29/January/2022

Exploring the Urban Model:

Employment, Housing, and Infrastructure.

Daniel M. Sturm,
London School of Economics

Kohei Takeda,
London School of Economics

Anthony J. Venables,
The Productivity Institute, University of Manchester
and Monash University

Abstract

The paper explores the properties of a modern urban model in which households' and firms' locations in the city are endogenously determined as functions of technology, preferences, and geography. This class of model provides insights into the factors that determine the shape and growth of cities. The paper increases understanding by studying the comparative static properties of the model, the effects of various policy interventions, and the circumstances under which different possible city types (mono- versus poly-centric) arise. It is a step to tailoring model structure to adequately describe real world cities of different sizes and types.

Keywords: Urban, agglomeration, transport, quantitative spatial model

JEL classification: O18, R1,R4

Acknowledgements: We gratefully acknowledge the support of the World Bank. The paper is a product of the Urban, Resilience and Land GP. Produced as a background paper for the World Bank's Sustainable Development Practise Group flagship report 'Pancakes to Pyramids; City Form to Promote Sustainable Growth', led by Somik Lall. It is part of a larger effort by the World Bank to provide open access to its research and make a contribution to development policy discussions around the world. Policy Research Working Papers are also posted on the web at http://www.worldbank.org/prwp . The author(s) may be contacted at tony.venables@economics.ox.ac.uk

Authors' addresses:

D. Sturm
London School of Economics
Houghton Street
London, UK
d.m.sturm@lse.ac.uk

K. Takeda
London School of Economics
Houghton Street
London, UK
k.takeda@lse.ac.uk

A. J. Venables,
The Productivity Institute
University of Manchester
Manchester, UK
tony.venables@economics.ox.ac.uk

1. Introduction

This paper explores the properties of a modern urban model in which households' and firms' locations in the urban area are endogenously determined as functions of technology, preferences, and geography. Such models have been applied to data (e.g. Ahlfeldt et al. 2015, Heblich et al. 2020, Redding and Rossi-Hansberg 2017, Monte et al. 2018) and used to simulate policy experiments (e.g. Tsivanidis 2019, Bird and Venables 2019, 2020). In their basic form the models have few ingredients (one sector, one factor) yet they can generate a wide range of possible outcomes. This is a strength of the models, but it is important to understand exactly what drives the results they generate. How do outcomes depend on modeling assumptions? How do parameters affect outcomes, both qualitatively and quantitatively? How does the predicted effect of a policy depend on parameters and modeling assumptions? What are the effects of relaxing some of the simplifying assumptions employed in the model? These questions need to be answered for the sake of transparency, particularly if the models are to be used in policy applications.

This paper is a first step to addressing these questions. It sets out a basic urban model and explores several issues. First, what are the comparative statics of the model with respect to key parameters? Second, what are the effects of various hypothetical policy changes, such as transport improvement and building controls? Third, how do results change under different modeling assumptions? Fourth, what different urban forms – for example mono- or poly-centric – are generated by the model? We address these questions by numerical simulation; clear analytical results are hard to obtain and are far from transparent in this environment of simultaneous co-location of firms and households in a 2-dimensional space.

The model has four main ingredients. The first is geography. Economic interactions take place within a geographical space, which we initially assume to be a 'featureless plain', with commuting and other transport costs a function of distance. This assumption serves to abstract from the idiosyncratic aspects of natural geography that shape real cities, and thereby enables the underlying economic mechanisms to come through more clearly. In the simplest cases it gives rise to circular monocentric cities, although this is not necessarily always the case. Monocentricity is replaced by polycentricity if the city is very large and, in these and other cases, the equilibrium may have rotational symmetry of lower order than the space in which it is modeled. In parts of the paper we employ a network approach and explore the implications of a road structure and optimal route choice, although restricting attention to networks that have radial symmetry.

The second ingredient is households who choose where to live and where to work, consuming floor-space and goods and services, and facing commuting costs. Decisions on where to live and work are discrete choices, to which 'noise' is added by assuming household specific preferences over housing and residential locations. This noise is one of the factors that underlies commuting flows, and also generates some degree of overlap between residential and commercial areas of the city, a key feature of real-world cities. There may also be place specific amenities, although in all cases reported in this paper utility derived from amenities is assumed to be the same in all locations in the city. In our benchmark case city population is assumed to be fixed, and we allow it to vary in response to utility levels in some of the experiments undertaken.

Third, there are perfectly competitive land developers (or builders) who convert land into structures and thence floor-space. The amount of floor-space constructed per unit land is decided by the profit maximizing decisions of these developers and depends, amongst other things, on land rent. Rents at each place are determined within the model and this sets an endogenous city edge, with land outside the city occupied by agriculture with exogenous land-rent. Urban construction takes place where land-rent exceeds this outside rent. Since rent varies within the city area so too does the amount of floor-space that is supplied per unit of land, this most easily thought of as building height. We allow there to be different construction technologies for residential and commercial structures, noting that technology may also vary across the types of structure needed for different commercial uses, i.e. office blocks versus factories. While focusing on building height, the model also contains the potential for infill, i.e. increasing the built area per unit of land at the cost of open or green space. We also note that the model is comparative static, so does not take into account the fact that construction is a sunk cost, and that decisions therefore depend on expected future rents.[1]

The final ingredients are firms who use labor, intermediates, and (commercial) floor-space to produce output. In the benchmark model this is a single composite good, perfectly freely tradable within and outside the city. This good has fixed relative price with respect to 'outside' goods (e.g. food, fuels, other manufactures) so can be transformed into consumption goods and intermediates at fixed terms of trade. Central to the model is the possibility that there are agglomeration economies, with productivity in each place depending positively on spillovers from nearby firms. In later sections of the paper we add a second productive sector which can be thought of as consumer services, including retail. This sector is monopolistically competitive and uses labor, floor-space and the perfectly competitive good (as an intermediate) to produce differentiated products which are sold entirely within the city, and are costly to deliver from firms to households.

