

NBER WORKING PAPER SERIES

THE ECONOMICS OF URBAN DENSITY

Gilles Duranton
Diego Puga

Working Paper 27215
<http://www.nber.org/papers/w27215>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2020

Puga gratefully acknowledges funding from the European Research Council under the European Union's Horizon 2020 Programme (ERC Advanced Grant agreement 695107 – dynurban) and from Spain's Ministry of Science and Innovation (grants ECO2016-80411-P and PRX19-00578), as well as the support and hospitality of the Wharton School's Department of Real Estate during his visit as Judith C. and William G. Bollinger Visiting Professor. We are grateful to Gordon Hanson, Enrico Moretti, Frédéric Robert-Nicoud, Timothy Taylor, and Heidi Williams for very helpful feedback and to Yan Hu, Claudio Luccioletti and Giorgio Pietrabissa for research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research. The replication files for this article are available at <http://diegopuga.org/data/density/> in addition to the American Economic Association Data and Code Repository.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Gilles Duranton and Diego Puga. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Economics of Urban Density
Gilles Duranton and Diego Puga
NBER Working Paper No. 27215
May 2020
JEL No. R12,R31,R32

ABSTRACT

Urban density boosts productivity and innovation, improves access to goods and services, reduces typical travel distances, encourages energy-efficient construction and transport, and facilitates sharing scarce amenities. However, density is also synonymous with crowding, makes living and moving in cities more costly, and concentrates exposure to pollution and disease. We explore the appropriate measurement of density and describe how it is both a cause and a consequence of the evolution of cities. We then discuss whether and how policy should target density and why the trade-off between its pros and cons is unhappily resolved by market and political forces.

Gilles Duranton
The Wharton School
University of Pennsylvania
3620 Locust Walk
Philadelphia, PA 19104
and NBER
duranton@wharton.upenn.edu

Diego Puga
Centro de Estudios Monetarios y Financieros (CEMFI)
Casado del Alisal 5,
28014 Madrid
Spain
and CEPR
diego.puga@cemfi.es

A Code repository is available at <http://diegopuga.org/data/density/>

1. Introduction

Everybody loves density. Economists like to model and quantify the many benefits of urban density. It boosts productivity and innovation, improves access to goods and services, reduces travel needs, encourages more energy-efficient buildings and forms of transport, and allows broader sharing of scarce urban amenities. Other social scientists and urban planners, along with many policymakers, share this fondness for density and would like to see it increase in cities everywhere, including the densest ones.

We share some of that enthusiasm, but we also recognize that high density is synonymous with crowding. Indeed, there is a meaningful trade-off between the benefits and costs of density, and it is not clear that these benefits and costs are appropriately weighted by either market or political forces. One reason for this is that the benefit-cost calculation looks very different for insiders, long settled in the city, compared with outsiders considering moving in. In addition, the benefits and costs often operate at very different spatial and temporal scales, so they are not necessarily incorporated by all city constituents.

Understanding density is also tricky because density is both a cause and a consequence of the evolution of cities. Anything that makes a city relatively more attractive (such as a productivity increase or improved amenities) draws population from other places, which puts upward pressure on house prices, which in turn translates into higher land prices. Faced with a higher price per unit area of land, developers opt to build with a greater capital-to-land ratio (essentially, taller buildings). Faced with a higher price per unit area of floorspace, residents opt for smaller residences. With people living on smaller dwellings in taller buildings, density increases. In this sense, density is a consequence of urban evolutions.

At the same time, density is also a cause of many significant changes happening in cities. On the production side, agglomeration economies make firms and workers more productive in dense urban environments than in other locations. The benefits of density for innovation through spillovers are harder to measure but also deemed substantial. On the consumption side, higher density brings many goods and services closer, lowering travel needs. Changes in the amount and form of transport and more energy-efficient construction allow density to mitigate total pollution, albeit concentrating exposure to it. Historically, greater exposure to pollution and disease have been some of the greatest hazards of dense urban environments, and while they have lost our attention, they remain relevant today. These pitfalls, together with greater crowding and congestion, more costly floorspace for residents and firms, and scarcer green space, imply that density also has downsides. The combined benefits and costs of higher density also lead to changes in the composition of cities, triggering changes in the quality and variety of goods and services that are available — amenities in particular.

In this paper, we discuss what economic researchers have learned about density and what we see as the most significant gaps in this understanding. We begin by describing how economic research measures density for empirical enquiries and how appropriate measurement is rapidly changing with increasingly detailed data. We then explore the benefits and costs of density, how the trade-off between them is resolved, and the welfare effects of how market and political forces affect density.

2. Measuring Density

Population or employment density is often used as a summary statistic to describe the spatial concentration of economic activity. In this context, density is commonly defined as the number of individuals per unit geographic area. Such “naive density” is easy to calculate. However, it may not appropriately reflect the density actually faced by the individual or firm at hand.

One problem is that economic units are traditionally defined as aggregates of administrative units. For example, us metropolitan areas are defined based on counties, but if a metro area includes some counties with substantial rural portions, such a calculation will understate the density experienced by most economic actors. In particular, the match between urban and county boundaries is systematically looser for younger and less dense metropolitan areas in the West. An extreme example is the metropolitan area of Flagstaff, Arizona, which includes the second-largest county in the country and expands across multiple national parks, monuments, and forests.

Data have now become available with much finer geographical detail than in the past. Traditional data from statistical agencies, which were previously aggregated into fairly large and often arbitrary administrative units, are now provided at a much finer spatial resolution. For instance, the us Census Bureau now routinely releases information for more than 200,000 “block groups” instead of 3,000 counties. Also, data such as property prices or retail locations that were hard or expensive to obtain have become more broadly accessible in many countries. A variety of new digital and pictorial trails has also become available, from cellphone data tracking the location of people to high-resolution satellite imagery or street-level photography.

These newly available data offer research opportunities, but also raise three questions concerning: 1) choice of scale, 2) using a single “index” measure of density, and 3) the appropriate variable of interest for measuring density. Let’s discuss these in turn.

The first issue is that choosing the appropriate scale at which to measure density is specific to the particular question being raised. Some agglomeration mechanisms rely on direct human interactions, which in turn suggest that effective density should be measured at a small spatial scale. In this symposium, Rosenthal and Strange discuss the literature about agglomeration economies from short-distance interactions. The study of urban travel may require the measurement of density within a five- to ten-kilometer radius to capture the distance within which most errands take place (Duranton and Turner, 2018). In contrast, the metropolitan level may be relevant to measure broad-based agglomeration effects happening through local labor markets. The choice of scale does not stop at the level of metropolitan areas. Another thread of literature, inspired by Krugman (1991), has considered the much longer distances at which physical goods, and intermediate inputs in particular, can be traded. Given our urban focus, we leave aside the concentration of economic activity at a regional scale.

The choice of scale requires data on density and its effects to match. For example, De la Roca and Puga (2017) and Henderson, Kriticos, and Nigmatulina (2020) have proposed measuring “experienced density” by counting population within a given radius around each individual. De la Roca and Puga (2017) then average this measure across individuals in each city, given that they do not observe the exact location within the city of employers in their wage regression. Such

experienced density, in addition to dealing with the uneven tightness of area boundaries, better captures how close the typical individual is to other people when population is unevenly distributed. To give an example at the country level, where boundaries are given, the United States has nearly nine times the population of Canada with a slightly smaller surface area, so its naive density is ten times higher. And yet, walking around cities and towns in both countries, one likely perceives similar concentrations of people nearby. Indeed, the average inhabitant in Canada has about 343,000 people living within a ten-kilometer radius, compared with about 306,000 in the United States.¹

Instead of concentrating attention on the immediate neighborhood, a spatial decay factor giving more weight to closer neighbors may also be used. It is also possible to measure population density for fine spatial units and then to take a population-weighted average for larger units that match the dependent variable. Ciccone and Hall (1996) provide an early example of this approach. Their productivity measure is at the state level, but they compute employment density at the county level before taking an employment-weighted average by state. This weighting avoids distorting the calculation of density in large states like Texas where there are vast rural portions but the population is highly concentrated in a small number of counties.

It is tempting for researchers to define the appropriate measure scale or density as the one that yields the largest or most statistically significant coefficient in the regression of interest, either implicitly or explicitly in a horse race across different measures. This temptation should be resisted. The largest or most significant coefficient may also be the one suffering from the worst identification problems.

