We focus on the implications of these idiosyncrasies for counterfactual outcomes. We compute the change in outcomes induced by a change in economic primitives from Υ to Υ' given idiosyncratic preferences $\boldsymbol{\nu}^I$.[36] In special cases, one can analytically characterize the uncertainty stemming from idiosyncrasies. In Appendix C.10, we describe a two-workplace economy. Let p denote the share of individuals expected to switch workplaces in response to a productivity change in the continuum model. In the model with a finite number of individuals, the variance of the share of individuals who switch workplaces is $\frac{p(1-p)}{I}$. In the general case, simulating realizations of $\boldsymbol{\nu}^I$ yields a distribution of counterfactual changes. The dispersion in counterfactual changes represents uncertainty stemming from individual idiosyncrasies. Although the analytical result describes a special case, it will turn out to approximate the uncertainty about the share of individuals switching workplaces in our Amazon HQ2 application well.

[36]This is akin to the computation of counterfactual changes holding a realization of firms' idiosyncratic unit costs fixed in Eaton, Kortum, and Sotelo (2013, Section 6.1).

5. APPLICATION TO AMAZON HQ2 IN LONG ISLAND CITY

We now examine the predictions of quantitative spatial models for the economic consequences of Amazon's aborted second headquarters (HQ2) in Long Island City, a controversial and widely discussed proposal. In November 2018, the company announced it would hire more than 25,000 employees in 4 million square feet of office space.[37] However, Amazon scrapped the project in February 2019 after facing a fierce backlash from local politicians and community members concerned about corporate subsidies and gentrification.

This large counterfactual change to a single workplace allows us to illustrate how choices of methods affect predictions. Section 5.1 shows how the geographic incidence of a productivity shock on local rents is governed by the parameterization of spatial linkages. Section 5.2 uses the model with a finite number of individuals to quantify the uncertainty about these counterfactual outcomes. Individual idiosyncrasies make neighborhood-level counterfactual changes considerably uncertain, even when examining a productivity shock shifting 1% of employment to a single tract.

5.1. *Geographic Incidence of Amazon HQ2*

We analyze the consequences of a productivity increase in the Long Island City tract that would cause employment to rise by 25,000 workers from its 2010 level. Figure 6A depicts the distribution of residences among workers employed in that tract in 2010. To contrast the predictions of the continuum model when using the covariates-based model and the calibrated-shares procedure, we find the productivity increases $\hat{A}_{n^*}^{\text{CBM}}$ and $\hat{A}_{n^*}^{\text{CSP}}$ that raise employment by 25,000 in the Amazon workplace tract n^*, holding all other economic primitives fixed.[38] Appendix F.1 reports additional results from using alternative values of the labor demand elasticity σ, from using a nested-logit demand system, and from the SVD-based and IFE specifications of the continuum model.

The counterfactual changes in residential populations are sensitive to the method used: while the covariates-based approach predicts small changes in all tracts, the calibrated-shares procedure predicts more varied and less spatially correlated changes. The covariates-based model's predicted changes in the number of residents in each tract, shown in Figure 6B, are modest in size and closely related to transit times to Long Island City. The calibrated-shares procedure's predicted changes, shown in Figure 6C, are very large for some residential tracts. The calibrated-shares procedure predicts increases in residents in neighborhoods scattered across the city, but little increase in residents in adjacent neighborhoods with short commutes to Long Island City. The calibrated-shares procedure's predictions are closely tied to the initial numbers of residents working in the treated tract: the spatial patterns in Figure 6A and Figure 6C are remarkably similar.

The incidence of Amazon HQ2 on residential rents predicted by the two methods differs considerably. The greater variation in changes in the number of residents predicted

[37]Amazon had requested proposals for HQ2 locations from cities and states across North America. In November 2018, it announced two winners, saying that it would hire more than 25,000 employees each in Long Island City in New York and Arlington in Virginia (Amazon (2018)). More details can be found in the Memorandum of Understanding between New York State, New York City and Amazon (https://d39w7f4ix9f5s9.cloudfront.net/4d/db/a54a9d6c4312bb171598d0b2134c/new-york-agreement.pdf). Berkes and Gaetani (2020) introduce a model that features productivity spillovers and skill heterogeneity to analyze Amazon HQ2. Following Berkes and Gaetani (2020), the "treated" tract in our counterfactual scenario is 36081000700.

[38]That is, we find the value of $\hat{A}_{n^*}^{\text{CBM}}$ such that the solution to equations (5)–(7) yields $\hat{\ell}_{kn}^{\text{CBM}}$ so that $\sum_k (\hat{\ell}_{kn^*}^{\text{CBM}} - 1)\ell_{kn^*}^{\text{CBM}} = 25{,}000$ and similarly for the CSP case.

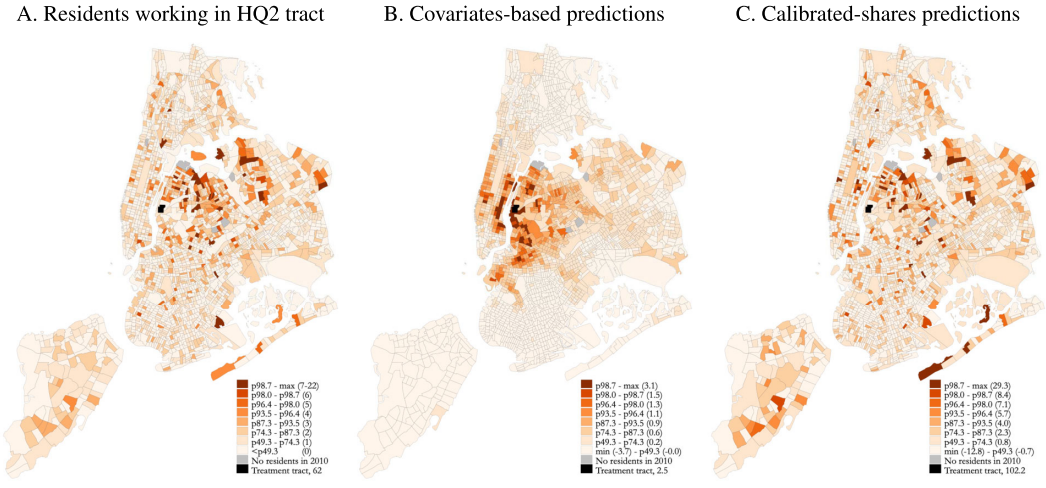

FIGURE 6.—Amazon HQ2 counterfactual change in residents. NOTES: Panel A depicts the number of residents in each tract who work in the treated Long Island City tract in the 2010 LODES data. The legend of panel A reports the percentiles corresponding to the integer number of residents, excluding the treatment tract. Panels B and C depict changes in the number of residents predicted by the covariates-based model and calibrated-shares procedure, respectively. The legend percentile cutoffs in panels B and C correspond to those in panel A. The covariates-based predictions for residents describe both the continuum model and the model with finite individuals because the change in the former equals the expected change in the latter.

by the calibrated-shares procedure is accompanied by greater variation in the predicted changes in real rents, $\hat{r}_k/\hat{P}$. The calibrated-shares procedure predicts rent increases of 1.6% to 10.4% for the top decile of tracts (Figure 7C). It suggests that in 12 census tracts, real rents would increase by at least 5%. By contrast, the covariates-based model predicts

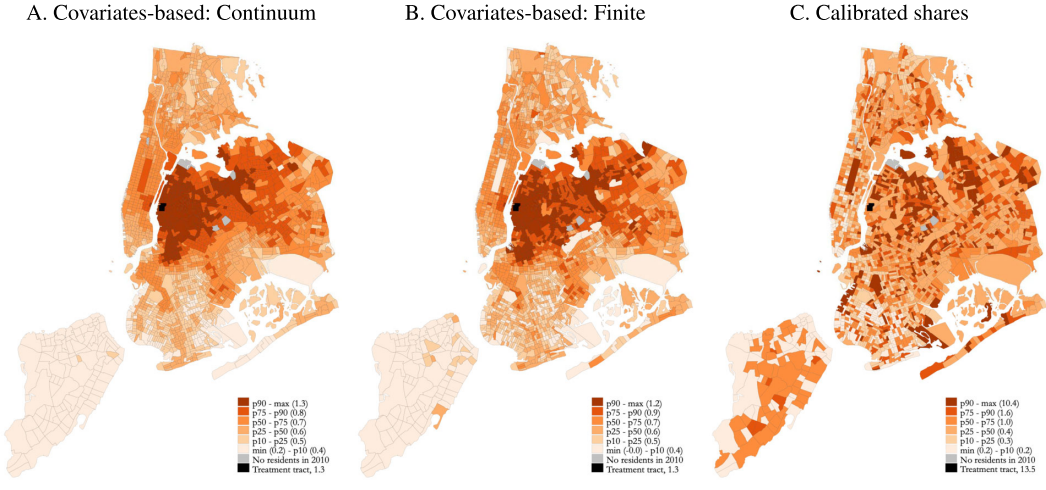

FIGURE 7.—Predicted changes in rents. NOTES: These maps depict percent changes in rents, $(\hat{r}_k/\hat{P} - 1) \times 100$. Panel A depicts the change in the covariates-based continuum model. Panel B depicts the mean change across 100 simulations of the model with finite individuals. Panel C depicts each tract's predicted rent change using the calibrated-shares procedure.

more modest increases of 0.8% to 1.3% for the top decile of affected tracts (Figures 7A and 7B).[39] The geographic incidence differs notably. Figure 7 shows that the calibrated-shares procedure predicts much larger rent increases in neighborhoods far from Amazon HQ2, such as Staten Island in the southwest corner of the map.[40] By contrast, the political opposition to Amazon HQ2 based on concerns about gentrification largely concerned housing prices in nearby constituencies (Goodman (2019)).

Unlike the changes in residents and rents, the changes in workers, wages, and welfare predicted by the two methods are quite similar. Appendix Figures F.1 and F.2 show the changes in workers and wages. These predictions are similar because the geographic pattern of employment changes in response to a productivity shock is closely tied to the initial number of workers employed in each tract, which the two methods agree upon. Similar to the special case presented in Appendix C.4 and summarized in Section 3.3, residential responses to a workplace productivity shock largely depend on the baseline share of a single residence-workplace pair, which differs between the two methods, whereas employment responses depend on the workplace's share of total employment. The citywide welfare consequence of the productivity shock, derived in Appendix C.7, is also very similar across methods. This occurs because the welfare effect of this single-tract employment boom is largely invariant to spatial linkages: computing counterfactual changes using a parameterization with no spatial linkages, $\delta_{kn} = 1\ \forall kn$, also yields a very similar welfare consequence.[41]

5.2. *Uncertainty About Counterfactual Predictions Induced by Idiosyncrasies*

Counterfactual predictions come with various forms of uncertainty. For example, to capture uncertainty about parameter values, researchers can bootstrap the counterfactual predictions of a continuum model. Beyond parameter uncertainty, when the number of individuals is finite, there is a distribution of equilibrium outcomes associated with the distribution of individuals' idiosyncratic preferences. To capture the uncertainty stemming from individual idiosyncrasies that accompanies a change in economic primitives, we simulate the counterfactual change from Υ to Υ' for 100 realizations of idiosyncratic preferences $\boldsymbol{\nu}^I$ and construct a 90% confidence interval by reporting the 5$^{\text{th}}$ and 95$^{\text{th}}$ percentiles of the distributions of changes.

The predicted tract-level changes in residents and rents are small relative to the idiosyncratic variation. Figure 8A plots the 90% confidence interval for each tract's predicted change in residents against its mean predicted change. Every tract's 90% confidence interval for the change in residents has a positive upper bound and a negative lower bound. In this sense, none of the predicted changes in residents are distinguishable from zero. Figure 8C depicts the results for rents. For about 15% of residential tracts, the predicted changes in real rents are distinguishable from zero.

Since the predicted changes in the number of workers for some tracts are larger, many of them are meaningfully large relative to the idiosyncratic variation. As the number of

[39]The covariates-based continuum model and model with a finite number of individuals predict very similar rent changes (compare panels A and B).

[40]Appendix Figure F.3 plots rent changes against distance to the treated tract for both methods.

[41]This finding does not imply that the two methods will always yield similar welfare conclusions. For counterfactual scenarios where spatial linkages matter for citywide welfare consequences, they may differ. For example, the citywide welfare benefit of faster commuting between two tracts that had zero commuters in the 2010 data would be zero under the calibrated-shares procedure but positive when using the baseline shares from the covariates-based parameterization.

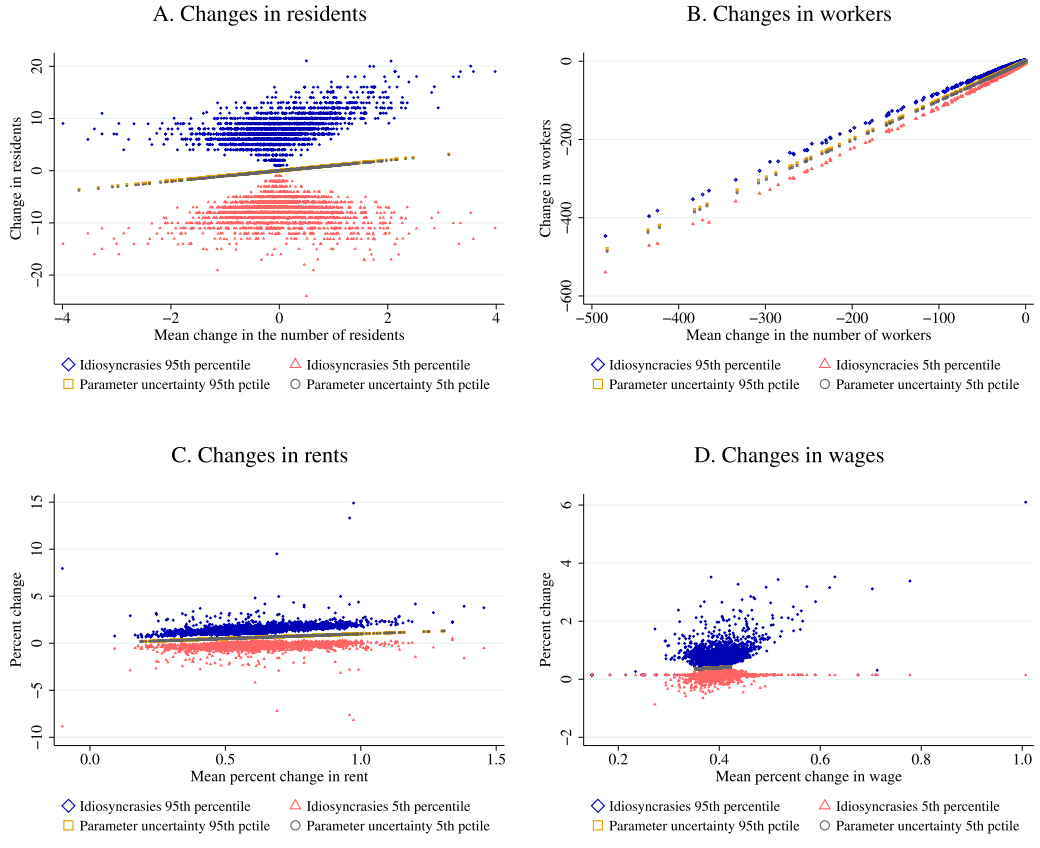

FIGURE 8.—Variation in counterfactual changes: Idiosyncrasies vs. parameter uncertainty. NOTES: These plots depict variation in counterfactual changes induced by individual idiosyncrasies in the model with a finite number of individuals alongside variation attributable to parameter uncertainty about $\{A_n\}$, $\{T_k\}$, and ϵ in the covariates-based continuum model. We summarize the former by plotting the 5th and 95th percentiles of predicted changes in quantities and prices across 100 simulations of the model with a finite number of individuals. In panels A, B, C, and D, the horizontal axes correspond to the values depicted in Figures 6B, F.1B, 7B, and F.2B, respectively. Panels B and D exclude the Amazon HQ2 workplace tract. In panel A, the 90% confidence interval for the change in residents includes zero for all tracts. In panel B, one outlier with an employment decline of 1468 is not depicted, and the 90% confidence interval for the change in employment includes zero for 421 of the 2142 non-Amazon workplaces. The 90% confidence interval for the change in employment for the Amazon workplace tract is 24,728 to 25,259. In panel C, real rent changes are $(\hat{r}_k/\hat{P}-1)\times 100$, and 316 out of 2160 origin tracts have a positive change in rents at the 5th percentile. In panel D, the real wage changes are $(\hat{w}_n/\hat{P}-1)\times 100$, and 1962 out of 2143 destination tracts have a positive change in wages at the 5th percentile. We quantify parameter uncertainty by producing 100 counterfactual predictions by sampling with replacement 2,488,905 individuals from the observed commuting matrix 100 times. For each of these 100 samples, we estimate the covariates-based continuum model and compute its counterfactual predictions. The plots depict the 5th and 95th percentiles of predicted changes in quantities and prices across these 100 bootstrap samples. The variation attributable to parameter uncertainty is considerably smaller than the variation attributable to individual idiosyncrasies: the bootstrap 5th and 95th percentiles are almost indistinguishable.