We show (numerically and graphically) equilibrium patterns of land-use across the city, looking at residential and employment densities, wages, land-rents, and building heights We investigate the way these vary with parameters of the model and may produce different urban configurations (mono- or polycentric). A number of hypothetical policy experiments are undertaken, illustrating the effects of transport improvements, height restrictions, and zoning controls. An illustration of the calibration of this model on data for Dhaka, Bangladesh, is given in Sturm et al. (2021), which discusses the data needed for such an application.

2. The benchmark model

In this section we set out the benchmark model – with a single production sector – first outlining the economic relationships and determination of equilibrium and then (sub-section 2.1) turning to details of implementation and functional form.

2.1 Model structure

Geography: The city occupies area within a geographical space which is divided into a large number of cells (745 in the simulations and figures that follow). We assume that each cell is a hexagon and the space is the lattice of such hexagons.[2] Cells will be referred to by subscripts i, j, and each of these cells has the same land area, $m_i = m$. The distance between the center points of two cells is d_{ij}, and this will underpin all spatial interactions. Within each cell we set $d_{ii} = 0$. The cells form an (approximate) circle and the city occupies a subset of these cells. The size of this subset (i.e. number of these cells occupied as places of residence, work, or both and hence inside the city) is endogenous, and this set of urban cells is denoted I.

---

[1] See Henderson et al. (2021) for a dynamic urban model in which formal sector construction is durable and incurs sunk costs, while informal sector construction is malleable, and achieves high population density by 'crowding' rather than by building tall.

[2] The hexagonal structure enables a considerably better approximation to a circular city than does a square grid.

Initially we work with Euclidean distance, and we choose units such that the radius of the geographical space is unity (so no city extends more than one-unit distance away from its center). In a later section we work with network distance (i.e. 'distance' formed by ease of movement between the centers of adjacent cells), this network approach allowing us to add transport systems (e.g. roads with different travel speeds) and optimal route choice.

Households: City-wide population is initially assumed to be exogenous and denoted $\bar{L}$. Each household is assumed to have one working member, to choose its consumption pattern, and to make discrete choices of where to live and work. The (indirect) utility function of a household that lives in cell *i* and works in cell *j* is

$$u_{ij} = \mathrm{u}(w_j, \tau_{ij}, q_{hi}; B_i) = w_j \tau_{ij} B_i / q_{hi}^{1-\alpha}, \tag{1}$$

where the final expression gives the Cobb-Douglas form that we take as our benchmark. The wage in cell *j* is denoted w_j and the denominator is the cost of living in cell *i*. The price of residential floor-space in cell *i* is q_{hi}, (where subscript *h* for household captures residential floor-space) and share $1-\alpha$ of income goes to housing. The remainder goes to a single consumption good which is freely tradable and will be used as numeraire. In the numerator, τ_{ij} captures commuting costs and is the fraction of utility remaining after commuting between cells *i* and *j*. This is $\tau_{ij} = 1$ for $i = j$, and $\tau_{ij} = \tau_{ji} < 1$ for $i \neq j$. It depends on distance, as will be discussed below. Cell *i* may also offer a distinct cell specific amenity to residents, the value of which is denoted B_i.

Each household chooses where to live and where to work, and these choices are captured by a discrete choice function giving the probability π_{ij} that an individual will live in *i*, and work in *j*. This choice function takes the form

$$\pi_{ij} = F(u_{ij}) \Big/ \sum\nolimits_i \sum\nolimits_j F(u_{ij}), \quad i, j \in I. \tag{2}$$

where *F* is an increasing function. This is rationalized by idiosyncratic preference shocks, exponential if the choice function is logit, iso-elastic if based on shocks following a Frechet distribution. Notice that, as written, this depends on utility u_{ij} but not on the number of jobs or houses in places *i* or *j*. We discuss this further in section 6, where we add the plausible assumption that location choices depend on the number of jobs in a place, as well as the wage and utility that each offers.[3] If the total city population of workers is $\bar{L}$ then the number living in cell *i* and the number working in cell *j* are respectively,

$$R_i = \bar{L} \textstyle\sum_j \pi_{ij}, \qquad L_j = \bar{L} \textstyle\sum_i \pi_{ij}, \qquad i, j \in I. \tag{3}$$

Production: There is a single perfectly tradable final good, with price unity (the numeraire). Firms are price-taking and perfectly competitive and have unit cost functions depending on wages w_i, the price of commercial floor-space q_{fi}, and productivity A_i. The equality of price (= 1) to unit cost is

$$1 = c(w_i, q_{fi}) / A_i = w_i^{\beta} q_{fi}^{\gamma} / A_i, \tag{4}$$

---

[3] Expressing the same point differently, expression (2) assumes the objects of choice are places in which individuals reside and work. If instead objects of choice are particular houses and jobs then the numbers of houses and jobs in each cell would show up in the choice of place.