The second problem is that any standard density measure tries to summarize a two-dimensional distribution (individuals within an area) with a single index number. However, other “shapes” of density may matter, and alternative characteristics of cities beyond just their population and land area can now be measured at a reasonable data cost. Such characteristics include the number of centers and subcenters, the mixture of land use, the compactness of development, and more. In a study of cities in India, Harari (2020) find that such variables may affect a wide range of urban outcomes.

The third consideration involves choosing the variable of interest to use when measuring density. Following the pioneering work of Ciccone and Hall (1996), much of the literature that seeks to quantify the effects of the concentration of economic activity on productivity has focused on population and employment density — a choice driven mostly by the easy availability of these data. However, a case can be made that the density of human capital (Moretti, 2004) or the density of business activity in the same sector of economic activity (Henderson, 1974, Moretti, 2019) might be more relevant variables. Moreover, as we discuss below, empirically separating agglomeration effects within and across sectors remains largely an open empirical question.

Along with these three challenges of scale, the appropriate index number, and the appropriate

¹We calculate experienced density using 2010 gridded population data at 3 arc-second resolution from WorldPop (2018). We first measure the number of people within a ten kilometres radius of each cell in the population grid. We then compute, for all grid cells (in the entire country in this example, or in each city when we consider us metropolitan areas below), the population-weighted average of this count of people within ten kilometres. Weighting by population is important, since otherwise we would be calculating population within ten kilometers of the average place instead of within ten kilometers of the average person.

variable of interest, new sources of data on location and economic activity keep opening new possibilities for analysis of density. Traditional sources of population data typically measure population at its place of residence, which can work fine if the analysis is done at the metropolitan level and most people live and work in the same area. However, once we start trying to measure the effects of the concentration of economic activity at a fine spatial scale, a nighttime measure of density based on residences may not match well with a daytime measure of density based on employment.

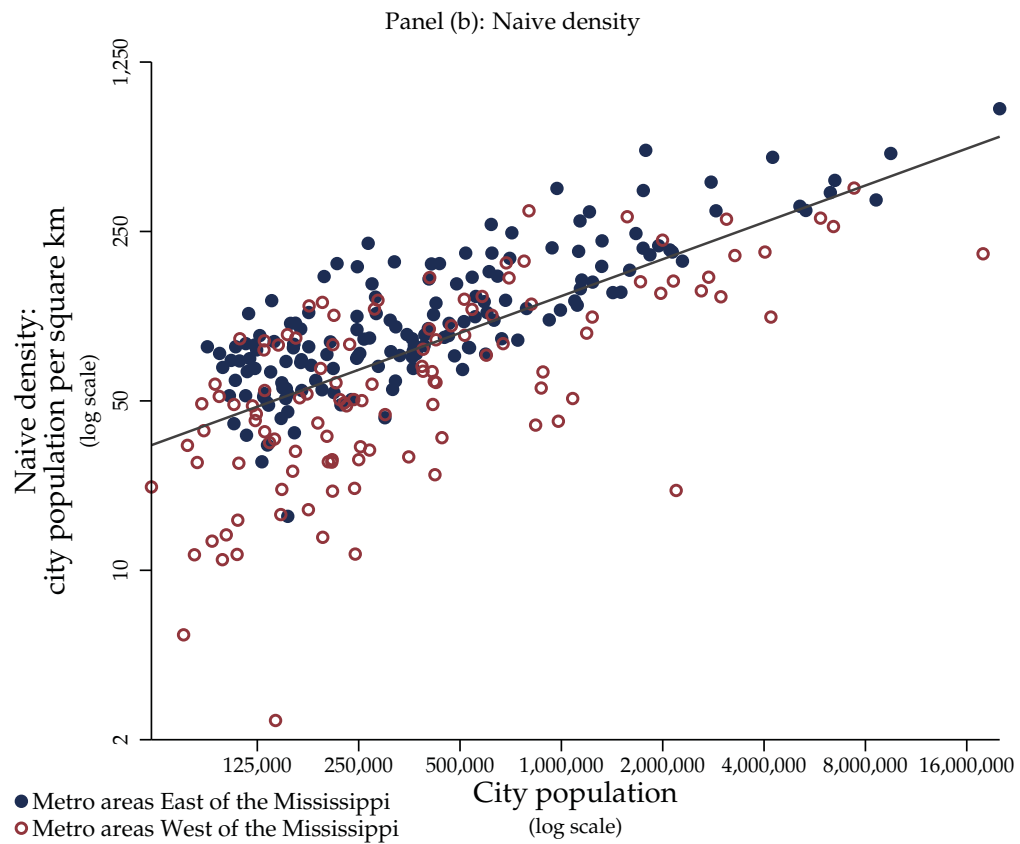

Cellphone data open the possibility of tracking people and measuring their location throughout the day (Kreindler and Miyauchi, 2019). Indeed, cellphone data even allow researchers to track those interactions directly, either by studying who is talking with whom on the phone (Büchel and Ehrlich, 2016, Büchel, Ehrlich, Puga, and Viladecans-Marsal, 2018) or by studying who is in the same building with whom at the same time (Atkin, Chen, and Popov, 2019).

Building-level data represent both day- and night-time density, and thus may offer a compromise. Daytime satellite imagery and, in some countries, official sources such as a land registry can provide detailed location data for individual buildings (Baragwanath-Vogel, Goldblatt, Hanson, and Khandelwal, 2020, Bellefon, Combes, Duranton, Gobillon, and Gorin, 2020). Information about built-up land is widely used to apportion population data measured at a broader scale to produce finely “gridded” population data (Leyk *et al.*, 2019). “Night-lights” satellite imagery offers the possibility of easier comparisons across countries and does not rely on the availability of more traditional administrative sources. However, it suffers from a range of measurement problems, notably the glow from bright sources of light — as discussed in this journal by Donaldson and Storeygard (2016). Building-level data, combined with population data, have the added benefit of being able to distinguish between population per unit of land area and population per unit of floorspace, which measures crowding more directly.

We have seen that most density measures either count individuals for comparable units or normalize this count by the geographical size of each unit. This raises a long-standing question: Should we measure the concentration of economic activity with its overall scale or its density? Induction suggests that, taken to extremes, neither density alone nor scale alone are particularly appealing. For instance, a highly concentrated but tiny cluster of economic activity is unlikely to generate strong agglomeration economies. On the other hand, workers located at the edge of large metropolitan areas are unlikely to benefit from their full scale in the job-matching process. The theoretical literature is mostly agnostic in the density versus scale debate. While the bulk of the work modeling the micro-foundations of agglomeration economies focuses on scale effects (Duranton and Puga, 2004), this is mostly a modeling choice, and it is easy to model agglomeration effects stemming from local density (Ciccone and Hall, 1996).

Empirically, the relationship between city density and city population is very tight, provided we measure density carefully enough. Panel (a) of Figure 1 plots for us metropolitan areas experienced density, measured as population within ten kilometers of the average resident, against total population. The implied elasticity is 0.51. If we use, instead, “naive density,” dividing total population by total land area within the official boundaries of the metropolitan areas, we find the same elasticity with respect to total population, 0.51, but the fit is poorer with an R^2 of 0.49 instead

Figure 1: Density vs. population for us metropolitan areas



of 0.76. This poorer fit is evident in panel (b), which also shows this is mostly because of artificially low densities in metropolitan areas with large rural portions in the western United States.²

3. The Benefits from Density

3.1 *Productivity Benefits*

Quantifying the productivity benefits from density has been a core theme in urban economics for several decades, and there is now broad consensus on their magnitude. The meta-analysis of Ahlfeldt and Pietrostefani (2019) suggests an elasticity of productivity with respect to density of 0.04, based on a citation-weighted average of estimates in the literature, on Combes, Duranton, Gobillon, Puga, and Roux (2012), and on an earlier meta-analysis (Melo, Graham, and Noland, 2009). Because the research on this topic has been reviewed carefully and extensively elsewhere (Rosenthal and Strange, 2004, Puga, 2010, Moretti, 2011, Combes and Gobillon, 2015), we focus this section on how such estimates are obtained and some recent developments.

Most estimates of the productive benefits from density are obtained by comparing productivity or earnings across spatial units with different densities. Early studies of productivity, starting with Sveikauskas (1975), studied average output per worker in cities. More recent studies rely on total factor productivity estimated from plant-level data to account for systematic differences in factor usage. In the case of earnings, firms must compare the wages they pay to the productivity benefits they receive when choosing a location. Both productivity and earnings are systematically higher in denser cities.

A concern when regressing productivity on density is that higher productivity in denser areas does not necessarily reflect a causal relationship. Instead, perhaps firms and workers are attracted to places with a strong but unobserved productivity advantage. Four strategies have been used to tackle this potential omitted variable problem. All of these approaches suggest that, while productivity-based sorting is a relevant concern, there is indeed a causal relationship in which greater urban density leads to higher productivity.