workers varies across tracts much more than the number of residents, so too do the widths of the confidence intervals. Figure 8B depicts these confidence intervals for the changes in the number of workers. For hundreds of workplace tracts, the 90% confidence interval

includes both negative and positive values.[42] For tracts predicted to lose many workers, the predicted change is sufficiently large relative to the uncertainty that even the 95th percentile of changes is negative. It turns out that a binomial distribution with I draws and p equal to the expected change in the share of employment (the analytical result for uncertainty for the two-workplace case mentioned in Section 4.5) is a good approximation of the 90% confidence intervals depicted in Figure 8B.[43] For the vast majority of workplace tracts, the uncertainty over real wages is sufficiently small that the 5th-percentile change for the real wage is greater than zero (see Figure 8D).[44]

The uncertainty stemming from individual idiosyncrasies in the model with finitely many individuals is larger than the variation in counterfactual predictions associated with parameter uncertainty in the continuum model in this setting. To compute the latter, we sample with replacement 2.5 million individuals from the observed commuting matrix 100 times. For each of these 100 samples, we estimate the covariates-based continuum model and compute its counterfactual predictions. Figure 8 shows that the variation in counterfactual changes induced by parameter uncertainty is much smaller than the variation induced by different realizations of idiosyncrasies.

We also investigate the uncertainty stemming from time-varying fundamentals by redoing the counterfactual analysis using 2008 data (instead of 2010 data). When using the covariates-based parameterization, the counterfactual predictions are not very sensitive to the choice of baseline year: all the real rent changes and more than 95% of the real wage changes produced by the 2008 estimated model are inside the 90% confidence intervals of the predictions from the 2010 estimated model depicted in Figure 8.[45]

As we move from tract-level outcomes to examining larger geographic units, the uncertainty about counterfactual outcomes induced by individual idiosyncrasies diminishes but remains substantial. If we aggregate tract-level changes into 20 bins based on distance to the Amazon HQ2 tract, the mean predicted changes in residents and rents decline with distance to the Amazon HQ2 tract. For all but the first and twentieth ventiles, the 90% confidence interval for the change in residents includes zero (see Figure F.7). For every ventile, the 90% confidence interval for the rent change includes the citywide change. The uncertainty about the change in real aggregate output stemming from individual idiosyncrasies is negligible.[46] When we apply the model with a finite number of individuals to 195 Neighborhood Tabulation Areas rather than tracts, the 90% confidence interval for the change in residents includes zero for 193 of the 195 NTAs. The confidence intervals for the predicted change in real rent include the citywide change for 132 NTAs (see Appendix Figure F.9).

The uncertainty about counterfactual tract-level changes in residents and rents is similar when residential locations are more substitutable and when labor demand is more

[42]While workplaces are generically substitutes for one another, the Amazon expansion alters the pattern of residential rents, which can cause some individuals to choose a new residence-workplace pair that includes a non-Amazon workplace. An (idiosyncratic) increase in labor supply can even reduce real wages in such workplaces, as shown in Figure 8D.

[43]For example, for $I = 2{,}488{,}905$ and the workplace with an expected decrease of 300 employees ($p \times I = 300$), the binomial distribution's 90% confidence interval for the employment decrease, $pI \pm 1.64\sqrt{p(1-p)I} \approx [272, 328]$, aligns closely with that in Figure 8B.

[44]If labor demand were perfectly elastic $\sigma = \infty$, equilibrium wages would be independent of residual labor supply and would exhibit no uncertainty: $w_n \propto A_n$.

[45]By contrast, the calibrated-shares procedure's predictions are much more sensitive to the use of 2008 versus 2010 data. The correlation between the real rent changes shown in Figure 7C and those obtained using the 2008 data is only 0.23 (0.74 for real wages).

[46]The 90% confidence interval for the increase in aggregate output is 0.41% to 0.43%.

elastic. Uncertainty is similar when we employ a nested-logit specification in which residential locations in the same NTA are closer substitutes (see Appendix Figure F.10). When labor demand is perfectly elastic ($\sigma = \infty$), wages are invariant to individual idiosyncrasies. Uncertainty about rent changes is similar to the uncertainty in the baseline case ($\sigma = 4$), though the mean rent increase is smaller so more tracts have confidence intervals that include zero (see Appendix Figure F.11).

The predicted consequences of the most closely followed corporate headquarters decision in recent memory for local economic outcomes are subject to considerable uncertainty stemming from individual idiosyncrasies. This notable uncertainty would be missed by the conventional continuum approach. Given the magnitude of uncertainty stemming from individual idiosyncrasies that we document in this case, we believe this uncertainty is likely relevant in many empirical applications that consider the geographic distribution of economic outcomes.

When should a researcher quantify the uncertainty about counterfactual predictions attributable to individual idiosyncrasies? Our results suggest the geographic scale of the predicted changes is key: individual idiosyncrasies induce little uncertainty about aggregate outcomes. When studying spatially fine outcomes, the answer depends on the empirical setting, such as the size and nature of the counterfactual shock and the expected number of individuals choosing a pair of locations. To assess whether they can rule out uncertainty attributable to individual idiosyncrasies, researchers can repeatedly draw finite labor allocations from the continuum model at counterfactual parameters Υ' to compute a distribution of differences between finite-sample counterfactual outcomes and the expected outcome at baseline parameters Υ. This distribution of differences is more dispersed than the distribution of counterfactual changes from the model with a finite number of individuals that fixes the vector of idiosyncratic preferences $\boldsymbol{\nu}^I$.[47] If the dispersion of the former is small, then the uncertainty stemming from individual idiosyncrasies is small. For example, if the distribution of continuum-model differences does not include zero, then the confidence interval from the model with finite individuals excludes zero. When uncertainty attributable to individual idiosyncrasies cannot be ruled out, computing and reporting this uncertainty will produce more informative predictions about local economic outcomes.

6. CONCLUSION

Economists increasingly use spatially fine data and quantitative spatial models to compute counterfactual general-equilibrium outcomes. The smaller the number of individuals behind each economic outcome reported in these granular settings, the less compelling the conventional modeling assumption that there is a continuum of individuals. These applied general-equilibrium models need to reliably predict economic changes if policymakers are to use them to inform their decisions (Kehoe (2005), Bryan, Glaeser, and Tsivanidis (2020), Adão, Arkolakis, and Esposito (2022)).

We document that some conventional methods perform poorly when applied to granular settings. Calibrating pair-specific parameters when there are many pairs of locations

[47]This distribution of differences is more dispersed than the distribution of counterfactual changes because an idiosyncratically high outcome at counterfactual parameters Υ' would also be higher at baseline parameters Υ given the same idiosyncratic preferences $\boldsymbol{\nu}^I$. This can be seen by contrasting Figure 9 in Dingel and Tintelnot (2021), which depicted such differences, to Figure 8 above. Sampling from the continuum model is computationally cheaper than computing counterfactual outcomes holding idiosyncratic preferences fixed in the model with a finite number of individuals.

relative to the number of decision makers fits the model to the idiosyncratic component of individual decisions. We recommend that researchers use Monte Carlo simulations to verify the finite-sample behavior of their procedures for producing counterfactual predictions in their empirical setting. We show how to use low-rank matrices to capture unobserved spatial linkages without overfitting when computing counterfactual outcomes using exact hat algebra.

Counterfactual outcomes for particular locations vary with individual idiosyncrasies. We introduce a model with a finite number of individuals in which equilibrium outcomes depend in part on the idiosyncratic component of individuals' choices. This estimated model yields a distribution of equilibrium outcomes when computing counterfactual scenarios. Evaluating the counterfactual consequences of Amazon's proposed second headquarters in Long Island City, we find that there is considerable uncertainty about most of the predicted changes in local outcomes. Since all empirical settings feature a finite number of individuals, the continuum assumption has been made in the interest of modeling convenience, not realism. The model with a finite number of individuals is tractable and features a covariates-based specification that can be estimated using the same data as its continuum counterpart. Thus, it can be applied in the same settings in which economists have thus far assumed a continuum.

APPENDIX A: MONTE CARLO SIMULATIONS

We conduct a Monte Carlo simulation to examine the finite-data performance of the covariates-based approach and the calibrated-shares procedure. In these simulations, the data-generating process is our estimated covariates-based model of New York City in 2010. We impose a counterfactual 9% increase in productivity ($\hat{A}_n = 1.09$) for the "treated" tract containing 200 Fifth Avenue, which generates an increase in employment matching the observed change from 2010 to 2012.[48] There are no changes in the productivity of other workplace tracts ($\hat{A}_n = 1$), no changes in land endowments ($\hat{T}_k = 1\ \forall k$), and no changes in commuting costs ($\hat{\hat{\delta}}_{kn} = 1\ \forall k, n$). In the limiting case, as $I \to \infty$, the calibrated-shares procedure would perfectly describe the changes in commuting flows associated with this productivity increase. Thus, any predictive failure when drawing finite-sample realizations from this data-generating process are due to problems stemming from the finite number of individuals.

A.1. *Monte Carlo Exercise: Continuum Changes*

In the continuum model with parameter values estimated in Section 3.2, the "true" change in outcomes is given by combining the relative changes in exogenous parameters and the baseline shares using equations (5)–(7).

We compute predicted changes using the covariates-based model and the calibrated-shares procedures in 100 simulations of finite-sample data. In each of the 100 simulations, we implement the following steps:

1. To generate "observed" baseline data, draw a finite-sample realization of the labor allocation from the parameterized continuum model.

[48] In 2011, Tiffany & Co. moved its corporate headquarters to 260,000 square feet of office space at 200 Fifth Avenue.

2. Estimate ϵ and other parameters of the covariates-based model using the "observed" realization. Compute baseline shares ℓ_{kn} and y_{kn} for the covariates-based model using these parameter estimates. Compute baseline shares ℓ_{kn} and $y_{kn} = \ell_{kn} w_n / \bar{\delta}_{kn}$ for the calibrated-shares procedure using the finite-sample realization of ℓ_{kn} drawn in the previous step and the "true" values of w_n and $\bar{\delta}_{kn}$ used in the data-generating process.
3. Using equations (5)–(7), compute the increase in productivity required to match the true change in employment for the "treated" tract for both the covariates-based approach and the calibrated-shares procedure.
4. Compute the predicted change in commuter counts $(\ell_{kn}(A') - \ell_{kn}(A))$ and rents $(\hat{r}_k / \hat{P} - 1)$ using both the covariates-based approach and the calibrated-shares procedure.
5. Regress the true change in commuters destined for the treated tract on the changes predicted by the covariates-based approach. Regress the true change in commuters destined for the treated tract on the changes predicted by the calibrated-shares procedure. Regress the true change in real rents on the changes predicted by the covariates-based approach and calibrated-shares procedure, respectively. Compute the mean squared error of each set of predictions.

The results of these 100 simulations are reported in Figure 2, Panel A of Table II, Figure A.1, and Table A.1. Figure A.1 shows that the calibrated-shares procedure predicts changes in rents poorly with the number of individuals in the data. Table A.1 shows that the slope coefficient becomes close to one when the number of individuals is two or three orders of magnitude larger.

A.2. *Monte Carlo Exercise: Finite-Sample Changes*

In any empirical application, the researcher does not observe changes in outcomes for a continuum of individuals. The observed changes in outcomes come from finite-sample draws from the pre- and post-shock data-generating processes. We simulate this setting by

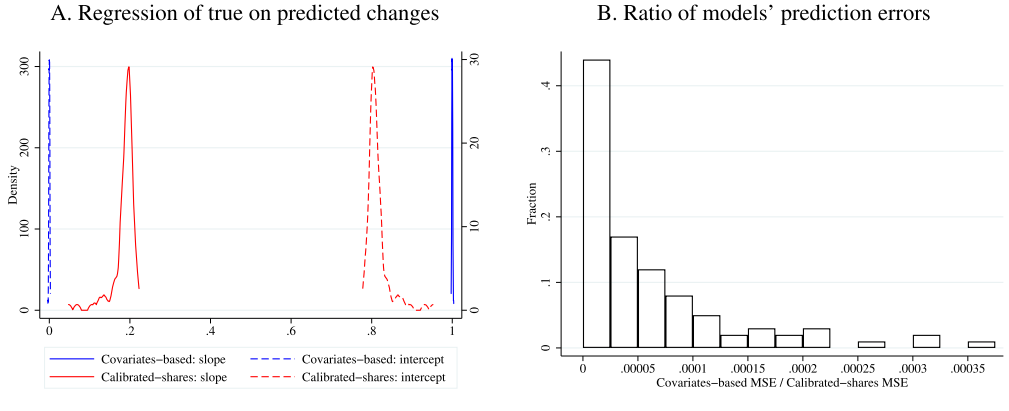

FIGURE A.1.—Monte Carlo: Calibrated-shares procedure predict rents poorly. NOTES: This figure describes predicted rent changes from the same simulations that produced the predicted commuter changes in Figure 2 and Panel A of Table II. This figure depicts the regression coefficients and mean squared errors from 100 simulations in which $I = 2{,}488{,}905$, the data-generating process is the estimated covariates-based model of New York City in 2010, and the counterfactual change is a 9% increase in productivity in one workplace tract. The predicted change in real rents is compared to the continuum-model change in real rents.