where the second equation gives the Cobb-Douglas form with shares of labor and floor-space respectively β, γ. $1-(\beta+\gamma)\geq 0$ is the share of the numeraire good used as an input, in the form of intermediates or capital equipment. Equation (4) can be thought of as implicitly giving the wage that can be offered by a firm in each cell, given productivity and the price of floor-space in that cell, i.e. the wage equation, $w_i = \left(A_i q_{fi}^{-\gamma}\right)^{1/\beta}$. This wage offer determines the attractiveness of a cell as a place of work, and hence determines the equilibrium distribution of production across the city. Thus, the number of employees that this wage offer generates is determined by household decisions (equations (1) - (3)), and it is this that determines the level of employment and production in each cell.

The productivity of firms in cell *i* is A_i, exogenous to the firm but, if there are agglomeration economies, depending on total employment in this and other cells. This agglomeration relationship takes the form

$$A_i = a_i f\left(\sum_j \theta_{ij} L_j\right). \tag{5}$$

There is a cell specific productivity parameter, a_i, which is set at unity in experiments that follow. Productivity spillovers enter through the weakly increasing function $f(.)$; employment in cell j is L_j, and $\sum_j \theta_{ij} L_j$ is a weighted sum of employment in cells *j*, where θ_{ij} captures the spillover interaction between cells. Thus, the productivity spillover from cell *i* to cell *j* depends on parameters θ_{ij} and the function $f(.)$.

Building and developers: Households and firms use floor-space, and the relationship between this and the underlying land area is determined by construction of buildings. We denote floor-space per unit land g_i and the price of floor-space q_i (ignoring, for the moment, the distinction between residential and commercial). Developers' profits per unit land in cell *i* are $r_i = q_i g_i - (g_i)^\nu \kappa$. The first term is revenue and the second is construction costs, with level parameter κ and increasing and convex in floor-space, $\nu > 1$; this convexity captures, for example, increasing marginal cost of building tall. Developers choose floor-space per unit land to maximize profits, giving first order condition

$$g_i = (q_i/\kappa\nu)^{1/(\nu-1)}\,. \tag{6}$$

Maximized profits per unit land are therefore

$$r_i = (\nu-1)\kappa(q_i/\kappa\nu)^{\nu/(\nu-1)}. \tag{7}$$

With free entry of developers this profit is captured entirely by land-owners, so r_i is land-rent. Equation (7) gives the relationship between land-rent and the price of floor-space, and (6) that between the price of floor-space and the quantity supplied. Together, equations (6) and (7) also give the relationship between land-rent and the supply of floor-space, $g_i = (r_i/\kappa(\nu-1))^{1/\nu}$, and hence the elasticity of supply of floor-space with respect to land-rent, $1/\nu$. Notice that, with iso-elastic construction cost function, land-rent and construction costs are respectively fractions $(1-1/\nu)$ and $1/\nu$ of revenue per unit land, $q_i g_i$. If ν is very large then creating a lot of floor-space (i.e. building tall) is prohibitively expensive; relatively little is therefore spent on construction, the share of land-rent in revenue earned on the property and site is high, and the elasticity of floor- space is low.

We allow residential and commercial buildings to have different technologies. The analytical apparatus above applies to both, but with use-specific parameters, $\nu^h, \nu^f, \kappa^h, \kappa^f$ and variables, $q_{hi}, q_{fi}, g_{hi}, g_{fi}$, where *h* denotes household and *f* denotes firm. In competitive equilibrium land-rent in a particular cell is the same for all users but the price of floor-space may be different for residential and commercial users, reflecting these differences in construction technology.

Equilibrium. We suppose that cell *i* is in the city if and only if $r_i \geq \bar{r}$ where $\bar{r}$ is outside (or agricultural) land-rent, fixed exogenously. Land-rents within the city will vary across cells, so this inequality defines the set of cells within the city, $I = \{i \mid r_i \geq \bar{r}\}$. In our simplest cases land-rent follows a monotonic gradient from the center, so the city forms an (approximate) circle.

For a given total population, $\bar{L}$, the model is closed by specifying land market clearing. For each cell *i* in the city this takes the form, in value terms,

$$m_i r_i = \left(1 - 1/\nu^h\right)(1-\alpha) R_i \sum_j \pi_{ij}\, w_j + \left(1 - 1/\nu^f\right)\gamma L_i w_i/\beta, \qquad i \in I. \tag{8}$$

The value of land supplied in cell *i* is $m_i r_i$. Land demand is from households and firms; the R_i households resident in cell *i* have total income $R_i \sum_j \pi_{ij}\, w_j$, of which fraction $(1-\alpha)$ is spent on floor-space, and fraction $\left(1 - 1/\nu^h\right)$ of this is land-rent. Analogously for firms; the value of their output is $1/\beta$ times their wage bill, fraction γ goes on floor-space, and $\left(1 - 1/\nu^f\right)$ of this to land-rent. Notice that in this specification landlords do not demand land, and therefore all their spending goes on the numeraire good.

For much of the paper we work with a closed city (fixed population). In some cases we endogenise population (open city), assuming an iso-elastic labor supply curve to the city, so supply is a weakly increasing function of urban utility. In what follows city population is held constant unless we state otherwise.