The first strategy uses instrumental variables when estimating the current density of an area. The usual instruments are historical measures of density (Ciccone and Hall, 1996) and land fertility (Combes, Duranton, Gobillon, and Roux, 2010). Both rely on differences in density being persistent over long periods, while the determinants of productivity have changed dramatically as the economy has evolved from being mostly agricultural to a concentration on manufacturing and then services. Another common instrument is land suitability for the construction of tall buildings (Rosenthal and Strange, 2008, Combes, Duranton, Gobillon, and Roux, 2010). A limitation of these approaches is that past populations and the nature of soils may affect current productivity through the persistence of productive infrastructure or the ease of building it.

²We calculate experienced density using 2010 gridded population data at 3 arc-second resolution from WorldPop (2018) as detailed in footnote 1 above. We calculate naive density using population and land area data from the 2010 US Census (US Census Bureau, 2011). To define cities, we use Metropolitan Statistical Area (MSA) and Consolidated Metropolitan Statistical Area (CMSA) definitions outside of New England and New England County Metropolitan Area (NECMA) definitions in New England, as set by the Office of Management and Budget on 30 June 1999. This defines 275 metropolitan areas in the conterminous United States.

A second strategy is to include either location or plant fixed effects in an attempt to capture any unobserved attributes that may have attracted more establishments to a given city (Henderson, 2003). Estimates are then identified from relating changes in productivity to changes in density over time, so the usefulness of this strategy is limited by the fact that relative changes in density tend to be small and slow (the same fact that makes the usual instruments relevant).

The third strategy is to find a quasi-experimental setting. For example, Greenstone, Hornbeck, and Moretti (2010) estimate changes in total factor productivity for incumbent plants in US counties that attracted a new plant investing over one million dollars. When compared to changes in total factor productivity for incumbent plants in runner-up counties that were being considered as an alternative location by the firm, the firm's final choice can be seen as an exogenous increase in density.

The final strategy is to impose more theoretical structure on the problem, as in Baum-Snow and Pavan (2012) or Ahlfeldt, Redding, Sturm, and Wolf (2015). The latter build a quantitative framework based on a canonical urban model and apply it to Berlin, Germany, as the Wall was built and then torn down.

Another important identification issue is sorting. Larger and denser metropolitan areas disproportionately attract more educated workers. While one can control for education and other observable characteristics, unobservable worker traits that affect productivity may differ systematically across cities. Here, following Glaeser and Maré (2001) and Combes, Duranton, and Gobillon (2008), the usual strategy is to introduce worker fixed-effects when relating individual earnings to density. The productivity benefits of density are then identified from the changes in earnings that a given worker experiences when changing work location.

Higher unobserved ability may be intrinsic to a worker due to natural talent or upbringing, but it may also be something that develops over time as the worker accumulates job experience. Separating the intrinsic and experience components of ability helps to evaluate the importance of sorting. It also allows us to study the extent to which the productive benefits of density can be absorbed almost immediately or instead accumulate gradually (Glaeser and Maré, 2001). De la Roca and Puga (2017) address this distinction by tracking the experience accumulated by workers in different locations. They then estimate an earnings regression where, in addition to incorporating worker fixed effects, they let the value of experience differ depending both on where it was acquired and where it is used. They conclude that workers across cities with different levels of density are not particularly different to start with; instead, working in different cities is mainly what makes their earnings diverge over time. They find that about one-half of the benefits of density are static and tied to currently working in a denser city. The other half accrues over time as workers accumulate more valuable experience in denser cities. Furthermore, workers take these dynamic gains with them when they relocate, which the authors interpret as evidence of important learning benefits to working in denser cities that get embedded in workers' human capital. These gains are stronger for those with higher initial ability.

Employing a similar strategy to look at the productivity of firms relative to density is difficult. Plant relocations are much less frequent than worker relocations. Which firms enter a market and which firms are able to survive may also be systematically different across more and less dense cities.

Combes, Duranton, Gobillon, Puga, and Roux (2012) develop a framework to distinguish between agglomeration and firm selection. The key insight of their model is that stronger selection in denser cities left-truncates the productivity distribution by removing the least productive firms. Stronger agglomeration instead right-shifts the productivity distribution by raising the productivity of all firms. If more productive firms benefit from density to a greater extent, this additionally dilates the productivity distribution. Using these insights, French establishment-level data, and a new quantile approach, Combes, Duranton, Gobillon, Puga, and Roux (2012) show that firm selection cannot explain spatial productivity differences. Instead, there are productivity benefits from density that are even greater for more productive firms. Gaubert (2018) argues that if, as shown by Combes, Duranton, Gobillon, Puga, and Roux (2012), more productive firms benefit more from density, they will also sort into denser environments to start with. Her results indicate that sorting reinforces agglomeration economies in explaining spatial productivity differences.

3.2 *Seeking the Sources of Productivity Benefits*

While urban economists broadly agree on the magnitude of the productivity benefits of density, the evidence distinguishing between possible sources is less solid. Duranton and Puga (2004) classify the mechanisms into three broad classes. First, a larger market allows for a more efficient *sharing* of local infrastructure, a variety of intermediate input suppliers, or a pool of workers. Second, a larger market also allows for better *matching* between employers and employees, or buyers and suppliers. Finally, a larger market can also facilitate *learning*, by facilitating the transmission and accumulation of skills or by promoting the development and adoption of new technologies and business practices.

On the empirical side, a widely used strategy to distinguish between mechanisms is to measure the geographical concentration of different sectors and regress this on proxies for the different mechanisms (Audretsch and Feldman, 1996, Rosenthal and Strange, 2001). Because plants in any given industry are similar in many dimensions, Ellison, Glaeser, and Kerr (2010) suggest instead studying which similarities across industries help to predict better co-agglomeration patterns. Both strategies rely on having good proxies for the underlying sources of agglomeration, and how one measures these can have an important effect on results. For instance, Overman and Puga (2010) suggest that, when measuring the importance of buyer-supplier relationships, one cannot just look at the value of input purchases, but instead should focus on purchases of crucial inputs whose production is geographically concentrated.

Instead of running a horse race between different agglomeration mechanisms, another possibility is to try to isolate a particular one. This approach is challenging because of behavior and outcomes that are difficult to track. Consider knowledge spillovers. Each of the links — like those from density to additional interactions, from interactions to information flows, and from information flows to innovation — is very difficult to trace and measure. In what has for a long time been arguably the best empirical evidence of knowledge spillovers. Jaffe, Trajtenberg, and Henderson (1993) show that an inventor of a patent is more likely to live in the same location as an inventor of a patent it cites than to live in same the location as an inventor of a similar matched patent it does not cite. However, as ingenious as this strategy is, it infers interactions from spatial proximity

and patents give only a very partial view of innovative activity. The strategy cannot show whether density increases interactions nor whether those interactions affect innovation more broadly.

One way to measure interactions is through survey data. Charlot and Duranton (2004^{b,a}) study the use of various communications technologies within and across firms in France. Their results suggest that dense urban environments result in more communication across firms and that greater communication results in higher wages, but they find little support for the hypothesis that density increases face-to-face interactions.

More recently, anonymized call detail records from cellphone operators allow measuring who interacts with whom on the phone. Büchel and Ehrlich (2016) use a major overhaul of public transport routes and schedules in Switzerland as a source of exogenous variation to show that proximity (measured by shorter travel times) does make interactions more likely. Interestingly, they find that, as in the model of Berliant, Reed, and Wang (2006), density facilitates meeting people, but this in turn makes people more choosy about with whom they interact. Thus, people in denser areas do not interact with more people but those with whom they interact are better matches. In addition, social networks in dense urban environments are less characterized by clustering into relatively isolated groups, likely facilitating more widespread information flows.

The idea that density facilitates the quality more than the quantity of matches is also present in labor markets. Dauth, Findeisen, Moretti, and Südekum (2018) use matched employer-employee data for Germany to show that high-quality workers (those who get high wages conditional on observables) are more likely to work for high-quality firms (those who pay high wages conditional on observables) in denser cities. This assortative matching reinforces the fact that high-quality workers and firms are also more likely to locate in denser cities.