TABLE A.1

CALIBRATED-SHARES PROCEDURE'S FINITE-SAMPLE PERFORMANCE PREDICTING RENT CHANGES.

I	2.5	5	12.5	25	50	125	250	2560
Calibrated-shares: slope	0.187	0.312	0.535	0.695	0.822	0.918	0.959	0.996
Calibrated-shares: intercept	0.814	0.689	0.465	0.306	0.178	0.082	0.041	0.004
Calibrated-shares: MSE	444.852	219.325	85.426	42.944	21.375	8.708	4.310	0.420

Note: This table describes predicted rent changes from the same simulations that produced the predicted commuter changes in Panel A of Table II.

drawing two labor-allocation realizations for each simulation, the first from the estimated model of New York City in 2010 and the second from that model with 9% higher productivity for the treated workplace tract. For each simulation, we execute six steps. The first four are identical to those described in the previous section, Appendix A.1. The last two are

5. Generate "observed" post-shock data by drawing a finite-sample realization of the labor allocation from the continuum model at the counterfactual productivity vector, which has 9% higher productivity for the treated workplace tract. Compute "observed" changes in commuter counts by subtracting the step-1 baseline realization of commuter counts from these realized post-shock commuter counts.
6. Regress the "observed" changes in commuter counts destined for the treated workplace tract on the changes predicted by the calibrated-shares procedure. Regress the "observed" changes in commuter counts destined for the treated workplace tract on the changes predicted by the covariates-based approach.

The results of these simulations are reported in Panel B of Table II and Figure A.2. The additional finite-sample noise raises the MSEs for both methods relative to the simulations described in Section A.1.

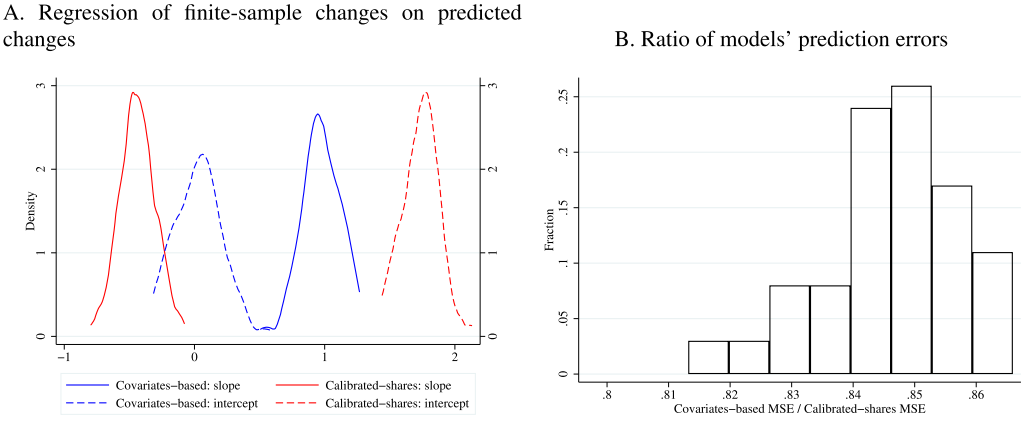

FIGURE A.2.—Calibrated-shares procedure overfits in finite-sample simulations. NOTES: This figure depicts the regression coefficients and MSEs from 100 simulations in which $I = 2{,}488{,}905$, the data-generating process is a finite-sample realization from the estimated covariates-based model of New York City in 2010, and a finite-sample realization from that model with a 9% increase in productivity in one workplace tract. "Covariates-based" means applying the parameterization described in Section 3.2 to the simulated data. "Calibrated-shares" means using the observed (simulated) shares in the exact hat algebra defined by equations (5)–(7).

A.3. *Monte Carlo Simulations With Unobserved Disutility,* $\lambda_{kn} \neq 1$

Next, we consider a data-generating process in which commuting costs have both an observed and unobserved component in order to examine the relative predictive performance of the covariates-based model and calibrated-shares procedure in granular settings. The commuting costs are now determined by

$$\ln \delta_{kn} = \ln \bar{\delta}_{kn} + \ln \lambda_{kn},$$

$$\ln \lambda_{kn} \overset{\text{iid}}{\sim} \mathcal{N}\left(0, \Lambda^2 \times \text{Var}(\ln \bar{\delta}_{kn})\right), \quad \Lambda \in \mathbb{R}^+.$$

The parameter Λ scales the standard deviation of the unobserved component of commuting costs, stated relative to the standard deviation of the observed component. The simulations shown in Table II are the case in which $\Lambda = 0$, so that $\lambda_{kn} = 1 \ \forall k, n$. When $\Lambda > 0$, the covariates-based model omits an unobserved component of commuting costs that the calibrated-shares procedure may, in principle, capture. Since $\ln \bar{\delta}_{kn}$ and $\ln \lambda_{kn}$ are orthogonal by assumption, the covariates-based model's predictions will be unbiased, but its forecast errors will increase with the magnitude of unobserved variation Λ. For $\Lambda = 0.1, 0.25, 0.5$, and 1.0, we compute 100 simulations using 100 realizations of the λ_{kn} matrix.

Tables A.2 and A.3 present the results of these simulations for continuum-model changes and finite-sample changes, respectively. As established in the $\Lambda = 0$ case, which appears in Table II, the calibrated-shares procedure is highly data demanding. For all values of Λ reported in Tables A.2 and A.3, the calibrated-shares procedure's predictions become unbiased (slope near one) only as the number of individuals becomes orders of magnitude larger than the true value. Since the calibrated-shares procedure can capture the unobserved component of commuting costs, its mean forecast error becomes smaller than that of the covariates-based model when the unobserved variation is sufficiently large. For the number of individuals in the empirical setting ($I = 2{,}488{,}905$), this occurs only when the unobserved component of commuting costs varies at least a quarter as much as the observed component ($\Lambda \geq 0.25$) when we evaluate continuum-model changes and more than half as much as the observed component ($\Lambda > 0.5$) when we evaluate finite-sample changes.

A.4. *Monte Carlo Simulations: Finite-Sample Changes With Fixed* ν^I

This Appendix revisits the Monte Carlo investigation of Section 3.4 using a data-generating process in which the economic primitives change from Υ to Υ' but the set of idiosyncratic preferences ν^I is fixed. This leverages the model introduced in Section 4 as the data-generating process. Contrasted with the finite-sample outcomes with iid-ν idiosyncrasies, these finite-sample simulations are substantially closer to the continuum-change outcomes.

Figure A.3 is the fixed-ν analogue of Figure 2, depicting the distributions of the regression coefficients and relative MSE for the covariates-based model and calibrated-shares procedure. The covariates-based model outperforms the calibrated-shares procedure. The regression coefficient results are very similar to Figure 2. The covariates-based model remains an unbiased predictor, while the calibrated-shares procedure has a median slope of 0.78. The ratio of mean squared errors is closer to one in these simulations because the noisy finite-sample outcomes mean that both approaches have larger forecast

TABLE A.2

MONTE CARLO SIMULATIONS WITH $\lambda_{kn} \neq 1$: REGRESSAND IS CONTINUUM CHANGE.

		Slope (Mean)		MSE (Mean)	
Λ	I	Covariates-Based	Calibrated-Shares	Covariates-Based	Calibrated-Shares
0	2.5	0.9984	0.7793	0.0014	0.2280
0	5	0.9992	0.8759	0.0007	0.1133
0	12.5	0.9997	0.9483	0.0003	0.0451
0	25	0.9998	0.9725	0.0001	0.0227
0	50	0.9999	0.9868	0.0001	0.0113
0	125	1.0000	0.9951	0.0000	0.0045
0	250	1.0000	0.9973	0.0000	0.0023
0	2560	1.0000	0.9997	0.0000	0.0002
0.1	2.5	0.9968	0.7869	0.0380	0.2269
0.1	5	0.9977	0.8810	0.0373	0.1129
0.1	12.5	0.9983	0.9501	0.0370	0.0453
0.1	25	0.9984	0.9731	0.0368	0.0226
0.1	50	0.9985	0.9865	0.0368	0.0113
0.1	125	0.9986	0.9944	0.0367	0.0045
0.1	250	0.9986	0.9971	0.0367	0.0023
0.1	2560	0.9986	0.9996	0.0367	0.0002
0.25	2.5	0.9957	0.8199	0.2405	0.2262
0.25	5	0.9961	0.9025	0.2399	0.1129
0.25	12.5	0.9969	0.9581	0.2395	0.0454
0.25	25	0.9971	0.9787	0.2393	0.0226
0.25	50	0.9971	0.9889	0.2393	0.0113
0.25	125	0.9972	0.9954	0.2393	0.0045
0.25	250	0.9972	0.9978	0.2393	0.0023
0.25	2560	0.9972	0.9997	0.2392	0.0002
0.5	2.5	0.9954	0.8923	1.1084	0.2271
0.5	5	0.9962	0.9451	1.1076	0.1130
0.5	12.5	0.9969	0.9782	1.1071	0.0448
0.5	25	0.9971	0.9877	1.1071	0.0223
0.5	50	0.9972	0.9944	1.1070	0.0113
0.5	125	0.9973	0.9980	1.1069	0.0045
0.5	250	0.9973	0.9986	1.1069	0.0022
0.5	2560	0.9973	0.9999	1.1069	0.0002
1	2.5	1.0057	0.9691	7.0808	0.2186
1	5	1.0063	0.9847	7.0806	0.1090
1	12.5	1.0065	0.9946	7.0809	0.0435
1	25	1.0069	0.9958	7.0802	0.0219
1	50	1.0068	0.9985	7.0805	0.0109
1	125	1.0070	0.9992	7.0803	0.0044
1	250	1.0071	0.9994	7.0803	0.0022
1	2560	1.0070	0.9998	7.0803	0.0002

Note: This table reports the mean outcomes across 100 simulated event studies for both the covariates-based model and the calibrated-shares procedure. The data-generating process is the covariates-based model estimated on data for New York City in 2010 augmented by an unobserved component of commuting costs of magnitude Λ (see text). The regressand is the continuum change in commuters. The value of I is stated in millions of individuals. The value labeled "2.5" million is in fact $I = 2{,}488{,}905$. The $\Lambda = 0$ results also appear in Panel A of Table II.

TABLE A.3

MONTE CARLO SIMULATIONS WITH $\lambda_{kn} \neq 1$: REGRESSAND IS FINITE-SAMPLE CHANGE.

		Slope (Mean)		MSE (Mean)	
Λ	I	Covariates-Based	Calibrated-Shares	Covariates-Based	Calibrated-Shares
0	2.5	0.9744	−0.4372	14.4853	17.1252
0	5	0.9981	0.2281	7.1769	8.5177
0	12.5	1.0060	0.6695	2.8879	3.4181
0	25	0.9939	0.8216	1.4331	1.6968
0	50	1.0035	0.9169	0.7146	0.8463
0	125	1.0034	0.9687	0.2872	0.3402
0	250	1.0007	0.9832	0.1432	0.1696
0	2560	0.9997	0.9979	0.0140	0.0166
0.1	2.5	0.9938	−0.3548	14.5123	17.1386
0.1	5	0.9901	0.2324	7.2340	8.5081
0.1	12.5	1.0004	0.6782	2.9092	3.4009
0.1	25	0.9994	0.8340	1.4768	1.7056
0.1	50	0.9978	0.9148	0.7557	0.8522
0.1	125	0.9932	0.9606	0.3234	0.3398
0.1	250	0.9977	0.9813	0.1795	0.1704
0.1	2560	0.9983	0.9979	0.0505	0.0164
0.25	2.5	0.9709	−0.1494	14.6285	17.0442
0.25	5	1.0028	0.3911	7.4322	8.5047
0.25	12.5	0.9930	0.7333	3.1229	3.4121
0.25	25	0.9940	0.8639	1.6771	1.6987
0.25	50	0.9915	0.9270	0.9570	0.8495
0.25	125	0.9947	0.9693	0.5244	0.3405
0.25	250	0.9987	0.9875	0.3820	0.1687
0.25	2560	0.9968	0.9977	0.2524	0.0166
0.5	2.5	0.9973	0.3116	15.6180	17.1638
0.5	5	1.0082	0.6617	8.4153	8.5620
0.5	12.5	1.0039	0.8589	4.0031	3.4027
0.5	25	0.9925	0.9212	2.5287	1.6891
0.5	50	0.9990	0.9616	1.8189	0.8509
0.5	125	0.9991	0.9883	1.3973	0.3351
0.5	250	0.9950	0.9911	1.2505	0.1701
0.5	2560	0.9968	0.9989	1.1195	0.0166
1	2.5	0.9945	0.7774	20.9828	17.1174
1	5	1.0069	0.9059	14.2812	8.4447
1	12.5	1.0080	0.9710	10.3390	3.3886
1	25	1.0083	0.9688	8.4068	1.7038
1	50	1.0082	0.9929	7.8493	0.8545
1	125	1.0054	0.9971	7.4572	0.3435
1	250	1.0059	0.9976	7.1958	0.1702
1	2560	1.0075	0.9996	7.0915	0.0163

Note: This table reports the mean outcomes across 100 simulated event studies for both the covariates-based model and the calibrated-shares procedure. The data-generating process is the covariates-based model estimated on data for New York City in 2010 augmented by an unobserved component of commuting costs of magnitude Λ (see text). The regressand is the simulation-specific change in commuters from a realization drawn from the model using pre-shock parameter values to a realization drawn from the model using post-shock parameter values. The value of I is stated in millions of individuals. The value labeled "2.5" million is in fact $I = 2{,}488{,}905$. The $\Lambda = 0$ results also appear in Panel B of Table II.

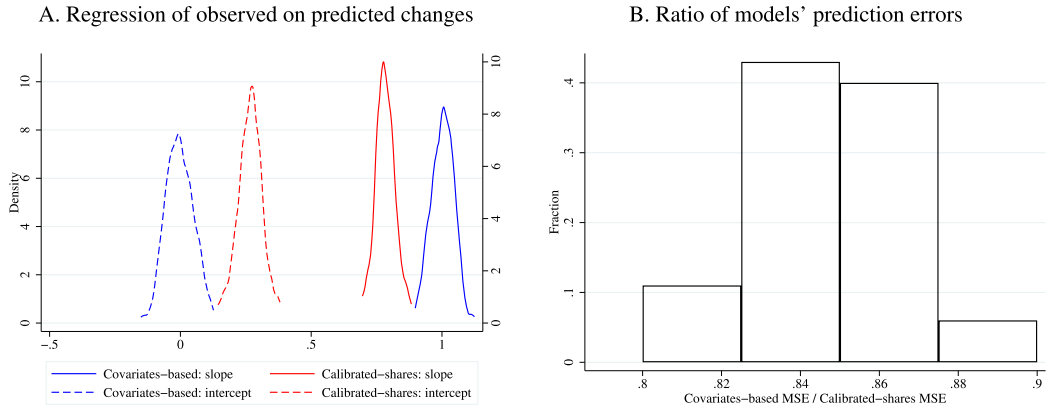

FIGURE A.3.—Calibrated-shares procedure performs poorly in fixed $-\nu^I$ simulations. NOTES: This figure depicts the regression coefficients and MSEs from 100 simulations in which $I = 2{,}488{,}905$, the data-generating process is the estimated covariates-based model of New York City in 2010 with fixed idiosyncratic preferences ν^I, and the counterfactual change is a 9% increase in productivity in one workplace tract.

errors than in the continuum-outcomes case. The MSE ratio in the median simulation is about 85%.