2.2 Model implementation

Commuting costs between two cells are a function of time or – given speed – distance between the cells. Our central case takes the exponential form $\tau_{ij} = \exp\left(-C d_{ij}\right) = c^{d_{ij}}$, implying that each unit increase in distance causes proportionate utility reduction of *C*.[4] This is best interpreted by noting that traveling from the center of our geographical space to the edge (a distance of unity) would leave the traveller with fraction $c \equiv \exp(-C)$ of utility. In the benchmark case that we use below we set *c* = 0.7.

Productivity, spillovers, and agglomeration economies: We assume that the functional form of equation (5) for productivity and spillovers is,

$$\theta_{ij} = \exp\left(-H d_{ij}\right) = \eta^{d_{ij}}, \;\text{ and }\; A_i = a_i f\left(\sum\nolimits_j \theta_{ij} L_j\right) = \max\left\{1, \bar{a}\left(\sum\nolimits_j \theta_{ij} L_j\right)^{\chi}\right\}. \tag{9}$$

Thus, spillover weights θ_{ij} decrease exponentially with distance, and $\eta \equiv \exp(-H)$ is the fraction of the spillover that is left at one-unit distance. A low value of parameter η means that effects are spatially concentrated, and in our benchmark case we set $\eta = 0.01$; thus, only 10% of any spillover effect remains at distance 1/10th of the way from center to edge of the space. These weights θ_{ij} combine to form an aggregate measure of cell i's proximity to employment and hence 'spillover-access', $\sum_j \theta_{ij} L_j$. χ is the elasticity of productivity with respect to this spillover-access and we set a benchmark value of $\chi = 0.1$, looking also at cases where this parameter is smaller or zero. We normalise by constant $\bar{a}$, set such that if the entire labor

---

[4] Some papers work with double log specification, $\ln\left(\tau_{ij}\right) = \delta \ln\left(d_{ij}\right)$, $\delta < 0$.

force were spread uniformly over the unit disc (the maximum city area), then the average spillover effect is such that productivity is unity. We also take this as a lower bound on possible productivity, hence the max statement.[5]

Preferences and location choice: We assume that the choice function (2) has function $F(.)$ iso-elastic, $F\left(u_{ij}\right) = u_{ij}^{\varepsilon}$, $\varepsilon > 1$. Micro-foundations of this iso-elastic form can be derived by assuming perfect household mobility, with the household utility subject to residence and place of work specific shocks from a Frechet distribution (see Ahlfeldt et al. 2015). The parameter ε gives the shape of this distribution, and higher values of ε concentrate the distribution, implying that individuals' choices are relatively less shaped by idiosyncratic shocks and more by the economic characteristics of the city contained in u_{ij}. The functional form of the indirect utility function u_{ij} is given in equation (1), and we assume that the amenity of each place, B_i, is equal to unity for all i.

Welfare analysis: The city generates income for its workers and landowners. We measure the real income of households as

$$Real\ wage \equiv \ Eu = \sum_{i} \sum_{j} u_{ij}\, \pi_{ij}. \tag{10}$$

This is the expected utility of a household in the city, with expectation taken prior to choice of location or to drawing any idiosyncratic preference shocks. Utility u_{ij} is, from eqn. (1), the wage deflated by the cost of living in the place of residence and by commuting costs, τ_{ij}. In what follows we refer to this expected value Eu as the real wage in the city.

The city generates rents in excess of the outside rent, and these accrue to landlords. This excess rent is spent entirely on the numeraire, so its value, expressed per person in the city, is

$$Excess\ rent\ per\ capita\ \equiv \textstyle\sum_i m_i (r_i - \bar{r})/\bar{L}. \tag{11}$$

In what follows we will refer the sum of the real wage and excess rent per capita as a measure of welfare. In cases where population is endogenous, we note that (intra-marginal) migrants to the city will receive a utility benefit from moving which is not included in this calculation, and neither is any change in agricultural productivity and incomes outside the city that might follow from population movement.

3. A benchmark city

We compute the equilibrium for different parameter values, taking as benchmark values those discussed above and summarized in table 1. These choices, and some of the relationships that the model generates on the basis of these values, are broadly in line with estimates from the literature. We vary some of these parameters in the following sections, and offer fuller discussion of these values as we do so.

---

[5] In equilibrium some places may have no employment. The max{} statement says that their productivity cannot fall below unity, the 'flat earth' average value, thereby preventing zero employment implying zero productivity.

Table 1: Benchmark parameters:					
Preferences	$\alpha = 0.75$;	$\varepsilon = 9.0$;	Residential construction	$\kappa^r = 1.0$;	$\nu^r = 1.67$;
Firms	$\beta = 0.6$;	$\gamma = 0.2$;	Commercial construction	$\kappa^f = 1.0$;	$\nu^f = 1.25$;
Productivity	$\eta = 0.01$;	$\chi = 0.10$;	Labor supply elasticity	0; closed. 2; 'open'	
Commuting:	utility left after 1 radius $c = 0.7$:		Outside rent	$\bar{r} = 1$	

Levels of employment and residential density are generated by the model and, using the parameters given in table 1, they generate a benchmark city with distribution of employment and residential population illustrated in Figure 1.1. The geographical space is a disk centerd at {0, 0}, covered by a lattice of 745 hexagonal cells. Those taking zero value (dark blue) are outside the city, those taking positive value are inside. Employment in the city (Figure 1.1a) is concentrated in the center because this has the best 'worker-access', an effect magnified by spatially concentrated agglomeration economies. There are residents everywhere in the city (1.1b) although relatively few in the center where they are priced out by firms and high land rents.