Modern cellphones can also provide information on users' locations, gathered from the identifier of the cell tower providing coverage to the user (stored by cellphone operators) or from location data collected by smartphone apps (purchased, combined, processed and resold by several private companies). These data can measure spatial proximity of users at a fine geographical scale and within a narrow period: for example, two people spending more than 15 minutes in the same coffee place within the same clock-hour. Atkin, Chen, and Popov (2019) use such data to study how chance meetings contribute to innovation. They isolate smartphone users who work in buildings belonging to tech companies in Silicon Valley and trace instances where the users are in the same place at the same time. They separate chance from planned meetings and show that chance meetings result in more patent citations across firms in different sectors whose workers had met by chance more often. Other interesting, new sources of data starting to be used to track specific agglomeration mechanisms include detailed input-output links between firms (Bernard, Moxnes, and Saito, 2019) and search and matching in job platforms (Marinescu and Rathelot, 2018).

3.3 Accessibility Benefits from Density

All else equal, having the same population of residents and establishments in a smaller area will reduce bilateral distances. However, shorter bilateral distances may encourage more trips, and more trips within a compact area may also make travel slower. What is the net effect of these influences? Using us travel survey data, Duranton and Turner (2018) estimate an elasticity of the

distance traveled by an individual driver with respect to the density of workers and residents within a ten kilometer radius around the driver's residence of -0.13 , which occurs despite a very small increase in the number of trips by this driver. Travel speed also declines with density with elasticity -0.11 , but because of reduced distances, total time spent traveling declines. After considering many alternatives, Duranton and Turner (2018) conclude that the density of resident population and employment within a five or ten-kilometer radius is the main local characteristic explaining distances traveled by local residents. Looking at credit card records for shopping and the purchase of personal services, Agarwal, Jensen, and Monte (2019) also document a decline in travel associated with a greater density of sales locations. Couture (2016) finds similar results when focusing more narrowly on restaurants.³

But the accessibility benefits from density cannot be captured by transport costs alone. The variety, prices, and quality of available goods and venues will all change with density. In turn, these changes affect the choices made by consumers. Regarding prices and quality, Handbury and Weinstein (2015) find that larger (and thus denser) cities do not have significantly different prices for the exact same grocery products. If prices for a certain type of product tend to be higher in large metropolitan areas, it is because consumers tend to buy higher quality varieties of the same product — like organic instead of regular eggs.

Regarding variety, Handbury and Weinstein (2015) find that the availability of grocery products, measured at the bar-code level, is much greater in larger cities. The count of restaurants accessible within a given travel time also increases with density. To assess the benefits of this expanded variety, Couture (2016) estimates the elasticity of substitution between restaurants, where a lower elasticity implies a greater willingness to pass many restaurants to access one's preferred choice. Couture's (2016) estimate of about nine for this elasticity is larger than the usual estimates for consumer goods but low enough to generate frequent trips beyond the closest restaurant and substantial welfare gains from restaurant density.

Couture's (2016) work provides an important bridge with earlier transportation research that attempts to model accessibility within a discrete choice framework (following the influential research of Ben-Akiva and Lerman, 1985). After making some distributional assumptions about the preference parameters, it becomes possible in this framework to recover accessibility for a given location from the consumers' choice set of destinations in this location and the costs of reaching these destinations. For many years, the application of this approach was limited by the paucity of data about the possible destinations and the cost of reaching them. New data from sources like Google Place and Google Maps have eased these constraints. However, one limitation of these accessibility measures is that they take as given both the location of the origin of trips and the set of destinations, whereas density matters partly because it changes the set of potential venues and, as a result, possibly alters the choice of residential location.

Another strand of the research literature provides full general equilibrium models which consider an explicit geography and can be quantified to estimate policy counterfactuals. Redding and

³While household travel for consumption is "local," it is not "extremely local." Using Yelp data for restaurant visits in New York City, Davis, Dingel, Monras, and Morales (2019) find that consumption is, by their metrics, about half as segregated as residences. Although transport frictions matter, they nonetheless find that social frictions play a bigger role.

Rossi-Hansberg (2017) provide an excellent guide to this literature. Among models that consider urban space and transport, Ahlfeldt, Redding, Sturm, and Wolf (2015) is a particularly accomplished contribution. They model the development of a city where residents choose their residence and workplace locations and use this to explore the benefits of density using historical variation in accessibility due to the Berlin Wall. These location choices are, to a large extent, guided by utility shocks over particular commuting routes. While this approach greatly simplifies the derivation of their model, it is limiting for current-day applications to the extent that commutes represent less than one-fifth of all trips and about one-fourth of the mileage for US drivers. A challenge for the future will be to harness the recent modeling advances in economic geography, while both keeping the versatility and ease of implementation of standard discrete choice approaches used in the transport literature and also making use of the much richer data now available to study urban travel.

Although we have discussed the productivity and accessibility benefits of density separately, they are interrelated. For example, the better accessibility of a denser urban environment may allow workers to search for better labor market opportunities (Manning and Petrongolo, 2017). In a study of relocating research and development establishments, Xiao and Wu (2020) find that researchers who end up with longer commutes to their workplace see a drop in patenting activity while those who get closer become more productive.⁴ One possible way to integrate the productivity benefits from agglomeration into a transport and accessibility framework is to compute density in a location as the sum of nearby employment discounted by the travel cost of accessing it. This follows the suggestion of Graham (2007) and is related to the gravity specifications of the spatial quantitative models reviewed by Redding and Rossi-Hansberg (2017). However, it is important to consider such density elasticities with care. For example, taking an elasticity of earnings with respect to density estimated from cross-city variation and then applying it to a change in “effective density” resulting from some expansion of transport infrastructure in one city may overestimate the actual gains from the project — for instance, if we are considering a new subway line while agglomeration benefits arise from input-output relations between firms unaffected by this line.

3.4 Other Benefits from Density: Innovation, Reduced Pollution, Amenities

For brevity, we limit our discussion here to three especially important benefits of density that seem to us ripe for additional study: innovation, reduced pollution, and access to amenities. For a literature review that includes other benefits, see Ahlfeldt and Pietrostefani (2019).

The extreme concentration of innovative activity is reviewed in Carlino and Kerr (2015). Moretti (2019) estimates an elasticity of the number of patents per innovator with respect to the number of innovators in the same city and field of innovation of about 0.07. This estimate is arguably a lower bound: for example, it ignores the effect of the concentration of innovators in the same field on the probability of innovating, and it ignores spillovers arising from other fields of research. In a prior

⁴In a very different context, Koster, Pasidis, and van Ommeren (2019) provide evidence about shopping externalities mediated by foot traffic. These shopping externalities are arguably about accessibility, since they arise from transport savings for customers when visiting multiple stores, but they end up affecting the productivity of stores as reflected in the rents they are willing to pay.

paper, Carlino, Chatterjee, and Hunt (2007) find a large elasticity of patenting with respect to urban density of about 0.20, reflecting both the higher productivity of research in denser places and the concentration of research inputs in these areas.

The link between density and pollution is also of particular importance. Residents in denser cities emit less greenhouse gas and fewer particulates than less dense cities (Glaeser and Kahn, 2008). This result is only in part due to transport. There are large differences in emissions related to home cooling, even after conditioning out climatic differences. However, we need to know how much of the lower energy consumption in denser places is a consequence of smaller dwellings or if there is something uniquely energy-efficient about greater density. At the same time, as Carozzi and Roth (2019) note, a higher concentration of population within a city may result in greater overall exposure to pollution, even with lower emissions per person. After instrumenting for urban density, they find an elasticity of exposure to particulates (2.5 micrometers or smaller) with respect to density of 0.13 for the United States.

The presence of consumption amenities in dense urban areas influences how one perceives the rising inequality of wages. The increased concentration of educated workers in a small number of increasingly attractive cities is a salient feature of the urban geography of the US and arguably of many other developed countries (Berry and Glaeser, 2005). If living in a dense area offers mainly negative “amenities,” like crime, then the increased concentration of skilled workers in increasingly expensive cities implies that inequality is less than wages suggest. If living in a dense area offers positive amenities, then inequality will exceed what wages suggest (Moretti, 2013). Diamond (2016) argues that, while rising skill premiums started the process of educated workers concentrating in dense urban settings, their presence then generated additional endogenous amenities, which she argues are central to reconciling observed changes in wages, rents, and the skill composition of residents across cities in the US between 1980 and 2000. In highly granular empirical work, Couture and Handbury (2019) provide direct evidence about the importance of local amenities to explain the return of young educated workers to higher density residential areas in American cities. In turn, the increased concentration of educated workers appears to foster the development of new local amenities. This recent work is in tension with more traditional estimations of the relationship between amenities and city size building on Roback (1982), which suggest only a weak relationship between city population and amenities (Albouy, 2008). Better knowledge about the formation of amenities in cities is undoubtedly a priority.