Fixing the idiosyncratic preferences reduces the excess churn in simulations with iid ν draws: across 100 simulations, the standard deviation of the total employment increase in the "treated tract" is about 185 in the iid-ν case and 48 in the fixed-ν case. Eliminating excess churn reduces the forecast error for both approaches, but their relative predictive performance is similar to the iid case. Across 100 simulations, the covariates-based model's average MSEs with iid and fixed idiosyncratic shocks are 14.49 and 1.24, respectively. The calibrated-shares procedure's average MSEs are 17.13 and 1.47, respectively.

APPENDIX B: EVENT STUDIES

This Appendix presents additional details regarding the event studies examined in Section 3.5.

B.1. *Two Employment Booms*

Figure B.1 depicts the time series of total employment in two New York City census tracts that contain 200 Fifth Avenue and 111 Eighth Avenue. Between 2010 and 2012, Tiffany & Co. and Google both substantially increased their employment in these two tracts. These tracts are two of the 83 examined in the event studies in Section 3.5.

B.2. *Event Studies Using Temporal and Geographic Aggregation*

Figure B.2 presents the results of pooling multiple years of data and aggregating the predicted changes in number of commuters to the booming workplace tracts from residential tracts up to residential NTAs.

B.3. *Estimating the Model Using Larger Spatial Units*

One may try to avoid the calibrated-shares procedure's overfitting problem by analyzing larger spatial units. Studying New York City's 195 Neighborhood Tabulation Ar-

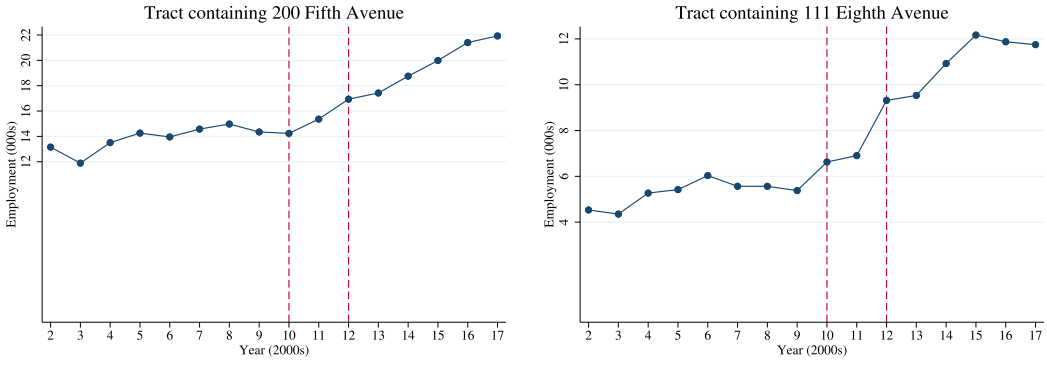

FIGURE B.1.—Employment in two of the event-study tracts. NOTES: This figure depicts the number of primary jobs held by New York City residents in tracts 36061005800 and 36061008300 in the LODES data.

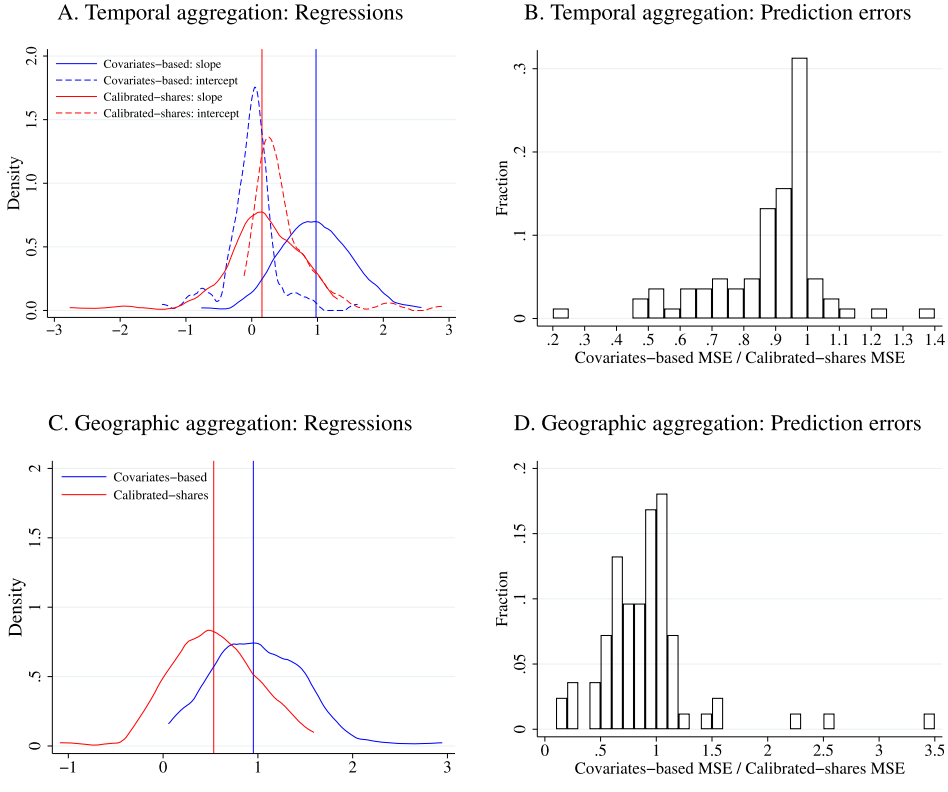

FIGURE B.2.—Models' predictive performance with temporal and geographic aggregation. NOTES: This figure contrasts the covariates-based model and calibrated-shares procedure's predictions for 83 tract-level employment booms when employing temporal aggregation or geographic aggregation. Panels A and B depict results when using baseline data for 2008–2010 instead of only 2010. Panels C and D depict predictive performance after aggregating the tract-level predictions underlying Panels A and B of Figure 3 up to 195 Neighborhood Tabulation Areas. We sum the predicted change in commuters to the booming workplace tract across all residential tracts in an NTA. In Panel C, only the slope coefficients are depicted. Vertical lines depict the medians of the slope coefficient distributions.

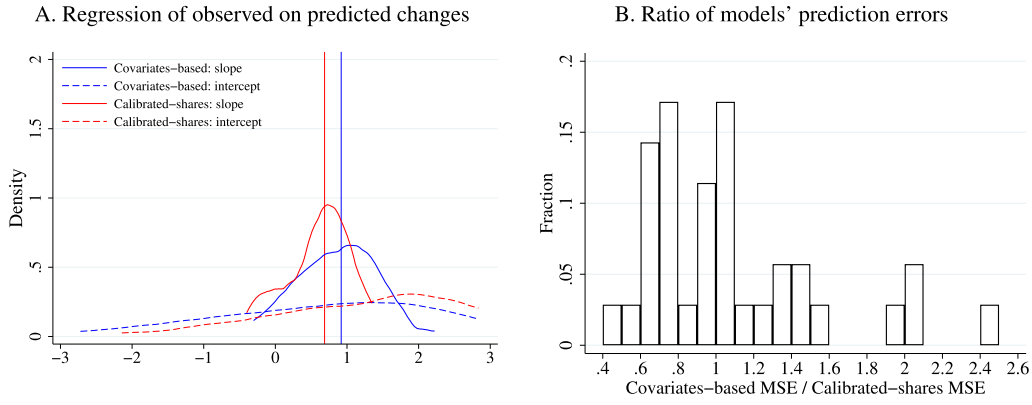

FIGURE B.3.—Models' predictive performance across 35 NTA-level employment booms. NOTES: This figure depicts predictive performance when fitting both approaches using data aggregated up to the 195 Neighborhood Tabulation Areas defined by the New York City Department of City Planning. In panel A, 10 intercept coefficients are not depicted for the covariates-based model. 5 intercept coefficients are not depicted for the calibrated-shares procedure. Panel A depicts the distributions of slope coefficients, and the vertical lines denote the median slope coefficients.

eas (NTAs) instead of tracts reduces the number of locations by an order of magnitude and the number of location pairs by two orders. To do so, we aggregate tract-to-tract commuter counts up to NTA-to-NTA flows, compute NTA-to-NTA commuting costs as the average of tract-to-tract commuting costs $\bar{\delta}_{kn}$, and compute NTA-level wages using employment-weighted averages of tract-level wages. We use both the covariates-based model and the calibrated-shares procedure to predict changes in NTA-to-NTA commuting flows for the 35 NTAs that had 2010–2012 employment growth of at least 12.5% and at least 2000 employees in 2010.

When analyzing these substantially larger spatial units, the covariates-based approach and the calibrated-shares procedure perform similarly. Figure B.3 shows that the covariates-based model's median slope coefficient is closer to one. The covariates-based model has a lower MSE than the calibrated-shares procedure in 18 of the 35 events.

B.4. *Noise Reduction via Low-Rank Approximation*

B.4.1. *The Approximation Procedure*

One method for guarding against idiosyncratic noise in the observed commuting matrix is to use values of ℓ_{kn} derived from a low-rank approximation of the matrix rather than the observed values when plugging ℓ_{kn} and $y_{kn} = \ell_{kn} w_n / \bar{\delta}_{kn}$ into equations (5)–(7). That is, for some particular rank $\mathfrak{r}$, replace the commuting matrix $\mathbf{L} = [\ell_{kn}]$ with

$$\tilde{\mathbf{L}} = \arg \min_{\mathbf{M} \in \mathbf{R}^{K \times N}} \|\mathbf{L} - \mathbf{M}\|_F^2, \quad \text{subject to } \operatorname{rank}(\mathbf{M}) \leq \mathfrak{r}, \tag{B.1}$$

where $\|\cdot\|_F$ is the Frobenius matrix norm. The solution to this problem can be computed from the singular value decomposition (SVD) of the original matrix $\mathbf{L}$. In particular, the optimal approximation $\tilde{\mathbf{L}}$ is given by the following algorithm (Hastie, Tibshirani, and Wainwright (2015)):

- Compute the SVD of $\mathbf{L} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$, where $\mathbf{U}$ is a $K \times K$ orthonormal matrix, $\mathbf{\Sigma}$ is a nonnegative $K \times N$ diagonal matrix with diagonal elements representing the singular values sorted in descending order, and $\mathbf{V}^T$ is an $N \times N$ orthonormal matrix.
- Keep only the first (largest) $\mathfrak{r}$ singular values in Σ, to obtain $\tilde{\mathbf{L}} := U\left(\begin{smallmatrix}\Sigma_{\mathfrak{r}} & 0\\ 0 & 0\end{smallmatrix}\right)V^T$, where $\Sigma_{\mathfrak{r}}$ is the upper left $\mathfrak{r} \times \mathfrak{r}$ block in Σ.

However, this approximation is imperfect for our purposes on two accounts: first, while the original matrix L is non-negative (as it consists of counts), there is no guarantee that the rank-r approximation is also non-negative. We therefore set all negative entries in $\tilde{\mathbf{L}}$ to 0. Additionally, $\tilde{\mathbf{L}}$ must be rescaled so that the sum across its entries matches the original; that is, we obtain $\tilde{\mathbf{L}}_{\mathfrak{r}}^{\text{SVD}}$ by replacing negative entries in the rank-r SVD approximation with zeros, then rescaling such that $\sum_{k,n}(\tilde{\mathbf{L}}_{\mathfrak{r}}^{\text{SVD}})_{kn} = \sum_{k,n}\ell_{kn}$.

Figure B.4 visualizes the approximated matrices derived from three choices of rank. alongside the observed 2010 data and the fitted values from the covariates-based and rank-1 interactive-fixed-effects specifications. (The IFE procedure is described in detail

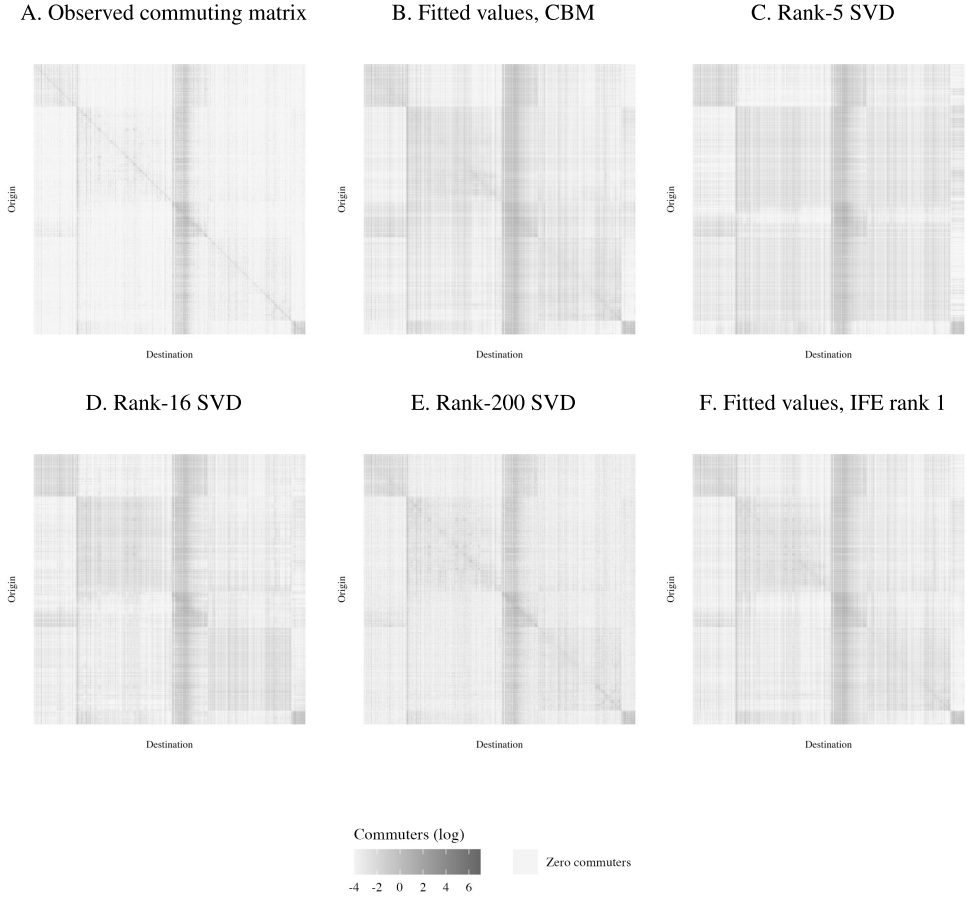

FIGURE B.4.—Approximations of the commuting matrix. NOTES: This figure depicts the observed 2010 commuting matrix, the fitted covariates-based model (CBM) of ℓ_{kn}, approximations derived from rank-5, rank-16, and rank-200 singular value decompositions, and a rank-1 interactive-fixed-effects specification. Rows and columns are ordered by FIPS codes, so that all but eight adjacent entries are in the same county.

in Appendix B.5.) The SVD-approximated commuting matrices preserve many features of the observed commuting matrix even under fairly stringent rank restrictions. In particular, workplaces with a large number of commuters are preserved, as are clusters of tracts with comparatively high or low commuting shares. Some features, however, are discarded by the approximation. Most notably, the higher values on the diagonal of the matrix (which corresponds with commuters who choose to live and work in the same tract) are missing from low-rank approximations (e.g., rank 5 and rank 16) but are contained in higher-rank approximations (e.g., rank 200). This is because the diagonal of a matrix is a fundamentally "high-rank" phenomenon and cannot in general be replicated by a low-rank approximation.