The city has radial symmetry, so details of it can be seen more clearly by taking a cut from the center, and values of different variables along such a cut are shown in Figure 1.2, which has distance from the CBD on the horizontal axis. Land-rent varies by a factor of approximately five (near five in the center, unity at the edge, the edge endogenously given at distance 0.7 radii from the center where $r_i - \bar{r}$). This variation of land-rent drives the construction of taller buildings in the center, panel (1.2.b).[6] The marginal cost of building tall is greater for residential than for commercial building ($\nu^r > \nu^f$), this giving different height profiles. Employment and residential densities (i.e. people per unit area) are given in 1.2.c.[7] These are based on a combination of building height (i.e. floor-space constructed) and the proportions of land in commercial vs residential use in each cell. As is apparent, employment is concentrated within 0.4 radii from the CBD. The concentration of employment in the CBD is driven by quite large and spatially concentrated productivity spillovers. Using the parameters of table 1 and the equilibrium distribution of employment, spillovers imply that productivity in the CBD is 27% higher than it is at the city edge (1.2.d), where it is set at base value of unity; the fact that there is a lower bound on productivity means that some employment continues to the city edge.

---

[6] The lines at the city edge are 'fat' because the hexagon lattice does not produce a perfect circle: the lines are averages over the set of edge points.

[7] These are densities, not total numbers of people at each distance. There is much more land area (more cells) in a circumference away from the CBD than in one close to the CBD; the number of people is density times area.