4. The Costs of Density

4.1 Land Prices, Housing Prices, Transport Costs, Congestion

Theory has long hypothesized that as population and density increase in a city, its benefits initially accumulate faster but eventually, its costs dominate (Henderson, 1974). Fujita and Thisse (2013) call this the “fundamental trade-off of spatial economics,” because it explains both the existence of cities and their finite sizes. However, compared to research on benefits of density, there is paucity of research on its costs, which Glaeser (2011) dubbed the “demons of density.”

As a starting point, density brings an increase in land prices. For French urban areas, Combes, Duranton, and Gobillon (2019) estimate an elasticity of land prices at the city center with respect to their population of about 0.30. These comparisons across cities are made for a central location to make sure we compare likes with likes. In and of itself, a higher price for land does not represent a cost for society, but more expensive land elicits various responses. Some of these responses do create social costs, such as the use of more expensive building technologies to build higher or longer and slower trips as residents move further out and roads get more congested. Let's explore these costs in turn.

More costly land provides incentives to build taller. Ahlfeldt and McMillen (2018) estimate an elasticity of building height with respect to land prices of 0.30 for residential buildings and 0.45 for commercial buildings in the city of Chicago circa 2000. Interestingly, this elasticity about doubled over 100 years, as technology made it easier to respond to high land costs by building taller. They note, however, that the elasticity of built-up floorspace with respect to land prices is only about one-third of the elasticity of building height because taller buildings are often surrounded by less tall buildings, open space, and roadway.

While taller buildings do provide more floorspace per unit of land, the marginal cost of floorspace increases with building height. Ahlfeldt and McMillen (2018) estimate elasticities of building cost per unit of floorspace with respect to building height ranging from 0.25 for small buildings to well above unity for skyscrapers. Tall buildings are not only costly to build; they also generate a range of recurring costs, including direct costs to their users. For example, Liu, Rosenthal, and Strange (2018) report that a typical tenant in a high-rise spends 23 minutes a day waiting for or riding elevators — about the same time as a typical one-way commute to work.⁵

A higher built-up density alleviates, but does not eliminate, the pressure created by higher land prices on the price of residences and offices. For French urban areas, Combes, Duranton, and Gobillon (2019) estimate an elasticity of housing prices at the city center with respect to their population of 0.11, compared with the aforementioned 0.30 for land.⁶ For US metropolitan areas, Duranton and Puga (2019) estimate a slightly lower elasticity of housing rents at the center of 0.07. Overall, higher demand for land at particularly desirable locations leads to an increase in floorspace density, higher prices for land and floorspace, and lower consumption of floorspace per person. These forces push towards an increase in human density per unit of land. Earlier, we provided an estimate of the elasticity of density with respect to population for US metropolitan areas of 0.51. In addition to lowering their housing consumption, residents also react to higher housing prices by moving to cheaper, less-accessible locations. When we estimate the elasticity of average distance to

⁵Liu, Rosenthal, and Strange (2018) also report some countervailing benefits. For instance, the top floors in tall buildings command higher rents than all but the street-level, suggesting that poorer accessibility relative to lower floors is more than offset by better views and perhaps more prestige.

⁶In their model, Combes, Duranton, and Gobillon (2019) show that the ratio of the land price elasticity to the housing price elasticity should be equal to the share of land in construction. For France, this ratio of $0.11/0.30 \approx 0.37$ is very close to the share of land in the construction of single-family homes.

the center with respect to city population, we get 0.30.⁷

The trade-off between housing costs and transport costs has been at the heart of land-use modeling since the pioneering work of Alonso (1964), Muth (1969), and Mills (1967). However, this early work used a monocentric model of cities, which both captures many essential features of actual cities and also has important shortcomings. Most notably, residents in the basic monocentric model only travel to commute to their job, and they all work at the center. However, because not all travel is travel to work and not all commutes reach the center of cities, average travel increases with a city's population far less than predicted by the monocentric model. Using transport data for US metropolitan areas, Duranton and Puga (2019) estimate that the elasticity of vehicle kilometers traveled with respect to the distance to the city of a resident household is only about 0.07. However, one key property of the monocentric model continues to hold. As households consider residences further away from the center, the lower price of the housing should be just offset by higher transport cost. Indeed, Duranton and Puga (2019) find that, just as predicted by the model, the elasticity of housing prices with respect to distance to the center is exactly the same as the elasticity of transport costs with respect to distance to the center, but with an opposite sign.

This literature suggests that cities that allow their urban fringe to expand may have more success in containing urban costs. Combes, Duranton, and Gobillon (2019) estimate the elasticity of land and housing prices at the center of French metropolitan areas with respect to either their density or their population. For housing prices, they estimate a density elasticity of 0.21 and a population elasticity of 0.11. For land prices, the density elasticity is 0.60 and the population elasticity 0.30. Since an increase in density is essentially an increase in population keeping built-up area constant, these differences indicate that if cities could only increase their population by increasing density, house prices would increase by twice as much in the long run, with even more substantial short-term price hikes.

As a city both gets denser and expands outwards, population growth also puts a strain on its infrastructure, and particularly its transport infrastructure. Urban travel gets slower as congestion worsens. Duranton and Puga (2019) estimate an elasticity of travel speed with respect to city population of -0.04 for US metropolitan areas using travel survey information. For cities in India, Akbar, Couture, Duranton, and Storeygard (2019) estimate the same elasticity using travel time data from Google Maps and obtain a similar figure of -0.05 .

When discussing the benefits of density, we included some endogenous changes in amenities, such as more consumption opportunities. Other amenity changes associated with density may instead constitute a cost. For instance, Glaeser and Sacerdote (1999) estimate that the elasticity of crime with respect to population for US cities is 0.16 if one focuses only on reported crime and 0.24 when one takes into account greater crime underreporting in larger cities. They find that, while the higher prevalence of crime-prone individuals in large cities plays an important role, almost

⁷To estimate the elasticity of average distance to the center with respect to city population, we first determine the location of the center of each metropolitan area from the location of its core municipality reported by Google Maps. We then compute for each metropolitan area, the population-weighted average distance to the center of its Census block groups using five-year 2008–2012 data from the 2012 American Community Survey obtained from the IPUMS-NHGIS project (Manson, Schroeder, Riper, and Ruggles, 2019). We finally regress the log of average distance on the log of population across metropolitan areas using ordinary least squares.

as important is the fact that higher urban density makes finding a victim for opportunistic crimes easier and catching criminals more difficult. However, it is intriguing to note that in Europe, larger cities tend to suffer less crime (Ahlfeldt and Pietrostefani, 2019).

We discussed earlier that density can also increase exposure to pollution from particulates, negatively affecting health. Historically, high density was also synonymous with frequent premature deaths caused by the poor hygiene of cities and the ease at which epidemics would propagate. Bairoch (1988) reports that early in the nineteenth century, rural youth were expected to live eight to twelve years longer than urban youth. In Europe and North America, urban life expectancy only overtook rural life expectancy after 1930. Urban planners tried to alleviate the burden of disease in cities not only by investing in water and sewage systems but also by building wider avenues and large urban parks and introducing regulations that limited overcrowding and improved air circulation and access to natural light (Colomina, 2019). If cities are not even denser today, it is partly a consequence of past diseases. And yet the lack of social distancing that cities promote — and which gives them so many advantages — also makes them more vulnerable to pandemics even today.

While the literature on urban costs remains limited, it offers three tentative conclusions. First, the various elasticities reported here provide support for a hill-shaped relationship between the net benefits of cities and their population scale and density. Second, the top of this hill is fairly flat, so that the costs of being moderately undersized or oversized are small. Finally, the downward-sloping part of the net benefits may eventually fall steeply and more so if cities cannot adjust at both the intensive (densification) and extensive (outwards expansion) margins.

4.2 Aggregating the Costs of Density and Population

Quantifying urban costs, in all their forms, is complicated. As one example, the multiple components of commuting costs are hard to observe and even harder to value (Small, 2012). Housing costs, despite being transfers from users to owners, are also expected to capitalize many other costs. Moreover, housing and transport costs vary across locations in a city.

To assess overall urban costs, the literature has developed three approaches. A first strategy uses a standard urban model that also includes agglomeration benefits. For example, Au and Henderson (2006) solve such a model to obtain an expression for average value added in a city as a function of its population. They also estimate the relationship between value added per worker and city population for Chinese cities during a period in which migration was greatly restricted and conclude that many of these cities were grossly undersized. The great advantage of this approach is that it requires little data — essentially just population and value added. However, it also has several drawbacks. The key fundamental relationship between agglomeration benefits and urban costs is expected to be hill-shaped, and the shape of the hill will be hard to estimate unless many cities are far from their optimal size, as in China in the 1990s. Also, it is unclear which urban costs are reflected in lower value added (for example, commuting costs paid in the time of workers will be missed, while the higher market activity of transit firms in congested cities may result in higher value added).