Table B.3 summarizes the sparsity generated by approximations derived from various ranks, including both the SVD-based procedure described above and the non-negative factorization procedure described in Section B.4.3.

B.4.2. *Selecting the Rank*

When applying any rank-restricted algorithm, the rank of the approximation is a key choice. Since we are using low-rank approximation to reduce idiosyncratic noise, the choice of the rank parameter $\mathfrak{r}$ balances the trade-off between the signal and noise in observed commuting flows. We present multiple procedures from which a researcher can select to choose the rank used; some are computationally and conceptually simple (drawing from the established norms in machine learning), while a Monte Carlo simulation is more complex and tailored to a particular context. We describe each procedure in turn and present results, including their suggested rank.

B.4.2.1. *Selecting the Rank by Simple Methods.* After computing the SVD of the commuting matrix (which takes only seconds), we obtain an ordered list of singular values $\sigma_1, \sigma_2, \ldots, \sigma_N$ with $\sigma_i \geq \sigma_j$ for $i < j$. These are depicted in a scree plot in Figure B.5, which normalizes the singular values by dividing each by the sum of all singular values (i.e., $\frac{\sigma_{il}}{\sum_i \sigma_i}$) and restricts attention to the 25 largest singular values. This normalization allows us to interpret each singular value as representing a fraction of the matrix's total variation (Cohen (2021, p. 499)).

Figure B.5 allows us to evaluate two distinct rank criteria. The first is an informal procedure known as the "elbow" method, which is a method of visual inspection that attempts to identify the singular values that are substantially larger than the rest. The second, the variance criterion, is numerically founded but nonetheless requires a cutoff chosen by the researcher (Cohen (2021, p. 496)). For the variance criterion, we compare the magnitude of each singular value with 0.5% of the sum of singular values; that is, we identify the set $\{\sigma_i | \sigma_i \geq 0.005 \sum_j \sigma_j\}$. This cutoff is shown with a dashed horizontal line in the scree plot. The two methods offer substantially different answers in our case, with the elbow method suggesting a lower rank of 4 and the 0.5% variance criterion suggesting a rank of 16. We select the more generous of these values and use rank 16 as our preferred specification in the main text.

B.4.2.2. *Selecting the Rank Using Monte Carlo Simulations.* Another method for selecting the rank is dependent on the research setting permitting Monte Carlo simulations. One can select the rank based on ranks' performance across Monte Carlo simulations (using regression coefficients and MSEs summarizing commuting-flow predictions). While substantially more computationally complex than the simple methods above, this procedure yields results that are more closely tied to the particular structure of the empirical setting.

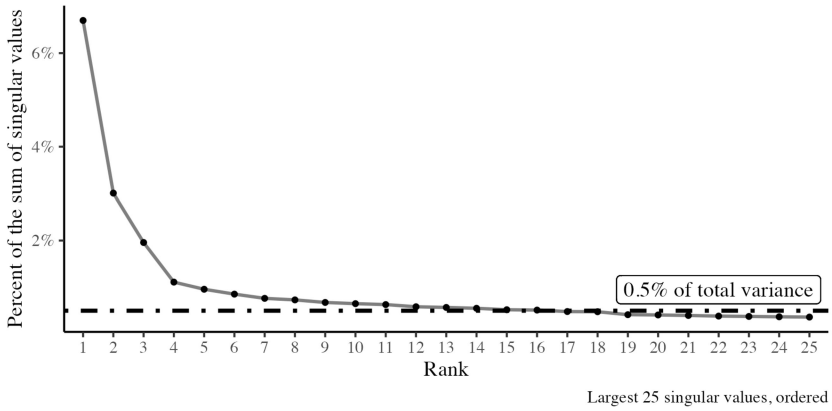

FIGURE B.5.—Scree plot (ordered singular values) for NYC 2010 commuting matrix. NOTES: This figure shows the magnitude of successive singular values for the NYC 2010 commuting matrix. The dashed line indicates the threshold of 0.5% of the sum of the singular values.

Starting from the steps outlined in Appendix A.1 across 100 Monte Carlo draws from the continuum model, we follow steps 1 through 5 with two exceptions. First, we consider only changes in commuter flows, not rents; and second, we substitute the SVD-approximated commuting counts $\tilde{\ell}_{kn}^{\text{SVD}}$ for the "observed" counts, ℓ_{kn}, which are drawn from the data-generating process. Note that this requires us to approximate the commuting matrix separately for each of the Monte Carlo draws. As in Appendix A.1, we estimate the commuting elasticity ϵ for each of the Monte Carlo draws. The inputs into this commuting elasticity estimation are the "observed" values as in the covariates-based model, not the SVD-approximated commuting matrix. Finally, we raise productivity in the same tract (containing 200 Fifth Avenue) as in the Appendix A.1 simulations.

Table B.1 reports the performance across ranks in these Monte Carlo simulations. Rank 18 performs best in these simulations based on MSEs and is nearly best based on regression coefficients. Rank 16, which is suggested by the variance criteria shown in Figure B.5, performs very similarly.

The lower panel of Table B.1 presents the event study performance (for the same 83 employment booms as in Section 3.5.3) for approximations derived from a wide variety of ranks. The rank-16 results are shown in Figure 4.

B.4.3. *Non-Negative Matrix Factorization*

While the previously outlined SVD procedure is one means of implementing noise reduction, truncating the commuting shares to be non-negative means that the resulting approximation is not necessarily low-rank. Another approach is to use an algorithm that imposes the nonnegativity constraint. That is, we find a solution to the problem

$$\tilde{\mathbf{L}}^{\text{NNMF}} = \arg \min_{\mathbf{M} \in \mathbf{R}^{K \times N}} \|\mathbf{L} - \mathbf{M}\|_F^2 \quad \text{subject to } \operatorname{rank}(\mathbf{M}) \leq \mathfrak{r}, \text{ and } [\mathbf{M}]_{kn} \geq 0. \tag{B.2}$$

Solving this problem is more difficult and, unlike in the unconstrained case in which the solution is uniquely determined by SVD, there may be multiple (local) optima. We use the Greedy Coordinate Descent algorithm from Hsieh and Dhillon (2011) with a fixed initialization as implemented in the Julia package NMF. This procedure produces low-rank approximated matrices that deliver very similar performance as the computationally

TABLE B.1

SUMMARY OF SVD MONTE CARLO AND EVENT STUDY RESULTS BY RANK.

	Rank																			
	1	2	3	4	5	6	8	10	12	14	15	16	18	20	50	100	500	1000	1500	Full Rank
Monte Carlo performance																				
Slope	1.03	1.05	1.04	1.04	1.02	1.01	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.99	0.91	0.78	0.78	0.78	0.78	0.78
Int.	−0.041	−0.059	−0.056	−0.049	−0.021	−0.012	0.001	−0.002	−0.002	−0.001	−0.001	0.001	0.003	0.014	0.116	0.268	0.271	0.271	0.272	0.272
MSE	0.1579	0.0461	0.0460	0.0460	0.0372	0.0357	0.0333	0.0322	0.0323	0.0320	0.0318	0.0315	0.0311	0.0311	0.0930	0.2241	0.2279	0.2280	0.2280	0.2280
Event study performance																				
Slope	0.73	0.70	0.71	0.80	0.83	0.86	0.85	0.83	0.83	0.82	0.82	0.81	0.80	0.79	0.62	0.32	−0.43	−0.47	−0.47	−0.46
Int.	0.06	0.14	0.14	0.09	0.08	0.08	0.10	0.13	0.13	0.14	0.14	0.15	0.16	0.17	0.30	0.51	0.80	0.82	0.82	0.82
MSE	10.53	10.38	10.37	10.29	10.27	10.26	10.27	10.30	10.32	10.32	10.38	10.40	10.45	10.48	10.96	11.71	12.94	13.23	13.35	13.39

Note: This table reports the predictive performance, measured by the average values of the regression slope and intercept coefficients and MSE, using SVD approximations of various ranks across two cases: (a) 100 Monte Carlo simulations, with each drawn from the parameterized continuum model with a 9% productivity shock to the tract containing 200 Fifth Avenue, where the performance is evaluated relative to the continuum model's predictions; and (b) the 83 event studies examined in Section 3.5.

TABLE B.2

SUMMARY OF NNMF EVENT STUDY RESULTS BY RANK.

	Rank														
	1	2	3	4	5	6	8	10	12	14	16	18	20	15	50
	Event study performance														
Slope	0.73	0.70	0.71	0.78	0.78	0.81	0.85	0.87	0.87	0.87	0.87	0.86	0.83	0.81	0.71
Int.	0.06	0.14	0.14	0.11	0.12	0.10	0.10	0.10	0.10	0.10	0.10	0.11	0.13	0.15	0.25
MSE	10.53	10.38	10.37	10.28	10.27	10.29	10.30	10.24	10.25	10.25	10.26	10.31	10.34	10.42	10.88

Note: This table reports the predictive performance, as measured by the average values of the regression slope and intercept coefficients and MSE, across the 83 event studies examined in Section 3.5 using various ranks of NNMF low-rank approximated commuting matrices.

simpler SVD procedure described above. The results for event studies are shown in Table B.2.

B.4.4. *Results*

Table B.1 reports the event study and Monte Carlo performance of the SVD-approximated matrices. Table B.2 reports the event study performance of the non-negative matrix factorization (NNMF) approximations across ranks. As the non-negative matrix factorization appears to provide no meaningful improvement over the simpler SVD approximation procedure, we do not present any further analysis via Monte Carlo simulations for that procedure. Finally, Table B.3 reports the proportion of the approximated commuting matrix's entries that are zero across ranks for both approx imation methods. While we used $\bar{\delta}_{kn}$ to calculate $y_{kn} = w_n \ell_{kn} / \bar{\delta}_{kn}$ in line with the rest of the paper, we have verified that the SVD approximation predictions are nearly identical when using $y_{kn} = w_n \ell_{kn}$.

B.5. *Interactive Fixed Effects*

B.5.1. *Modeling Unobserved Commuting Costs*

This section extends the covariates-based model to incorporate unobserved disutility components of commuting costs. We assume that λ_{kn} has a low-rank factor structure.

TABLE B.3

PERCENTAGE (%) OF APPROXIMATED COMMUTING MATRIX ENTRIES EQUAL TO ZERO.

	Rank																			
	1	2	3	4	5	6	8	10	12	14	15	16	18	20	50	100	500	1000	1500	2143
SVD	0	1.9	2	6	11	12	15	15	16	17	17	17	17	18	22	26	37	41	42	85
NNMF	0	0.4	0.1	3	2.6	3.9	2.5	2	1.7	1.5	1.4	1.3	1	1.1	0.7					

Note: This table reports the proportion of entries in the approximated commuting matrices, across ranks, that are equal to zero. These values are shown for both the SVD and the NNMF procedures, though we omit high-rank approximations via NNMF because of computational burden. The SVD procedure results in many more zeros, particularly for larger ranks, than the NNMF, largely due to the generation of negative values by SVD prior to our truncating them to zero.

Specifically, we parameterize the overall commuting costs δ_{kn} as

$$\delta_{kn} = \bar{\delta}_{kn} \times \underbrace{\exp\left(\psi_k' \gamma_n\right)}_{=\lambda_{kn}^{\text{IFE}}},$$

where ψ_k and γ_n are $R \times 1$ vectors that give rise to an R-dimensional factor structure. The log likelihood function then becomes

$$\ln \mathcal{L}^{\text{IFE}} = \sum_{k,n} \ell_{kn} \ln\left[\frac{w_n^{\epsilon}\left(r_k^{\alpha} \bar{\delta}_{kn} \exp\left(\psi_k' \gamma_n\right)\right)^{-\epsilon}}{\sum_{k',n'} w_{n'}^{\epsilon}\left(r_{k'}^{\alpha} \bar{\delta}_{k'n'} \exp\left(\psi_{k'}' \gamma_{n'}\right)\right)^{-\epsilon}}\right]. \tag{B.3}$$

In practice, to implement the estimation of the parameters, we follow Chen, Fernández-Val, and Weidner (2021), who estimate interactive fixed effects in a Poisson model. Here, multiple practical difficulties arise. Most critically, the likelihood is no longer convex; in particular, there is the possibility of multiple local optima. To guard against this, we use multiple randomized starting points and select the parameter estimates that produce the highest likelihood. Another difficulty is that the computational cost of this procedure rises rapidly with the factor dimension R. We consider $R \leq 3$. Note that in all specifications, additive fixed effects are already included to obtain estimates for w_n^{ϵ} and r_k^{α}; that is, R gives the dimension of the "purely" interactive fixed effects.[49]

In the main text, we show the estimates from the model with $R = 1$ for two reasons. First, we take advantage of the relative computational ease of the $R = 1$ case to estimate from a large number of starting points and find a consistent convergence pattern that suggests our estimate is the true optimum. Second, the eigenvalue ratio test suggested by Chen, Fernández-Val, and Weidner (2021) indicates that $R = 1$ is the optimal factor structure among the set of options $\{1, 2\}$. Possible choices of R span the range from the covariates-based specification (when $R = 0$) to the calibrated-shares procedure (the full-rank case).[50] The event study results for rank $R = \{2, 3\}$ are very similar to those for rank $R = 1$ and are shown below.

B.5.2. *Estimation Results*

Table B.4 presents the estimates from maximizing the log likelihood function in Eq. (B.3) using the algorithm outlined by Chen, Fernández-Val, and Weidner (2021). We show results for $R = 0, 1, 2, 3$, where $R = 0$ is identical to the covariates-based model and the OLS column is included for the sake of comparison. The magnitude of the commuting elasticity ϵ diminishes with the rank of the factor structure. As expected, the model fit metrics increase with the factor structure rank, as higher ranks correspond to more free parameters to attain better (in-sample) fit.

[49] This differs from the notation in Chen, Fernández-Val, and Weidner (2021), who include the additive fixed effects in the factor dimension R. Our purely interactive factor dimension is, in their notation, R_2.

[50] Consider the case of a 2×2 commuting matrix A. This matrix can be decomposed into two 2×2 matrices (e.g., AI). However, there are many such decompositions, illustrating a failure of identification when fitting the commuting matrix with full-rank interactive fixed effects. Notice also that simply limiting the dimensions of the decomposition is not sufficient. In the 2×2 case, the only matrices A that can be written as the product of a 2×1 with a 1×2 vector are those with rank 1. As our commuting matrix is full rank, this suggests why a very high rank (and accompanying identification issues) would be needed to reproduce the calibrated-shares procedure using only high-rank interactive fixed effects.

TABLE B.4

COMMUTING ELASTICITY ESTIMATES FROM IFE SPECIFICATION.