Figure 1.1a: Employment

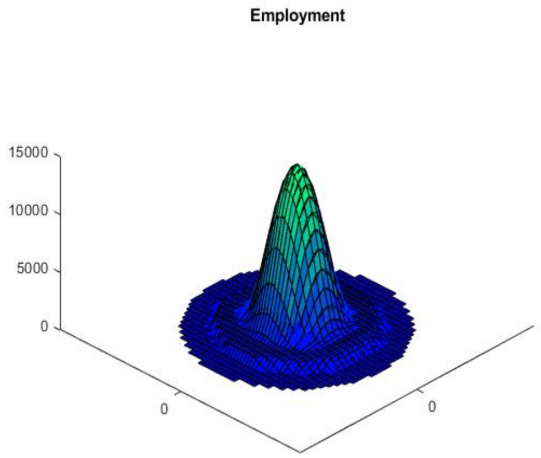

Figure 1.1b: Residential population

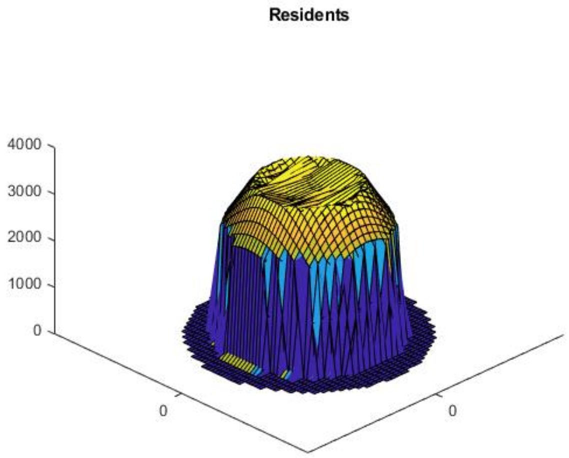

Figure 1.2: Urban density in the benchmark city

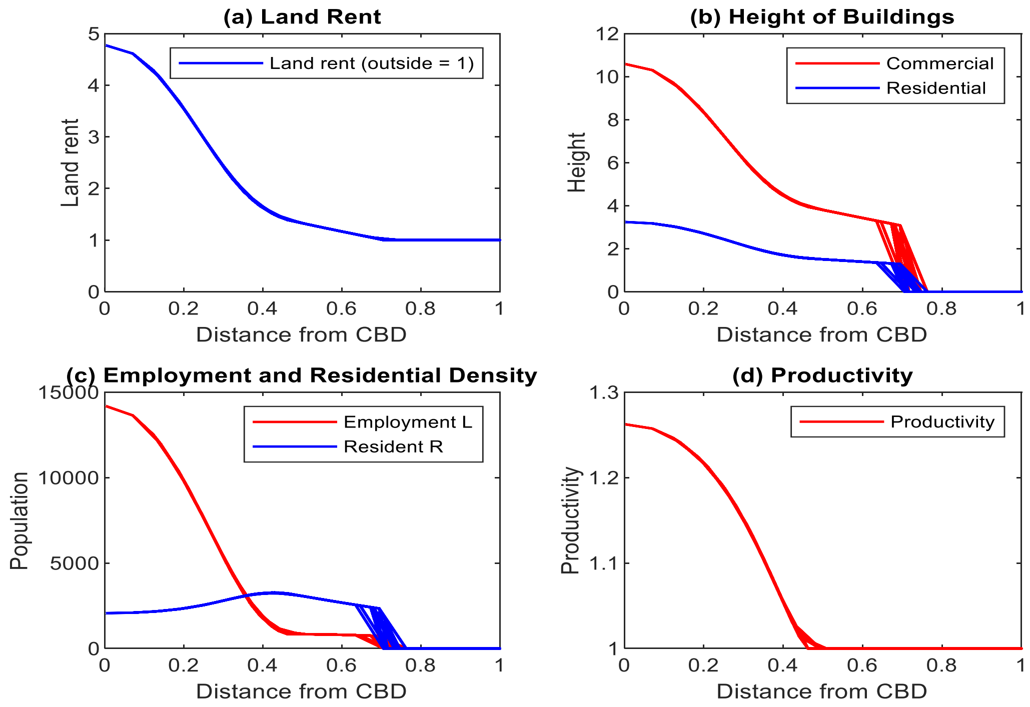

4: The shape of the city

The benchmark city described above is strongly monocentric, with a center that is well-defined, tall, and around which employment is concentrated. Several forces create this relatively concentrated center. One is spatially concentrated agglomeration economies in production, such that firms benefit from proximity to each other. A second is ‘worker-access’; given commuting costs, firms benefit from locating in a place that is relatively easily reached by commuters. The third is ‘market access’; if there are costs of shipping goods, firms benefit from locating where they can be easily reached by consumers.

In the remainder of this section we look at the first two of these forces (sections 4.1 and 4.3). The last of these effects is switched off in the benchmark model as output is assumed to bear no transport costs, but we return to it in an extension of the model in section 7. The strength of each of these forces depends also on the population of the city, as we show in section 4.2.

4.1 Agglomeration economies

The city without agglomeration economies (benchmark city, but with $\chi = 0.0$) is illustrated in figure 2.1. The city is now a flat dome, rather than the previous spike (the vertical axes of figure 2.1 are at larger scale than in figure 1.1). Even without agglomeration economies, the CBD has highest rent, as it has the best worker-access; however, land rent in the center is 60% higher than at the edge (compared to a factor of five times in the benchmark) and buildings are correspondingly less tall. Absent the proximity benefits of agglomeration firms locate close to workers, as illustrated in 2.1.c, although there is still a substantial amount of commuting because of workers' idiosyncratic preferences for residence/ employment pairs. The total area occupied by the city (as indicated by the city edge at around 0.7) is very similar to that in the benchmark case. This is perhaps surprising, as the city appears less dense. However, while employment occupies more space (lower land rents mean less tall buildings), the overall effect is driven principally by residential land use close to the center. Rents at each place are lower so residential occupation at each place is less dense; but this is offset by a much higher proportion of residential accommodation being located near the center, where rents are relatively high.

Further insight on the role of agglomeration economies is given in figure 2.2, in which the horizontal axis in each panel is the measure of agglomeration economies, χ. All variables on the vertical axis are expressed relative to the benchmark equilibrium, where $\chi = 0.1$. Thus, the point $\chi = 0.1$ is the benchmark city described in Figure 1, while the dome of Figure 2.1 is where $\chi = 0.0$.

Aspects of the city shape are illustrated in Figure 2.2a. The maximum (i.e. CBD) value of land rent increases sharply with the strength of agglomeration. The convexity of the relationship is due to the externalities associated with agglomeration economies and firm location. Larger χ means not only that firms already near the CBD get higher productivity, but also that firms move closer to the CBD, so the effect is magnified.[8] Average population density varies little because, as noted above, the city area changes little. Commuting distances (average and hence also total) increase with agglomeration, although only by 9% across the range of χ values given in the figure. The effect only becomes apparent at relatively high values of χ, in line with our discussion of the convexity of the maximum rent schedule.

Panel (b) gives real income measures. Peak productivity is 26% higher with agglomeration economies of $\chi = 0.1$ than at zero, and once again this is both a direct effect and amplification due to movement of firms towards the CBD. The benefits of this productivity accrue proportionately more to land-rents than to real wages, the former increasing by 20% and the latter 73% across the range $\chi = [0, 0.1]$. Notice that the relatively dampened wage effect is not due to labor entering the city from outside – the city is 'closed' in this experiment. It is simply that, within the city, competition for the immobile factor (land near the CBD) raises its price and hence the share of city income going to land-rent.

Productivity spillovers in the model are captured by two parameters, χ and η, the latter capturing the spatial range of spillovers. Increasing this range (increasing the fraction of spillover effect remaining at each distance)

---

[8] At even higher values of χ employment concentrates further – providing the same iso-elastic construction technology continues to enable building ever taller.

has a direct positive effect on productivity, but reduces the spatial concentration of the city, reducing rents, height, and employment near the CBD, this partially offsetting the direct positive productivity effect.