The second approach models the choices of a consumer who needs to decide on a residential location and asks how much more costly it would be to achieve the same level of utility at the same location should the city become denser or grow in population. Using this approach, Combes, Duranton, and Gobillon (2019) assume that households have free mobility and leverage the insight that house prices will capitalize transport costs and amenities. As a result, the elasticity of urban costs with respect to city population turns out to be equal to the elasticity of house prices at the center of cities with respect to their population multiplied by the share of housing in household expenditure. As mentioned earlier, Combes, Duranton, and Gobillon (2019) also estimate that the elasticity of housing prices with respect to city density is 0.11 and fairly stable over the range of city populations observed in France. The share of income devoted to housing increases with urban population, from about 16 percent in a city with 100,000 inhabitants to 39 percent in a city like Paris. Taken together, these figures are indicative of urban cost elasticities associated with a greater density ranging from 0.03 for smaller cities to 0.08 for cities with more than ten million inhabitants. The main drawbacks of this approach are that it relies heavily on the free-mobility condition to simplify a wide array of changes associated with greater density, that it considers only monetary costs, and that it ignores endogenous amenities (whether positive or negative).

A third approach, developed by Duranton and Puga (2019), models the various costs of cities explicitly and estimates the parameters associated with these costs.⁸ One advantage of this approach is that the key urban costs elasticity can be estimated based on equations of the model at three different levels of aggregation and using three different sources of variation. These approaches amount to estimating the assumed commuting cost equation (using within-city variation in travel distance across individuals), the spatial equilibrium within each city (using within-city variation in house prices across locations), and the spatial equilibrium across cities (using cross-city variation in city-center house prices). All three approaches result in a similar elasticity of urban costs with respect to city population of about 0.07. These urban costs are then further amplified by congestion with a population elasticity that they estimate at about 0.04. The main drawback relative to the previous approaches is that the modeling and data demands are even greater.

5. Getting Closer to Optimal Density?

5.1 *The Unhappy Welfare Economics of Density*

When considering the benefits and costs from density in a location, firms and workers choose based on their private benefits and costs, not on the social benefits and costs. There are two wedges between private and social that tend to push toward suboptimally low levels of density. The agglomeration wedge refers to the fact that firms and workers consider the agglomeration spillovers they may receive from others nearby, but not the agglomeration spillovers they may provide to others. Another wedge arises from the capitalization of land prices. When the land is not owned by local residents, a fraction of the net benefits from density are transferred away as rents to absentee

⁸Desmet and Rossi-Hansberg (2013) propose a related approach with a quantified model. The model is then used to assess the effects of shutting down various forms of heterogeneity across cities, rather than exploring the costs and benefits of increased density or rising population.

landowners who benefit from agglomeration without contributing to it. On the other side, there is a congestion wedge pushing toward suboptimally high levels of density because the marginal cost of congestion exceeds its average.

The overall effect of these three wedges is ambiguous. We did report above that with respect to density, the congestion elasticity is estimated to be higher than the agglomeration elasticity. However, the smaller agglomeration elasticity pertains to the labor income of residents choosing a location, whereas the larger congestion elasticity pertains to their travel costs, which are much smaller than labor income. Much less is known about the wedge from land capitalization. Here, the key issue is not who owns the land, but whether agents making decisions about local density bear the full costs and benefits of such decisions. To complicate the welfare economics problem further, the development of high density over a sufficient spatial scale is subject to “all-or-nothing” decisions: That is, no firm may want to move to a newly-developed location unless other firms are expected to move as well. Large-scale development also often requires coordination among developers (Henderson, 1974, Henderson and Mitra, 1996), but the market for large-scale development is absent or limited in most of the world and it remains limited in the United States. This coordination failure implies that there might be too few communities and, as a result, they may be overly dense.

Putting together the near-absence of a market to provide density at scale and the various externalities associated with location decisions, it seems unlikely that the factors will precisely counterbalance each other in ways that cause market forces to provide optimal density.

5.2 *The Unhappy Politics of Density*

Almost everywhere in the world, land use is heavily regulated with a view to determine overall density as well as specific types of density (of people, jobs, shops, green space, etc.).⁹ A commonly heard criticism is that land-use policies tend to deliver suboptimally low levels of density. For example, many land-use policies aim to reduce densities through instruments such as minimum lot sizes, maximum floor-area ratios, or single-family residential zoning. Such policies are particularly prevalent in the United States, where land zoned for detached single-family homes accounts for 94 percent of all land zoned for residential use in San José, 81 percent in Seattle, 75 percent in Los Angeles, and 70 percent in Minneapolis, although only 36 percent in Washington DC, and 15 percent in New York (Badger and Bui, 2019).

Many reasons have been suggested for restrictive zoning: 1) a fear by the rich that poorer residents will free-ride on public amenities (in particular, high-quality public schools) by consuming a small quantity of housing in a rich jurisdiction (Tiebout, 1956, Fischel, 1987); 2) a fear by risk-averse incumbents that less restrictive zoning would harm property values (Fischel, 2001); 3) the possibility that costs of increasing density are more short-term and highly local, while the benefits may take more time to accrue and diffuse across the metropolitan area; 4) when there are strong preferences for particular locations, incumbent residents can act as monopolies restricting entry (Ortalo-Magné

⁹Given our focus, we do not discuss policy interventions that try to get firms or people to relocate over large distances, even if the density of origin and destination can be quite different. For a starting point to this work, see the papers on place-based policies in this issue, including Bartik on us policies and Overman and von Ehrlich on European policies. See Duranton and Venables (2018) for detailed discussions of place-based policies in developing countries.

and Prat, 2011, Hilber and Robert-Nicoud, 2013); and 5) incumbent residents seeking to limit entry into particularly productive cities, thus maximizing their own welfare at the expense of potential newcomers (for a model, see Duranton and Puga, 2019).

Overall, the main cost of overly restrictive land-use regulations for society may result from a spatial misallocation of population. Using very different models to quantify the social losses from excessive regulation, Hsieh and Moretti (2019) and Duranton and Puga (2019) both suggest that relaxing planning regulations in the three most productive US cities to the median level might generate large aggregate real gains of about 8 percent. How much of these gains can be realized would depend greatly on how rapidly urban costs increase as some cities grow well beyond their currently observed sizes. Nevertheless, these quantitative assessments strongly indicate that observed densities are far from optimal — to low in some places and too high in others.

6. Conclusion

Over the last ten years, the study of urban density has been revitalized by the arrival of new fine-grained data. We are increasingly able to observe key links such as face-to-face interactions for learning spillovers. Granular data about job searches and matching in cities or trades between firms are also increasingly available. Significant modeling advances have also taken place during the last decade. A new generation of general-equilibrium urban models has come of age, and their main novelty lies in their ability to handle the heterogeneity we observe in the distribution of jobs and residences. New models have been developed to distinguish between the agglomeration, selection, and sorting effects of density; model job changes within and between cities; provide better estimates of the costs of density; and so on. There is less to report on the front of causal identification during the last ten years. There has been a lot of empirical work around the issues surrounding urban density. However, it pushed more-or-less along the same lines as previously, with a continued emphasis on instrumental-variable estimations, the use of difference-in-difference after a plausibly exogenous shock, and the exploitation of spatial discontinuities.

Thus, as we look forward to future progress on the economics of urban density, our wish list includes novel data explorations providing a richer set of facts related to the manifestations of density, models that integrate urban mobility and consider the dynamics of buildings and construction, and rising empirical standards in the identification of causal effects.

As we read one last time the preprint version of this article while confined in response to the COVID-19 pandemic, the costs we incur and the benefits we receive by seeking proximity during normal times in dense urban environments have become even more prominent. The streets are free from congested vehicle traffic and the sky is unusually clear from pollution. At the same time, we miss the ideas that often arise from serendipitous encounters with our colleagues and the concentration and sanity of separate office and home environments. For many, the sudden drop in economic activity has brought much deeper troubles.