		MLE			
	OLS	$R=0$	$R=1$	$R=2$	$R=3$
Commuting cost	−2.307	−7.986	−7.177	−6.654	−6.36
(pseudo-) R^2	0.561	0.662	0.685	0.694	0.701
Location pairs	690,673		4,628,880		
Commuters	2,488,905		2,488,905		

Note: All specifications include residence fixed effects and workplace fixed effects. The columns under "MLE" present the results from the maximum likelihood estimation described in the text, with R denoting the interactive fixed effects factor structure rank. The "OLS" column presents the results of estimating the log version of equation (2) by ordinary least squares, omitting observations where $\ell_{kn}=0$. The model-fit statistic is pseudo-R^2 for MLE and R^2 for OLS.

B.5.3. *Event Study Results for Rank* $R=\{2,3\}$

In Figure B.6, we display the event study performance for IFE factor structure ranks 2 and 3. The extreme similarity in performance across the ranks suggests that rank 1 is sufficient to capture most of the expected gains from using IFE over the CBM.

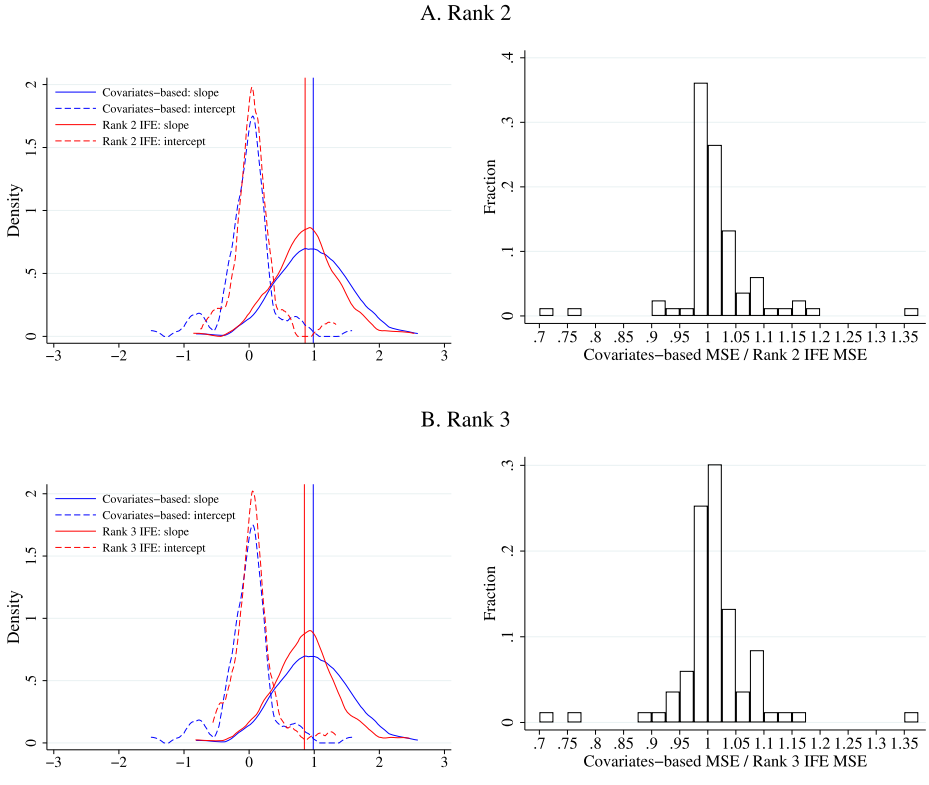

FIGURE B.6.—Event study performance: varying factor structure rank, IFE. NOTES: This figure summarizes the predictive performance of the covariates-based model and the ranks 2 and 3 IFE specifications, analogous to Figure 5. Panel A depicts the distributions of slope coefficients, and the vertical lines denote the median slope coefficients.

B.6. *Tests of in-Sample Fit*

In this appendix section, we assess the in-sample fit of different model parameterizations. The χ^2 test contrasts the expected and observed frequencies of outcomes and compares this difference to a critical value from the test-statistic distribution. Since in our application the expected frequencies of outcomes are small, the test statistic may not follow the asymptotic χ^2 distribution. We therefore construct the test statistic's distribution using a parametric bootstrap.

The test statistic is $\chi^2 = \sum_{kn} \frac{(\ell_{kn}^{\text{obs}} - \mathrm{E}[\ell_{kn}])^2}{\mathrm{E}[\ell_{kn}]}$, where ℓ_{kn}^{obs} is the observed number of individuals who reside in k and work in n in the data and $\mathbb{E}[\ell_{kn}]$ is the expected number of commuters (given by ℓ_{kn} in equation (2)). We construct the model-implied distribution of this test statistic from a parametric bootstrap in the following way:

1. We calculate the test statistic using

$$\chi^{2\,\text{simulated}} = \sum_{kn} \frac{\left(\ell_{kn}^{\text{sim}} - \mathbb{E}[\ell_{kn}]\right)^2}{\mathbb{E}[\ell_{kn}]}, \quad \ell_{kn}^{\text{sim}} \overset{\text{iid}}{\sim} \text{Multinomial}\big(I, \{p_{kn}\}\big),$$

   where $p_{kn} \equiv \Pr(U_{kn}^i > U_{k'n'}^i \ \forall (k', n') \neq (k, n))$ is the model-implied probability of choosing kn defined in equation (8) and $I = 2{,}488{,}905$ is the total number of commuters.
2. We repeat the first step 1000 times to obtain the parametric-bootstrap-based distribution $\hat{f}(\chi^{2\,\text{simulated}})$.

The χ^2 test statistic for the covariates-based model is extremely large (calculated as 8.445e299, though numerical error may be sizable). This very large value of the test statistic is driven by outlier observations for which the model's expected number of commuters is extremely small (i.e., nearly equal to zero). By comparison, the test statistics for the interactive-fixed-effects specifications are much smaller: 1.006e7, 1.268e7, and 8.962e6 for ranks $R = 1$, 2, and 3, respectively. Nonetheless, all these test statistics exceed the 99$^{\text{th}}$ percentiles of the distributions of the χ^2 test statistics shown in Figure B.7A.

By contrast, the calibrated-shares procedure has perfect in-sample fit by construction and a χ^2 test statistic of zero. This fit is too good to be true: the test statistic of zero for the observed data is lower than all simulated values in the distribution generated by the parametric bootstrap.

We also assess the in-sample fit produced by the rank-restricted singular value decomposition (SVD) introduced in Section 3.5.3. One issue for this specification is that the rank-restricted SVD assigns zeros to some residence-workplace pairs that have positive counts in the data. A zero-probability event occurring rejects the parameterized model. We calculate the mean squared error (MSE) as an alternative test statistic that sidesteps the divide-by-zero issue. Specifically, we calculate $\text{MSE} = \frac{1}{K \times N} \sum_{kn} (\ell_{kn}^{\text{obs}} - \mathbb{E}[\ell_{kn}])^2$, where K is number of residence tracts and N is the number of workplace tracts. The parameterized bootstrapped distributions of MSE test statistics are shown in Figure B.7B. The approximated commuting matrix derived from a rank-110 SVD would lie within the 99-percent confidence interval, whereas higher-rank approximations exhibit an in-sample fit that is too good to be true.

B.7. *Spatially Correlated Idiosyncratic Preferences*

This appendix section relaxes the assumption that idiosyncratic preferences are independent across granular spatial units. In particular, preferences for residence-workplace

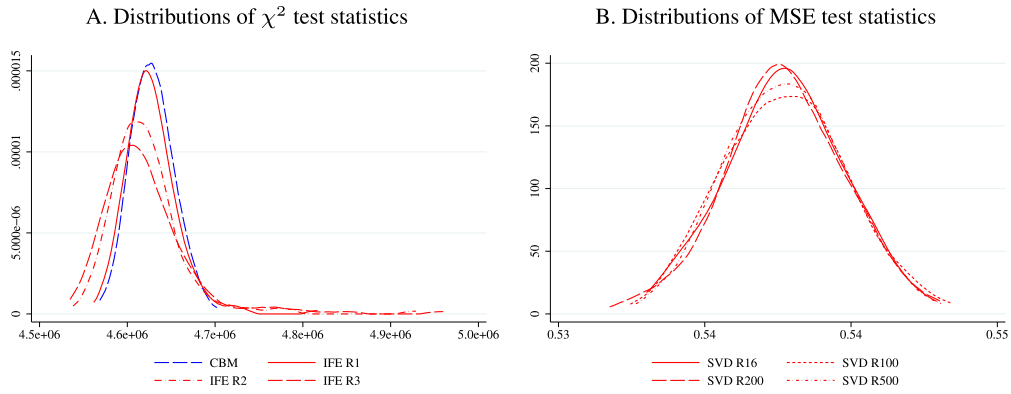

FIGURE B.7.—Distributions of χ^2 and MSE test statistics. NOTES: The test statistic in the left panel is $\chi^2 = \sum_{kn} \frac{(\ell_{kn}^{\text{obs}} - \mathrm{E}[\ell_{kn}])^2}{\mathrm{E}[\ell_{kn}]}$, where ℓ_{kn}^{obs} is the number of commuters from k to n and $\mathbb{E}[\ell_{kn}]$ is the expected value of commuters from k to n (given by ℓ_{kn} in the continuum model's equation (2)). We depict the distributions between their 0.5th and 99.5th percentiles. The test statistics using the commuter counts in the 2010 LODES data are 8.445e299, 1.006e7, 1.268e7, and 8.962e6 for the CBM, rank-1 IFE, rank-2 IFE, and rank-3 IFE, respectively. The test statistic in the right panel is $\text{MSE} = \frac{1}{K \times N} \sum_{kn} (\ell_{kn}^{\text{obs}} - \mathbb{E}[\ell_{kn}])^2$, where ℓ_{kn}^{obs} is the number of commuters from k to n, $\mathbb{E}[\ell_{kn}]$ is the expected value of commuters, K is number of residence tracts, and N is the number of workplace tracts. We depict the distributions between their 0.5th and 99.5th percentiles. The test statistics using the commuter counts in the 2010 LODES data are 1.467, 0.5725, 0.3581, and 0.1799 for the rank-16, rank-100, rank-200, and rank-500 singular value decompositions, respectively.

pairs have a nested-logit structure, such that individual i's indirect utility of residing in k and working in n is

$$\overset{\circ}{U}{}_{kn}^{i} = \underbrace{\overset{\circ}{\epsilon} \ln\left(\frac{w_n}{r_k^{\alpha} P^{1-\alpha} \delta_{kn}}\right)}_{\equiv \overset{\circ}{U}_{kn}} + \overset{\circ}{\nu}{}_{kn}^{i}, \tag{B.4}$$

where $\overset{\circ}{U}_{kn}$ denotes the mean utility of choice kn and the parameter $\overset{\circ}{\epsilon}$ governs the importance of mean utility relative to idiosyncratic preferences in the nested-logit specification. We allow the idiosyncratic preferences $\overset{\circ}{\nu}{}_{kn}^{i}$ to be correlated across residences within Neighborhood Tabulation Areas (NTAs). We assume that the vector of idiosyncratic preferences has a cumulative distribution $F(\overset{\circ}{\nu}{}^{i}) = \exp\{-\sum_z (\sum_{(k,n)\in B_z} e^{-\overset{\circ}{\nu}{}_{kn}^{i}/\zeta})^{\zeta}\}$, where B_z denotes the set of residence-workplace kn pairs with the residence k in NTA z, and ζ governs the correlation of idiosyncratic preferences among alternatives in the set B_z. Whereas draws of idiosyncratic preferences between residence-workplace pairs with residences in distinct NTAs remain uncorrelated, the within-residence-NTA correlation of idiosyncratic preferences is $1 - \zeta^2$. Given Eq. (B.4) and this distribution of idiosyncratic preferences,

the fraction of people residing in k and working in n is

$$\mathbb{P}_{kn} = \underbrace{\frac{w_n^{\mathring{\epsilon}/\zeta}\left(r_k^{\alpha}\delta_{kn}\right)^{-\mathring{\epsilon}/\zeta}}{\sum_{(k',n')\in B_z} w_{n'}^{\mathring{\epsilon}/\zeta}\left(r_{k'}^{\alpha}\delta_{k'n'}\right)^{-\mathring{\epsilon}/\zeta}}}_{=\mathbb{P}_{kn|B_z}} \times \underbrace{\frac{\left[\sum_{(k',n')\in B_z} w_{n'}^{\mathring{\epsilon}/\zeta}\left(r_{k'}^{\alpha}\delta_{k'n'}\right)^{-\mathring{\epsilon}/\zeta}\right]^{\zeta}}{\sum_{z'}\left[\sum_{(k',n')\in B_z} w_{n'}^{\mathring{\epsilon}/\zeta}\left(r_{k'}^{\alpha}\delta_{k'n'}\right)^{-\mathring{\epsilon}/\zeta}\right]^{\zeta}}}_{=\mathbb{P}_{B_z}},$$

where $\mathbb{P}_{B_z}$ denotes the probability of residing in NTA z and $\mathbb{P}_{kn|B_z}$ denotes the probability of selecting residence-workplace pair kn among residents of NTA z.

Under the assumption of nested-logit preferences and no unobserved disutility ($\delta_{kn} = \bar{\delta}_{kn}\ \forall kn$), the counterfactual-baseline ratio of commuters is

$$\hat{\ell}_{kn} = \left(\frac{\hat{w}_n}{\hat{r}_k^{\alpha}\hat{\bar{\delta}}_{kn}}\right)^{\mathring{\epsilon}/\zeta} \times \frac{\left[\sum_{(k',n')\in B_z}\left(\frac{\hat{w}_{n'}}{\hat{r}_{k'}^{\alpha}\hat{\bar{\delta}}_{k'n'}}\right)^{\mathring{\epsilon}/\zeta} \times \mathbb{P}_{k'n'|B_z}\right]^{\zeta-1}}{\sum_{z'}\mathbb{P}_{B_{z'}}\left[\sum_{(k',n')\in B_{z'}}\left(\frac{\hat{w}_{n'}}{\hat{r}_{k'}^{\alpha}\hat{\bar{\delta}}_{k'n'}}\right)^{\mathring{\epsilon}/\zeta} \times \mathbb{P}_{k'n'|B_{z'}}\right]^{\zeta}} \quad \text{if } \ell_{kn} > 0.$$

This specification yields the following log likelihood function:

$$\ln\mathcal{L} \equiv \sum_{k,n}\ell_{kn}\ln\mathbb{P}_{kn} = \sum_{k,n}\ell_{kn}\ln\left[\frac{w_n^{\mathring{\epsilon}/\zeta}\left(r_k^{\alpha}\delta_{kn}\right)^{-\mathring{\epsilon}/\zeta} \times \left[\sum_{(k',n')\in B_z} w_{n'}^{\mathring{\epsilon}/\zeta}\left(r_{k'}^{\alpha}\delta_{k'n'}\right)^{-\mathring{\epsilon}/\zeta}\right]^{\zeta-1}}{\sum_{z'}\left[\sum_{(k',n')\in B_{z'}} w_{n'}^{\mathring{\epsilon}/\zeta}\left(r_{k'}^{\alpha}\delta_{k'n'}\right)^{-\mathring{\epsilon}/\zeta}\right]^{\zeta}}\right].$$