Figure 2.1: Urban density with no agglomeration economies

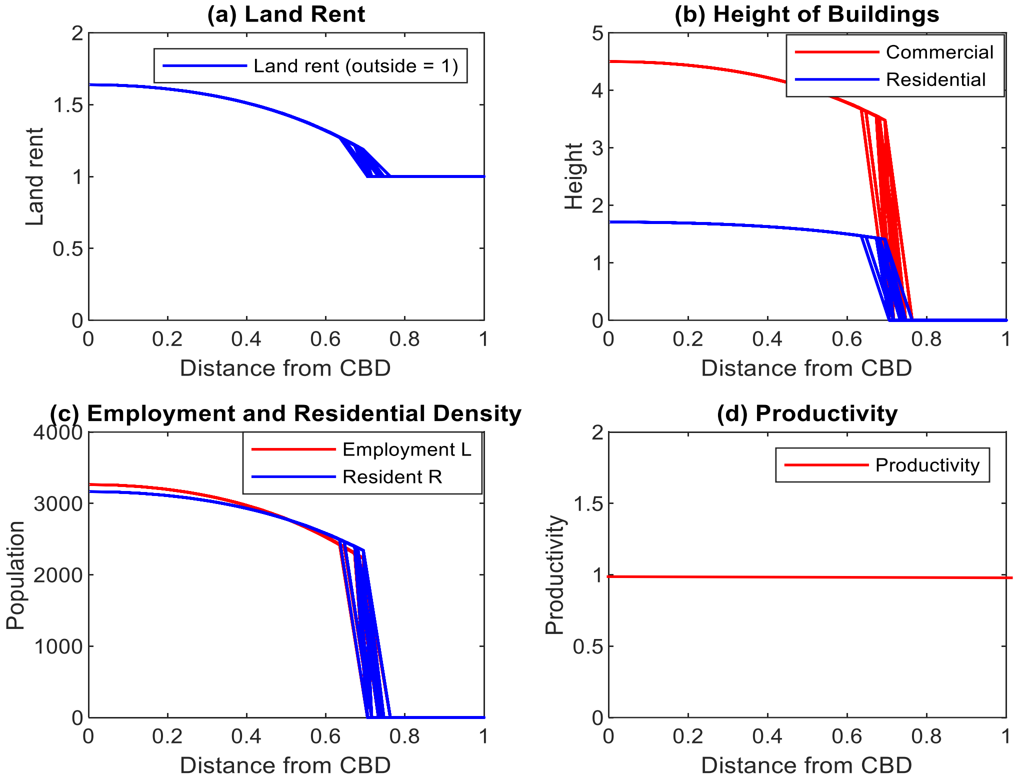

Figure 2.2: The strength of agglomeration economies

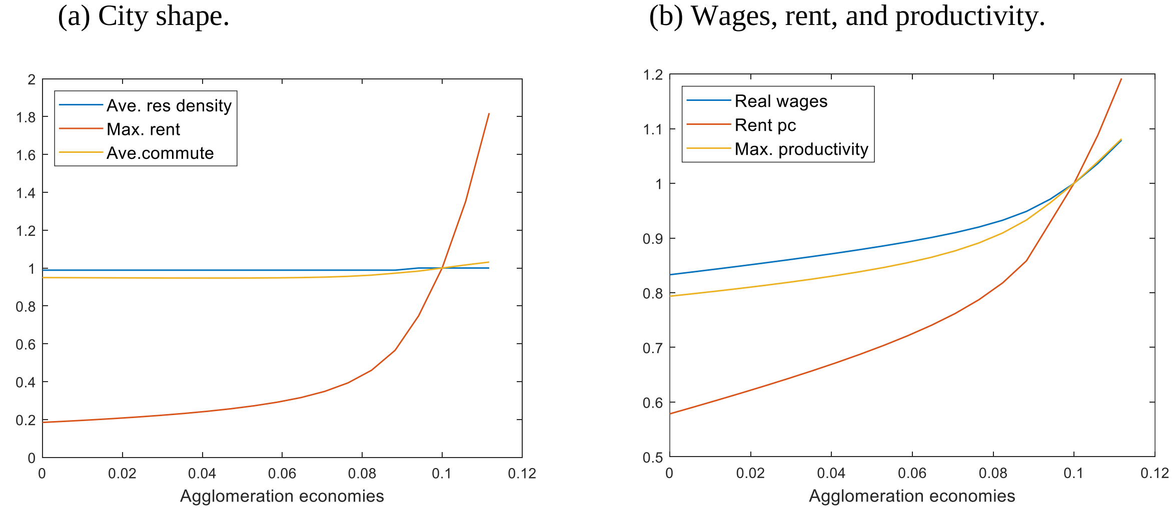
(a) City shape.

(b) Wages, rent, and productivity.

4.2 Population

Large cities – other things equal – have higher city center rents and higher CBD productivity than do small ones. We draw this out and provide illustrative quantification of effects by varying the population of our benchmark city (with agglomeration economies returned to benchmark value $\chi = 0.1$). Results are indicated in the panels of Figure 3. The horizontal axis is population relative to that in the benchmark, so covers a near four-fold variation in city size, from 50% up to 180% of the benchmark.

Panel (a) indicates the large variation in city center land-rents, reflecting high employment density and causing increasing building height. The area of the city grows more or less in proportion to population, so average residential density is approximately constant. The combination of larger city area and central concentration of employment means that the average distance commuted goes up by more than 20%.

Panel (b) gives productivity and income effects. With population at 1.8 times the benchmark, productivity at the CBD is around 3% higher than it is with population 50% that of the benchmark. Nominal wages move broadly in line with productivity, but real wages face the penalty of higher average commuting costs and rents in the larger city. In the case illustrated they fall by around 4%, while average rent per capita increases by around 6%.[9] Aggregate welfare per capita (real wages plus land-rent) is lower in the larger city, since rents amount to less than 1/5th of city income. The effect of increasing population is to reduce the real income of workers as the benefits of agglomeration are transferred to landlords.