Beyond the temporary quietness, the immediate prominence of the costs and benefits of density, and the impact of the emerging economic crisis, what will be the long-run consequences of this virus for our densest cities? Pandemics have hit cities the hardest for centuries, and cities have adapted

and been shaped by them — from investments in water and sewage systems to prevent cholera, to urban planning to reduce overcrowding and improve air circulation and access to sunlight in response to tuberculosis. Maybe temporary social distancing measures will also leave a permanent footprint on cities, for instance in the form of more space for pedestrians and bicycles or a gain of outdoor versus indoor leisure environments. But the idea that this pandemic will change cities forever is likely an overstretch. Cities are full of inertia and this crisis has stressed both the costs and benefits of density. Confinement is forcing us to see both the advantages and the great limitations of online meetings relative to the more subtle and unplanned in-person exchanges. It has made us realize that many tasks are impossible to do from home. At schools and universities, the haphazard transition to online courses may speed up their development, or it may delay it as many students have become frustrated by losing aspects of a full educational experience. For a while, some people may try to avoid dense cities for fear of contagion, but others may be drawn to them seeking work opportunities in difficult times. Perhaps one persisting lesson is that the cost of the pandemic has so far been associated more with urban inequalities than with urban density. While the consequences are hardest for lower-income households and minorities, they affect us all in profound ways.

References

- Agarwal, Sumit, J. Bradford Jensen, and Ferdinando Monte. 2019. The geography of consumption. Processed, Georgetown University.
- Ahlfeldt, Gabriel, Stephen J. Redding, Daniel Sturm, and Nikolaus Wolf. 2015. The economics of density: Evidence from the Berlin wall. *Econometrica* 83(6): 2127–2189.
- Ahlfeldt, Gabriel M. and Daniel P. McMillen. 2018. Tall buildings and land values: Height and construction cost elasticities in Chicago, 1870–2010. *Review of Economics and Statistics* 100(5): 861–875.
- Ahlfeldt, Gabriel M. and Elisabetta Pietrostefani. 2019. The economic effects of density: A synthesis. *Journal of Urban Economics* 111: 93–107.
- Akbar, Prottoy A., Victor Couture, Gilles Duranton, and Adam Storeygard. 2019. Mobility and congestion in urban India. Processed, University of Pennsylvania.
- Albouy, David. 2008. Are big cities really bad places to live? Improving quality-of-life estimates across cities. Working Paper 14472, National Bureau of Economic Research.
- Alonso, William. 1964. *Location and Land Use; Toward a General Theory of Land Rent*. Cambridge, MA: Harvard University Press.
- Atkin, David, Keith Chen, and Anton Popov. 2019. The returns to face-to-face interactions: Knowledge spillovers in Silicon Valley. Processed, MIT.
- Au, Chun-Chung and J. Vernon Henderson. 2006. Are Chinese cities too small? *The Review of Economic Studies* 73(3): 549–576.
- Audretsch, David B. and Maryann P. Feldman. 1996. R&D spillovers and the geography of innovation and production. *American Economic Review* 86(3): 630–640.

- Badger, Emily and Quoc Trung Bui. 2019. Cities start to question an American ideal: A house with a yard on every lot. *The New York Times* 18 June.
- Bairoch, Paul. 1988. *Cities and Economic Development: From the Dawn of History to the Present*. Chicago: University of Chicago Press.
- Baragwanath-Vogel, Kathryn, Ran Goldblatt, Gordon Hanson, and Amit K. Khandelwal. 2020. Mixing satellite imagery to define urban markets: An application to India. *Journal of Urban Economics* (forthcoming).
- Baum-Snow, Nathaniel and Ronni Pavan. 2012. Understanding the city size wage gap. *Review of Economic Studies* 79(1): 88–127.
- de Bellefon, Marie-Pierre, Pierre-Philippe Combes, Gilles Duranton, Laurent Gobillon, and Clément Gorin. 2020. Delineating urban areas using building density. *Journal of Urban Economics* (forthcoming).
- Ben-Akiva, Moshe E. and Steven R. Lerman. 1985. *Discrete Choice Analysis: Theory and Application to Travel Demand*. Cambridge, MA: MIT Press.
- Berliant, Marcus, Robert R. Reed, III, and Ping Wang. 2006. Knowledge exchange, matching, and agglomeration. *Journal of Urban Economics* 60(1): 69–95.
- Bernard, Andrew B., Andreas Moxnes, and Yukiko U. Saito. 2019. Production networks, geography, and firm performance. *Journal of Political Economy* 127(2): 639–688.
- Berry, Christopher R. and Edward L. Glaeser. 2005. The divergence of human capital levels across cities. *Papers in Regional Science* 84(3): 407–444.
- Büchel, Konstantin and Maximilian V. Ehrlich. 2016. Cities and the structure of social interactions: Evidence from mobile phone data. Discussion Paper 1608, University of Bern.
- Büchel, Konstantin, Maximilian V. Ehrlich, Diego Puga, and Elisabet Viladecans-Marsal. 2018. Calling from the outside: The role of networks in residential mobility. Processed, CEMFI.
- Carlino, Gerald A., Satyajit Chatterjee, and Robert M. Hunt. 2007. Urban density and the rate of invention. *Journal of Urban Economics* 61(3): 389–419.
- Carlino, Gerald A. and William R. Kerr. 2015. Agglomeration and innovation. In Gilles Duranton, J. Vernon Henderson, and William Strange (eds.) *Handbook of Regional and Urban Economics*, volume 5A. Amsterdam: Elsevier, 349–404.
- Carozzi, Felipe and Sefi Roth. 2019. Dirty density: Air quality and the density of American cities. Discussion Paper 1635, Centre for Economic Performance.
- Charlot, Sylvie and Gilles Duranton. 2004a. Cities and workplace communication: Some quantitative French evidence. *Urban Studies* 43(8): 1365–1394.
- Charlot, Sylvie and Gilles Duranton. 2004b. Communication externalities in cities. *Journal of Urban Economics* 56(3): 581–613.
- Ciccone, Antonio and Robert E. Hall. 1996. Productivity and the density of economic activity. *American Economic Review* 86(1): 54–70.
- Colomina, Beatriz. 2019. *X-ray Architecture*. Zürich: Lars Müller Publishers.

- Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon. 2008. Spatial wage disparities: Sorting matters! *Journal of Urban Economics* 63(2): 723–742.
- Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon. 2019. The costs of agglomeration: House and land prices in French cities. *Review of Economic Studies* 86(4): 1556–1589.
- Combes, Pierre-Philippe, Gilles Duranton, Laurent Gobillon, Diego Puga, and Sébastien Roux. 2012. The productivity advantages of large cities: Distinguishing agglomeration from firm selection. *Econometrica* 80(6): 2543–2594.
- Combes, Pierre-Philippe, Gilles Duranton, Laurent Gobillon, and Sébastien Roux. 2010. Estimating agglomeration effects with history, geology, and worker fixed-effects. In Edward L. Glaeser (ed.) *Agglomeration Economics*. Chicago, IL: Chicago University Press, 15–65.
- Combes, Pierre-Philippe and Laurent Gobillon. 2015. The empirics of agglomeration economies. In Gilles Duranton, J. Vernon Henderson, and William Strange (eds.) *Handbook of Regional and Urban Economics*, volume 5. Amsterdam: Elsevier, 247–348.
- Couture, Victor. 2016. Valuing the consumption benefits of urban density. Processed, University of California Berkeley.
- Couture, Victor and Jessie Handbury. 2019. Urban revival in America. Processed, University of California Berkeley.
- Dauth, Wolfgang, Sebastian Findeisen, Enrico Moretti, and Jens Südekum. 2018. Matching in cities. Discussion Paper 13347, Centre for Economic Policy Research.
- 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.
- De la Roca, Jorge and Diego Puga. 2017. Learning by working in big cities. *Review of Economic Studies* 84(1): 106–142.
- Desmet, Klaus and Esteban Rossi-Hansberg. 2013. Urban accounting and welfare. *American Economic Review* 103(6): 2296–2327.
- Diamond, Rebecca. 2016. The determinants and welfare implications of us workers' diverging location choices by skill: 1980–2000. *American Economic Review* 106(3): 479–524.
- Donaldson, Dave and Adam Storeygard. 2016. The view from above: Applications of satellite data in economics. *Journal of Economic Perspectives* 30(4): 171–198.
- Duranton, Gilles and Diego Puga. 2004. Micro-foundations of urban agglomeration economies. In J. Vernon Henderson and Jacques-François Thisse (eds.) *Handbook of Regional and Urban Economics*, volume 4. Amsterdam: Elsevier, 2063–2117.
- Duranton, Gilles and Diego Puga. 2019. Urban growth and its aggregate implications. Processed, CEMFI.
- Duranton, Gilles and Matthew A. Turner. 2018. Urban form and driving: Evidence from US cities. *Journal of Urban Economics* 108: 171–191.
- Duranton, Gilles and Anthony J. Venables. 2018. Place-based policies for development. Working Paper 24562, National Bureau of Economic Research.