We sequentially estimate the parameters of the nested-logit model (Train (2009)). In the first step, we maximize the log likelihood with respect to $\mathring{\epsilon}/\zeta$, residence fixed effects, workplace fixed effects, and nest-specific fixed effects. Given the fixed effects, the estimate of $\mathring{\epsilon}/\zeta$ will be equal to the estimate of ϵ in the logit case. Note that the residence fixed effects are unique only up to a nest-specific normalization. We normalize the residence fixed effects relative to $r_k^{-\alpha\mathring{\epsilon}/\zeta}$ for some residence k in each NTA z. We index the baseline residence within a NTA as $\bar{k}(z)$. Once we impose the normalization, we have residence

fixed effects $(\frac{r_k}{r_{\bar{k}(z)}})^{-\alpha\mathring{\epsilon}}$. We can therefore write

$$\mathbb{P}_{kn|B_z} = \frac{\left[\left(\frac{r_k}{r_{\bar{k}(z)}}\right)^{-\alpha\mathring{\epsilon}} w_n^{\mathring{\epsilon}} \bar{\delta}_{kn}^{-\mathring{\epsilon}}\right]^{1/\zeta}}{\sum_{(k',n')\in B_z}\left[\left(\frac{r_{k'}}{r_{\bar{k}(z)}}\right)^{-\alpha\mathring{\epsilon}} w_{n'}^{\mathring{\epsilon}} \bar{\delta}_{k'n'}^{-\mathring{\epsilon}}\right]^{1/\zeta}} = \frac{\left[\left(\frac{r_k}{r_{\bar{k}(z)}}\right)^{-\alpha\mathring{\epsilon}} w_n^{\mathring{\epsilon}} \bar{\delta}_{kn}^{-\mathring{\epsilon}}\right]^{1/\zeta}}{I_z}, \quad \text{and}$$

$$\mathbb{P}_{B_z} = \frac{r_{k(z)}^{-\alpha\mathring{\epsilon}} I_z^{\zeta}}{\sum_{z'} r_{\bar{k}(z')}^{-\alpha\mathring{\epsilon}} I_{z'}^{\zeta}},$$

where I_z is the inclusive value of nest z. In the second step, we compute the inclusive values using the estimates from the first step. Having the estimates of the inclusive values, $\hat{I}_z$, we maximize the following log likelihood function with respect to $r_{\bar{k}(z)}$:

$$\ln \tilde{\mathcal{L}} \equiv \sum_z \left[\sum_{(k,n)\in B_z} \ell_{kn}\right] \ln \mathbb{P}_{B_z} = \sum_z \left(\left[\sum_{(k,n)\in B_z} \ell_{kn}\right] \times \ln\left[\frac{r_{\bar{k}(z)}^{-\alpha\mathring{\epsilon}} \hat{I}_z^{\zeta}}{\sum_{z'} r_{\bar{k}(z')}^{-\alpha\mathring{\epsilon}} \hat{I}_{z'}^{\zeta}}\right]\right).$$

Note that we cannot simultaneously identify $r_{\bar{k}(z)}$ and parameter ζ because the number of parameters exceeds the number of alternatives. We assume that the value of parameter ζ is known, in order to estimate $r_{\bar{k}(z)}$ up to a proportionality coefficient. We compute the equilibrium prices and the remaining economic primitives $\{T_k\}$ and $\{A_n\}$ from the estimated fixed effects and the market-clearing conditions.

We estimate the model for two values of parameter ζ: 0.25 and 0.75. We estimate each stage of the model by using the Poisson pseudo maximum likelihood estimator described in Section 3.2. When ζ is 0.75, the elasticity parameter $\mathring{\epsilon}$ is $5.99 (= 7.99 \times \zeta)$. The implied productivities are unchanged, and the implied land endowments have a correlation of 0.99 with the land endowments in the baseline model. When ζ is 0.25, the elasticity parameter $\mathring{\epsilon}$ is 2.0. The implied productivities are unchanged, and the implied land endowments have a correlation of 0.61 with the land endowments in the baseline model.

As long as the nested-logit model and the logit-model agree on the unconditional choice probabilities (i.e., they imply the same $\mathbb{P}_{k'n'}$), we note that in the nested-logit model there is greater within-nest substitution in response to a change in commuting costs than in the simple logit model. This can be seen in the second of the following three elasticities, noting that $\mathbb{P}_{k'n'|B_z} > \mathbb{P}_{k'n'}$:

$$\frac{\partial \ln \mathbb{P}_{kn}}{\partial \ln \delta_{k'n'}} = \begin{cases} -\epsilon(1 - \mathbb{P}_{k'n'|B_z}) - \mathring{\epsilon}\,(\mathbb{P}_{k'n'|B_z} - \mathbb{P}_{k'n'}) & \text{if } k' = k, n' = n \text{ and } (k,n) \in B_z, \\ \epsilon\left(\zeta \mathbb{P}_{k'n'} + (1-\zeta)\mathbb{P}_{k'n'|B_z}\right) & \text{if } (k',n') \neq (k,n), \left(k',n'\right) \in B_z \\ & \text{and } (k,n) \in B_z, \\ \mathring{\epsilon}\,\mathbb{P}_{k'n'} & \text{if } \left(k',n'\right) \notin B_z \text{ and } (k,n) \in B_z. \end{cases}$$

We use the model to compare the predictive performance of the covariates-based model and the calibrated-shares procedure in 83 employment booms that are detailed in Section 3.5. Figure B.8 contrasts the predictions of the two procedures for each value of

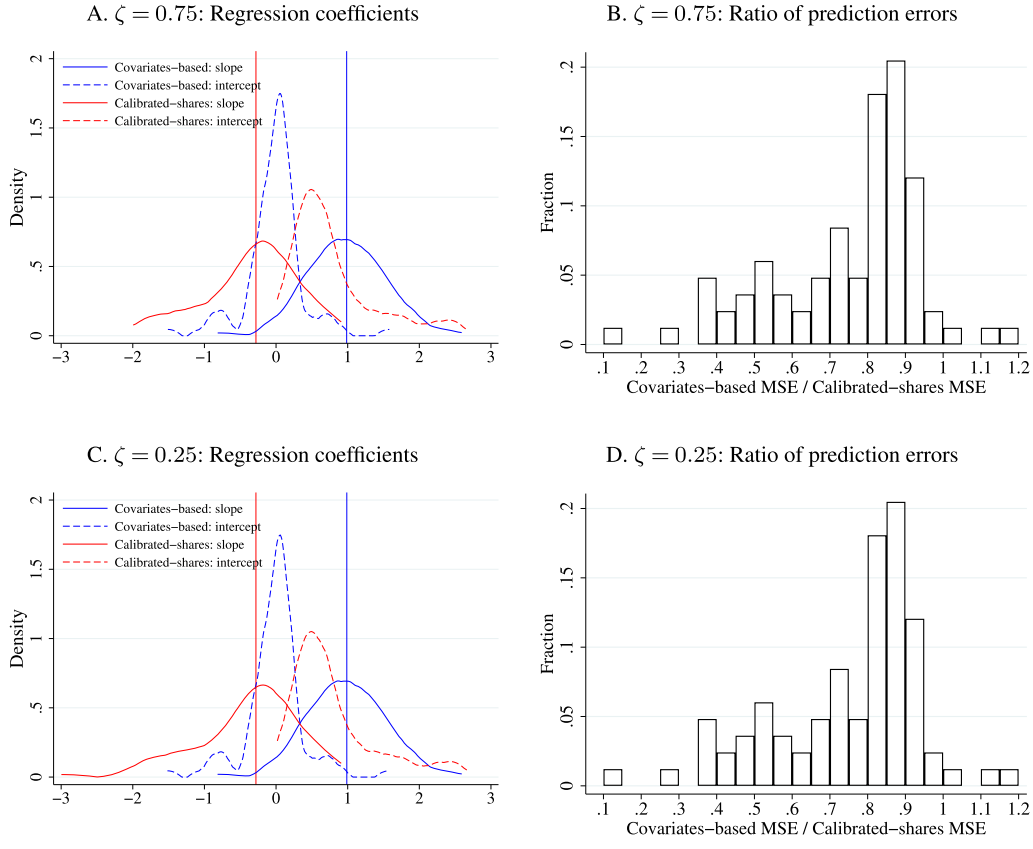

FIGURE B.8.—Comparison of methods' predictive performance with nested-logit ν^i. NOTES: The pairs of panels in this figure depict results analogous to those in Figure 3 for the nested-logit model introduced in this Appendix with nesting parameter $\zeta = 0.75$ (Panels A and B) and $\zeta = 0.25$ (Panels C and D). The results for $\zeta = 0.75$ and $\zeta = 0.25$ are not numerically identical, but they typically coincide to two decimal places and are visually indistinguishable.

the parameter ζ. As the figure shows, the results are robust to the alternative assumptions about preferences. Even when idiosyncratic preferences are highly correlated within residence NTAs, the covariates-based model considerably outperforms the calibrated-shares procedure.

B.8. *Alternative Labor Demand Elasticities*

We compare the predictive performance of the covariates-based model and the calibrated-shares procedure for varying values of σ, the local labor demand elasticity. For both $\sigma = 1.1$ and $\sigma = \infty$, the covariates-based model has a lower mean squared error than the calibrated-shares procedure in 80 of the 83 events.[51] The covariates-based model outperforms the calibrated-shares procedure regardless of the value of σ because their predictions are not very sensitive to σ. Varying σ does not meaningfully change the slope, intercept, or MSEs. As an example, Figure B.9 shows that the slope coefficients obtained

[51]Note that the specification with $\sigma = 1.1$ does not satisfy the sufficient condition for uniqueness.

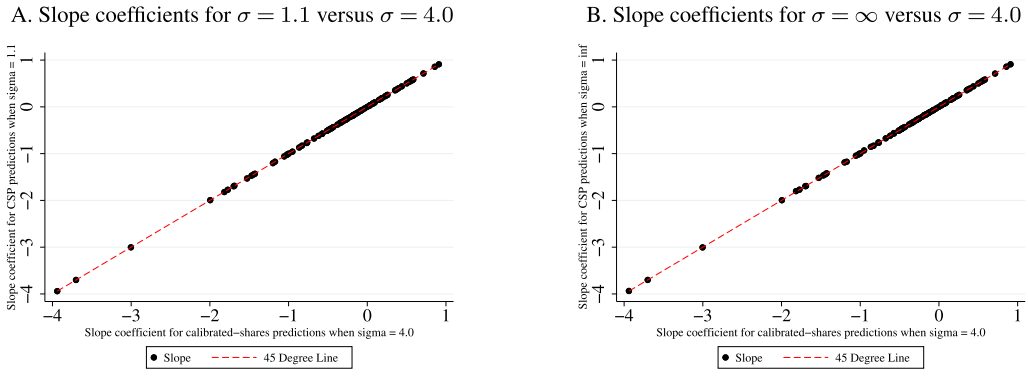

FIGURE B.9.—Event-study slope coefficients by local labor demand elasticity σ. NOTES: The slope coefficients come from regressions of the observed change in commuters on the predicted change in commuters for each of the 83 tract-level employment booms.

when using $\sigma = 1.1$ or $\sigma = \infty$ are very similar to those for $\sigma = 4.0$ in the calibrated-shares procedure.

B.9. *Local Increasing Returns to Scale*

This appendix section details the process for computing counterfactual predictions with local increasing returns and shows the results of such an exercise. In short, the predictions from both the covariates-based model and the calibrated-shares procedure are not sensitive to local increasing returns over the range of relevant values.

Per Appendix C.11.4, we assume that the increasing returns are external to the firm and that the production function in location n is given by $q_n = \bar{A}_n L_n^{\eta+1}$. This implies the goods-market-clearing condition is

$$\bar{A}_n \left(\sum_k \ell_{kn} / \bar{\delta}_{kn} \right)^{1+\eta} = \frac{w_n^{-\sigma} \bar{A}_n^{\sigma} \left(\sum_k \ell_{kn} / \bar{\delta}_{kn} \right)^{\eta\sigma}}{P^{1-\sigma}} Y.$$

This can be rearranged to solve for w_n:

$$w_n = \bar{A}_n^{\frac{\sigma-1}{\sigma}} \left(\sum_k \ell_{kn} / \bar{\delta}_{kn} \right)^{\frac{\sigma\eta-1-\eta}{\sigma}} P^{\frac{\sigma-1}{\sigma}} Y^{\frac{1}{\sigma}}.$$

To express the counterfactual outcomes via exact hat algebra, take ratios to obtain:

$$\hat{w}_n = \hat{\bar{A}}_n^{\frac{1}{\eta+1}} \left(\sum_k \hat{y}_{kn} \frac{y_{kn}}{\sum_{k'} y_{k'n}} \right)^{1-\frac{\sigma}{(\sigma-1)(\eta+1)}} \hat{P}^{\frac{1}{\eta+1}} \hat{Y}^{\frac{1}{(\sigma-1)(\eta+1)}}.$$

This expression for changes in wages generalizes equation (5) to the case of local increasing returns. Equations (6) and (7) are unchanged. The change in the CES price index

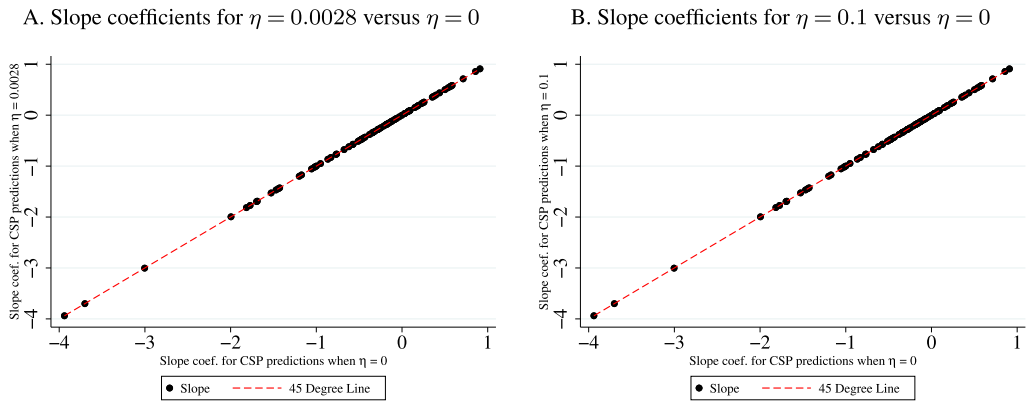

FIGURE B.10.—Event-study slope coefficients by local increasing returns η. NOTES: The slope coefficients come from regressions of the observed change in commuters on the predicted change in commuters for each of the 83 tract-level employment booms. The baseline case, $\eta = 0$, is the model without increasing returns.

is

$$\hat{P} = \left[\sum_n \left(\frac{\hat{w}_n}{\hat{\bar{A}}_n \hat{L}_n^{\eta}} \right)^{1-\sigma} \frac{\sum_k y_{kn}}{Y} \right]^{\frac{1}{1-\sigma}}.$$

In the event studies of 83 local employment booms, the covariates-based model outperforms the calibrated-shares procedure regardless of the value of η; varying η does not meaningfully change the slope, intercept, or MSEs. The differences in predicted commuting flows typically differ only on the order of 10^{-2} relative to the baseline, even when using values of η larger than the sufficient condition for the equilibrium uniqueness. As an example, Figure B.10 shows that the slope coefficients obtained when using $\eta = 0.0028$ or $\eta = 0.1$ are very similar to those for $\eta = 0$ in the calibrated-shares procedure.