- Ellison, Glenn, Edward L. Glaeser, and William Kerr. 2010. What causes industry agglomeration? Evidence from coagglomeration patterns. *American Economic Review* 100(3): 1195–1213.
- Fischel, William A. 1987. *The Economics of Zoning Laws: A Property Rights Approach to American Land Use Controls*. Baltimore (MD): Johns Hopkins University Press.
- Fischel, William A. 2001. *The Homevoter Hypothesis*. Cambridge, MA: Harvard University Press.
- Fujita, Masahisa and Jacques-François Thisse. 2013. *Economics of Agglomeration: Cities, Industrial Location, and Globalization*. Cambridge: Cambridge University Press.
- Gaubert, Cecile. 2018. Firm sorting and agglomeration. *American Economic Review* 108(11): 3117–3153.
- Glaeser, Edward L. 2011. *Triumph of the City: How Our Greatest Invention Makes Us Richer, Smarter, Greener, Healthier, and Happier*. London: MacMillan.
- Glaeser, Edward L. and Matthew E. Kahn. 2008. The greenness of cities: Carbon dioxide emissions and urban development. *Journal of Urban Economics* 67(3): 404–418.
- Glaeser, Edward L. and David C. Maré. 2001. Cities and skills. *Journal of Labor Economics* 19(2): 316–342.
- Glaeser, Edward L. and Bruce Sacerdote. 1999. Why is there more crime in cities? *Journal of Political Economy* 107(S6): S225–S258.
- Graham, Daniel J. 2007. Agglomeration, productivity and transport investment. *Journal of Transport Economics and Policy* 41(3): 317–343.
- Greenstone, Michael, Richard Hornbeck, and Enrico Moretti. 2010. Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings. *Journal of Political Economy* 118(3): 536–598.
- Handbury, Jessie and David E. Weinstein. 2015. Goods prices and availability in cities. *Review of Economic Studies* 82(1): 258–296.
- Harari, Mariaflavia. 2020. Cities in bad shape: Urban geometry in India. *American Economic Review* (forthcoming).
- Henderson, J. Vernon. 1974. The sizes and types of cities. *American Economic Review* 64(4): 640–656.
- Henderson, J. Vernon. 2003. Marshall’s scale economies. *Journal of Urban Economics* 53(1): 1–28.
- Henderson, J. Vernon, Sebastian Kriticos, and Jamila Nigmatulina. 2020. Measuring urban economic density. *Journal of Urban Economics* (forthcoming).
- Henderson, J. Vernon and Arindam Mitra. 1996. The new urban landscape: Developers and edge cities. *Regional Science and Urban Economics* 26(6): 613–643.
- Hilber, Christian and Frédéric Robert-Nicoud. 2013. On the origins of land use regulations: Theory and evidence from us metro areas. *Journal of Urban Economics* 75(1): 29–43.
- Hsieh, Chang-Tai and Enrico Moretti. 2019. Housing constraints and spatial misallocation. *American Economic Journal: Macroeconomics* 11(2): 1–39.

- Jaffe, Adam B., Manuel Trajtenberg, and Rebecca Henderson. 1993. Geographic localization of knowledge spillovers as evidenced by patent citations. *Quarterly Journal of Economics* 108(3): 577–598.
- Koster, Hans R.A., Ilias Pasidis, and Jos van Ommeren. 2019. Shopping externalities and retail concentration: Evidence from Dutch shopping streets. *Journal of Urban Economics* 114: 103–194.
- Kreindler, Gabriel E. and Yuhei Miyauchi. 2019. Measuring commuting and economic activity inside cities with cell phone records. Processed, Harvard University.
- Krugman, Paul R. 1991. *Geography and Trade*. Cambridge, MA: MIT Press.
- Leyk, Stefan, Andrea E. Gaughan, Susana B. Adamo, Alex de Sherbinin, Deborah Balk, Sergio Freire, Amy Rose, Forrest R. Stevens, Brian Blankespoor, Charlie Frye, Joshua Comenetz, Alessandro Sorichetta, Kytt MacManus, Linda Pistolesi, Marc Levy, Andrew J. Tatem, and Martino Pesaresi. 2019. The spatial allocation of population: A review of large-scale gridded population data products and their fitness for use. *Earth System Science Data* 11(3): 1385–1409.
- Liu, Crocker H., Stuart S. Rosenthal, and William C. Strange. 2018. The vertical city: Rent gradients, spatial structure, and agglomeration economies. *Journal of Urban Economics* 106: 101–122.
- Manning, Alan and Barbara Petrongolo. 2017. How local are labor markets? Evidence from a spatial job search model. *American Economic Review* 107(10): 2877–2907.
- Manson, Steven, Jonathan Schroeder, David Van Riper, and Steven Ruggles. 2019. *Integrated Public Use Microdata Series, National Historical Geographic Information System: Version 14.0*. Minneapolis: IPUMS.
- Marinescu, Ioana and Roland Rathelot. 2018. Mismatch unemployment and the geography of job search. *American Economic Journal: Macroeconomics* 10(3): 42–70.
- Melo, Patricia C., Daniel J. Graham, and Robert B. Noland. 2009. A meta-analysis of estimates of urban agglomeration economies. *Regional Science and Urban Economics* 39(3): 332–342.
- Mills, Edwin S. 1967. An aggregative model of resource allocation in a metropolitan area. *American Economic Review Papers and Proceedings* 57(2): 197–210.
- Moretti, Enrico. 2004. Estimating the social return to higher education: Evidence from longitudinal and repeated cross-sectional data. *Journal of Econometrics* 121(1): 175–212.
- Moretti, Enrico. 2011. Local labor markets. In Orley Ashenfelter and David Card (eds.) *Handbook of Labor Economics*, volume 4. Amsterdam: Elsevier, 1237–1313.
- Moretti, Enrico. 2013. Real wage inequality. *American Economic Journal: Applied Economics* 5(1): 65–103.
- Moretti, Enrico. 2019. The effect of high-tech clusters on the productivity of top inventors. Processed, University of California Berkeley.
- Muth, Richard F. 1969. *Cities and Housing*. Chicago: University of Chicago Press.
- Ortalo-Magné, François and Andrea Prat. 2011. The political economy of urban growth: Homeownership versus affordability. Processed, University of Wisconsin.
- Overman, Henry G. and Diego Puga. 2010. Labour pooling as a source of agglomeration: An empirical investigation. In Edward L. Glaeser (ed.) *Agglomeration Economics*. Chicago, IL: Chicago University Press, 133–150.

- Puga, Diego. 2010. The magnitude and causes of agglomeration economies. *Journal of Regional Science* 50(1): 203–219.
- Redding, Stephen J. and Esteban Rossi-Hansberg. 2017. Quantitative spatial economics. *Annual Review of Economics* 9(0): 21–58.
- Roback, Jennifer. 1982. Wages, rents, and the quality of life. *Journal of Political Economy* 90(6): 1257–1278.
- Rosenthal, Stuart S. and William Strange. 2004. Evidence on the nature and sources of agglomeration economies. In J. Vernon Henderson and Jacques-François Thisse (eds.) *Handbook of Regional and Urban Economics*, volume 4. Amsterdam: Elsevier, 2119–2171.
- Rosenthal, Stuart S. and William C. Strange. 2001. The determinants of agglomeration. *Journal of Urban Economics* 50(2): 191–229.
- Rosenthal, Stuart S. and William C. Strange. 2008. The attenuation of human capital spillovers. *Journal of Urban Economics* 64(2): 373–389.
- Small, Kenneth A. 2012. Valuation of travel time. *Economics of Transportation* 1(1): 2–14.
- Sveikauskas, Leo. 1975. Productivity of cities. *Quarterly Journal of Economics* 89(3): 393–413.
- US Census Bureau. 2011. *Intercensal Estimates of the Resident Population for Counties and States: April 1, 2000 to July 1, 2010*. Washington, DC: US Census Bureau.
- Tiebout, Charles M. 1956. A pure theory of local expenditures. *Journal of Political Economy* 64(5): 416–424.
- WorldPop. 2018. *Global High Resolution Population Denominators Project*. Southampton: WorldPop (<https://www.worldpop.org>). Funded by The Bill and Melinda Gates Foundation (OPP1134076).
- Xiao, Hongyu and Andy Wu. 2020. Commuting and innovation: Are closer inventors more productive? Processed, University of Pennsylvania.