REFERENCES

Adão, Rodrigo, Costas Arkolakis, and Federico Esposito (2022), "General Equilibrium Effects in Space: Theory and Measurement." Working paper. [0411,0438]

Adão, Rodrigo, Arnaud Costinot, and Dave Donaldson (2024), "Putting Quantitative Models to the Test: An Application to the U.S.-China Trade War." *The Quarterly Journal of Economics*, 140 (2), 1471–1524. [0411]

Ahlfeldt, Gabriel M., Stephen J. Redding, Daniel M. Sturm, and Nikolaus Wolf (2015), "The Economics of Density: Evidence From the Berlin Wall." *Econometrica*, 83 (6), 2127–2189. [0410,0411,0413,0415,0417,0430]

Allen, Treb, Costas Arkolakis, and Xiangliang Li (2016), "Optimal City Structure." Working paper. [0414]

Allen, Treb, Costas Arkolakis, and Xiangliang Li (2023), "On the Equilibrium Properties of Network Models With Heterogeneous Agents." Working Paper 27837, NBER. [0413]

Amazon (2018), "Amazon Selects New York City and Northern Virginia for New Headquarters." Press release. [0433]

Armenter, Roc and Miklós Koren (2014), "A Balls-and-Bins Model of Trade." *American Economic Review*, 104 (7), 2127–2151. [0411]

Arzaghi, Mohammad and J. Vernon Henderson (2008), "Networking off Madison Avenue." *Review of Economic Studies*, 75 (4), 1011–1038. [0410]

Bai, Jushan (2009), "Panel Data Models With Interactive Fixed Effects." *Econometrica*, 77 (4), 1229–1279. [0409,0427]

Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen (2014), "High-Dimensional Methods and Inference on Structural and Treatment Effects." *Journal of Economic Perspectives*, 28 (2), 29–50. [0421]

Berkes, Enrico and Ruben Gaetani (2020), "Income Segregation and the Rise of the Knowledge Economy." Working paper. [0433]

Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa (2019), "A Distributional Framework for Matched Employer Employee Data." *Econometrica*, 87 (3), 699–739. [0409]

Brinkman, Jeffrey and Jeffrey Lin (2022), "Freeway Revolts! The Quality of Life Effects of Highways." *The Review of Economics and Statistics*, 1–45. [0430]

Bryan, Gharad, Edward Glaeser, and Nick Tsivanidis (2020), "Cities in the Developing World." *Annual Review of Economics*, 12, 273–297. [0438]

Carvalho, Vasco M. and Basile Grassi (2019), "Large Firm Dynamics and the Business Cycle." *American Economic Review*, 109 (4), 1375–1425. [0411]

Chen, Mingli, Iván Fernández-Val, and Martin Weidner (2021), "Nonlinear Factor Models for Network and Panel Data." *Journal of Econometrics*, 220 (2), 296–324. [0409,0427,0453]

Cohen, Mike X. (2021), *Linear Algebra: Theory, Intuition, Code*. sincXpress. [0426,0449]

Conlon, Christopher, Julie Holland Mortimer, and Paul Sarkis (2023), "Estimating Preferences and Substitution Patterns From Second-Choice Data Alone." Working paper. [0425]

Correia, Sergio, Paulo Guimarães, and Tom Zylkin (2020), "Fast Poisson Estimation With High-Dimensional Fixed Effects." *The Stata Journal*, 20 (1), 95–115. [0417]

Costinot, Arnaud and Andrés Rodríguez-Clare (2014), "Trade Theory With Numbers: Quantifying the Consequences of Globalization." In *Handbook of International Economics*, Vol. 4, 197–261, Elsevier. Chapter 4. [0413]

[CDGH+] Couture, Victor, Jonathan I. Dingel, Allison E. Green, Jessie Handbury, and Kevin R. Williams (2022), "JUE Insight: Measuring Movement and Social Contact With Smartphone Data: A Real-Time Application to COVID-19." *Journal of Urban Economics*, 127. [0410]

Davis, Donald R., Jonathan I. Dingel, Joan Monras, and Eduardo Morales (2019), "How Segregated Is Urban Consumption?" *Journal of Political Economy*, 127 (4), 1684–1738. [0415,0416]

Davis, Morris A. and Francois Ortalo-Magne (2011), "Household Expenditures, Wages, Rents." *Review of Economic Dynamics*, 14 (2), 248–261. [0416]

Dekle, Robert, Jonathan Eaton, and Samuel Kortum (2008), "Global Rebalancing With Gravity: Measuring the Burden of Adjustment." *IMF Staff Papers*, 55 (3), 511–540. [0413]

di Giovanni, Julian, Andrei A. Levchenko, and Isabelle Mejean (2014), "Firms, Destinations, and Aggregate Fluctuations." *Econometrica*, 82 (4), 1303–1340. [0411,0432]

Diamond, Rebecca, Tim McQuade, and Franklin Qian (2019), "The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality: Evidence From San Francisco." *American Economic Review*, 109 (9), 3365–3394. [0430]

Dingel, Jonathan I. and Felix Tintelnot (2026), "Supplement to 'Spatial Economics for Granular Settings'." *Econometrica Supplemental Material*, 94, https://doi.org/10.3982/ECTA19350. [0409,0411]

Dingel, Jonathan I. and Felix Tintelnot (2021), "Spatial Economics for Granular Settings." Working paper. [0438]

Donaldson, Dave and Adam Storeygard (2016), "The View From Above: Applications of Satellite Data in Economics." *Journal of Economic Perspectives*, 30 (4), 171–198. [0410]

Eaton, Jonathan, Samuel Kortum, and Sebastian Sotelo (2013), "International Trade: Linking Micro and Macro." In *Advances in Economics and Econometrics: Tenth World Congress*, Vol. 3, Cambridge University Press. [0411,0432]

Eldén, Lars (2007), *Matrix Methods in Data Mining and Pattern Recognition*. Society for Industrial and Applied Mathematics. [0425]

Ellison, Glenn and Edward L. Glaeser (1997), "Geographic Concentration in U.S. Manufacturing Industries: A Dartboard Approach." *Journal of Political Economy*, 105 (5), 889–927. [0411]

Foschi, Andrea, Christopher L. House, Christian Proebsting, and Linda L. Tesar (2023), "Labor Mobility and Unemployment Over the Business Cycle." *AEA Papers and Proceedings*, 113, 590–596. [0415]

Gabaix, Xavier (2011), "The Granular Origins of Aggregate Fluctuations." *Econometrica*, 79 (3), 733–772. [0411,0429,0432]

Gandhi, Amit, Zhentong Lu, and Xiaoxia Shi (2023), "Estimating Demand for Differentiated Products With Zeroes in Market Share Data." *Quantitative Economics*, 14 (2), 381–418. [0411]

Gaubert, Cecile and Oleg Itskhoki (2021), "Granular Comparative Advantage." *Journal of Political Economy*, 129 (3). [0410,0411,0432]

Goodman, J. David (2019), "Why Amazon Is Caught in an Unexpected Brawl in New York." New York Times, 12 February. [0435]

Graham, Matthew R., Mark J. Kutzbach, and Brian McKenzie (2014), "Design Comparison of LODES and ACS Commuting Data Products." Working Papers 14–38, Center for Economic Studies, U.S. Census Bureau. [0415]

Guimarães, Paulo, Octávio Figueiredo, and Douglas Woodward (2003), "A Tractable Approach to the Firm Location Decision Problem." *The Review of Economics and Statistics*, 85 (1), 201–204. [0417]

Hastie, Trevor, Robert Tibshirani, and Jerome Friedman (2009), *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Vol. 2. Springer. [0421]

Hastie, Trevor, Robert Tibshirani, and Martin Wainwright (2015), *Statistical Learning With Sparsity: The Lasso and Generalizations*. CRC Press. [0447]

Heblich, Stephan, Stephen J. Redding, and Daniel M. Sturm (2020), "The Making of the Modern Metropolis: Evidence From London." *The Quarterly Journal of Economics*, 135 (4), 2059–2133. [0414,0430]

Hodrick, Robert J. and Edward C. Prescott (1997), "Postwar U.S. Business Cycles: An Empirical Investigation." *Journal of Money, Credit and Banking*, 29 (1), 1–16. [0409]

Holmes, Thomas J. and Holger Sieg (2015), *Structural Estimation in Urban Economics*, Vol. 5, 69–114. Elsevier. [0410]

[HNPS+] Hortaçsu, Ali, Olivia R. Natan, Hayden Parsley, Timothy Schwieg, and Kevin R. Williams (2023), "Demand Estimation With Infrequent Purchases and Small Market Sizes." *Quantitative Economics*, 14 (4), 1251–1294. [0411]

Hsiao, Allan (2022), "Coordination and Commitment in International Climate Action: Evidence From Palm Oil." Working paper. [0409]

Hsieh, Cho-Jui and Inderjit S. Dhillon (2011), "Fast Coordinate Descent Methods With Variable Selection for Non-Negative Matrix Factorization." In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1064–1072. [0450]

Kalouptsidi, Myrto (2014), "Time to Build and Fluctuations in Bulk Shipping." *American Economic Review*, 104 (2), 564–608. [0409]

Kehoe, Timothy J. (2005), "An Evaluation of the Performance of Applied General Equilibrium Models on the Impact of NAFTA." In *Frontiers in Applied General Equilibrium Modeling: in Honor of Herbert Scarf*, 341–377, Cambridge University Press. [0411,0438]

Kehoe, Timothy J., Pau S. Pujolàs, and Jack Rossbach (2017), "Quantitative Trade Models: Developments and Challenges." *Annual Review of Economics*, 9 (1), 295–325. [0411]

Krebs, Oliver and Michael Pflüger (2023), "On the Road (Again): Commuting and Local Employment Elasticities in Germany." *Regional Science and Urban Economics*, 99, 103874. [0414]

Kreindler, Gabriel and Yuhei Miyauchi (2023), "Measuring Commuting and Economic Activity Inside Cities With Cell Phone Records." *The Review of Economics and Statistics*, (WP2020-006), 1–48. [0410]

Krusell, Per and Anthony A. Smith (1998), "Income and Wealth Heterogeneity in the Macroeconomy." *Journal of Political Economy*, 106 (5), 867–896. [0429]

McFadden, Daniel L. (1974), "Conditional Logit Analysis of Qualitative Choice Behavior." In *Frontiers in Econometrics* (Paul Zarembka, ed.), 105–142, Academic Press. [0417]

McFadden, Daniel L. (1978), "Modelling the Choice of Residential Location." In *Spatial Interaction Theory and Planning Models*, North Holland. [0417]

Mogstad, Magne, Joseph P. Romano, Azeem M. Shaikh, and Daniel Wilhelm (2024), "Inference for Ranks With Applications to Mobility Across Neighbourhoods and Academic Achievement Across Countries." *The Review of Economic Studies*, 91 (1), 476–518. [0411]

Monte, Ferdinando, Stephen J. Redding, and Esteban Rossi-Hansberg (2018), "Commuting, Migration, and Local Employment Elasticities." *American Economic Review*, 108 (12), 3855–3890. [0411,0414,0416]

Owens, Raymond III, Esteban Rossi-Hansberg, and Pierre-Daniel Sarte (2020), "Rethinking Detroit." *American Economic Journal: Economic Policy*, 12 (2), 258–305. [0414,0415,0430]

Panigrahi, Piyush (2022), "Endogenous Spatial Production Networks: Quantitative Implications for Trade and Productivity." Working paper. [0411]

Perez-Cervantes, Fernando (2016), "Insurance Against Local Productivity Shocks: Evidence From Commuters in Mexico." Working Papers 2016–19, Banco de México. [0414]

Proost, Stef and Jacques-François Thisse (2019), "What Can Be Learned From Spatial Economics?" *Journal of Economic Literature*, 57 (3), 575–643. [0407]

Quan, Thomas W. and Kevin R. Williams (2018), "Product Variety, Across-Market Demand Heterogeneity, and the Value of Online Retail." *The RAND Journal of Economics*, 49 (4), 877–913. [0411]

Redding, Stephen J. and Esteban Rossi-Hansberg (2017), "Quantitative Spatial Economics." *Annual Review of Economics*, 9 (1), 21–58. [0407,0411,0430]

Rosenthal, Stuart S. and William C. Strange (2020), "How Close Is Close? The Spatial Reach of Agglomeration Economies." *Journal of Economic Perspectives*, 34 (3), 27–49. [0410]

Rossi-Hansberg, Esteban, Pierre-Daniel Sarte, and Raymond Owens (2010), "Housing Externalities." *Journal of Political Economy*, 118 (3), 485–535. [0410]

Rutherford, Thomas F. (1995), "Constant Elasticity of Substitution Functions: Some Hints and Useful Formulae." Notes prepared for GAMS General Equilibrium Workshop held December, 1995, in Boulder Colorado. [0414]

Sanders, Bas (2024), "Counterfactual Sensitivity in Quantitative Trade and Spatial Models." [0418]

Schoefer, Benjamin and Oren Ziv (2022), "Productivity, Place, and Plants." *The Review of Economics and Statistics*, 1–46. [0411]

Severen, Christopher (2023), "Commuting, Labor, and Housing Market Effects of Mass Transportation: Welfare and Identification." *The Review of Economics and Statistics*, 105 (5), 1073–1091. [0414,0415]

Silva, J. M. C. Santos and Silvana Tenreyro (2006), "The log of Gravity." *The Review of Economics and Statistics*, 88 (4), 641–658. [0417]

Sotelo, Sebastian (2019), "Practical Aspects of Implementing the Multinomial PML Estimator." Working paper. [0417]

Train, Kenneth (2009), *Discrete Choice Methods With Simulation*. Cambridge University Press. [0457]

Tsivanidis, Nick (2023), "Evaluating the Impact of Urban Transit Infrastructure: Evidence From Bogota's TransMilenio." Working paper. [0415]

Udell, Madeleine and Alex Townsend (2019), "Why Are Big Data Matrices Approximately Low Rank?" *SIAM Journal on Mathematics of Data Science*, 1 (1), 144–160. [0425]

Waddell, Sonya Ravindranath and Pierre Daniel Sarte (2016), "From Stylized to Quantitative Spatial Models of Cities." *Economic Quarterly*, 169–196. [0414]

Weintraub, Gabriel Y., C. Lanier Benkard, and Benjamin Van Roy (2008), "Markov Perfect Industry Dynamics With Many Firms." *Econometrica*, 76 (6), 1375–1411. [0429]

Zárate, Román David (2023), "Spatial Misallocation, Informality and Transit Improvements: Evidence From Mexico City." Working paper. [0415]

Co-editor Dave Donaldson handled this manuscript.

Manuscript received 11 January, 2021; final version accepted 11 December, 2025; available online 22 December, 2025.

The replication package for this paper is available at https://doi.org/10.5281/zenodo.17841577. The Journal checked the data and codes included in the package for their ability to reproduce the results in the paper and approved online appendices. Given the highly demanding nature of the algorithms, the reproducibility checks were run on a simplified version of the code, which is also available in the replication package.

All authors assume responsibility for all aspects of the paper.