Figure 3: City population

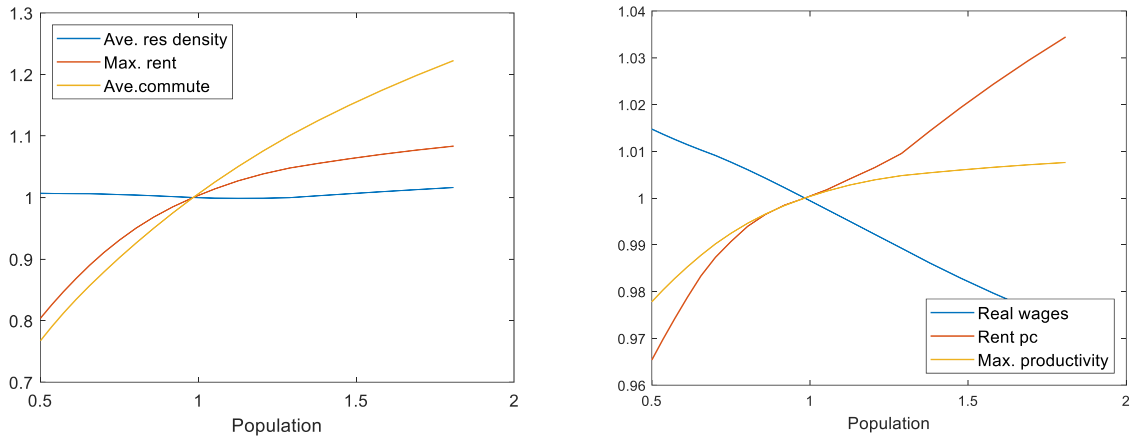

(a) City shape:

(b) Wages, rent, and productivity

Two important points follow from these numbers. First, the tension between urban increasing returns to scale (agglomeration) and diminishing returns (scarce land) give a combined negative effect of scale. Even with the high agglomeration parameter we assume in the benchmark case ($\chi = 0.1$), the effect of increasing city size is to reduce real income.

Second, the relatively small productivity effect arises because increasing city population by a factor of four does not raise employment density, even in the center, by this amount; the central employment area grows outwards, as well as upwards. Econometric estimation of agglomeration economies should, ideally, control

---

[9] The city is closed, so this near constancy of real wages is not due to elastic supply of migrants. It is because the competing forces of agglomeration benefits and commuting plus rent costs come close to cancelling out.

for this. A naïve regression of productivity on city size would fail to control for any of this spreading effect, and so would not capture the elasticity χ, underestimating it in this case by a factor of four.

4.3 Commuting costs

If commuting costs are the only driver of spatial concentration – the benchmark model with agglomeration economies switched off – then the relationship between these costs and the centrality of the city is inverse U-shaped, implying that central tendency and height of the dome are greatest at intermediate values of commuting costs. Essentially, there is no geography at all if commuting costs are zero, or if they are prohibitive so each place has no interaction with any other.

In the presence of agglomeration economies low commuting costs have the dual effects of enabling both employment concentration and residential dispersion. This is illustrated in Figures 4.1 and 4.2, on which the horizontal axis is travel cost per unit distance, relative to that in the benchmark.[10] Thus, the left-hand end of the horizontal axis is 20% lower commuting costs between all pairs of cells in the city.

Figure 4.1: Commuting costs with a fixed population

(a) City shape: (b) Wages, rent, and productivity

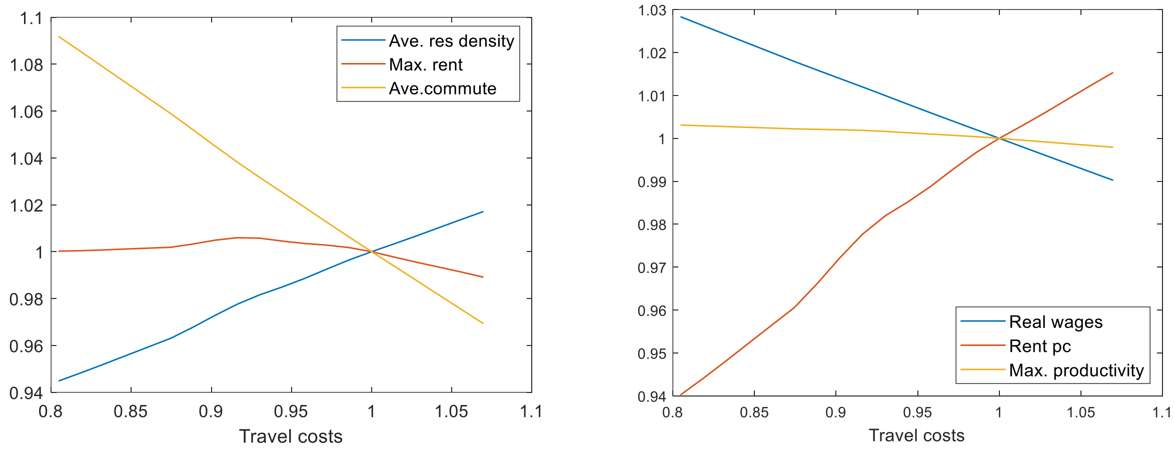

Figure 4.1 is for a closed city, i.e. with population held constant at the benchmark level. Panel (a) indicates that the 20% reduction in travel costs leads to a 10% increase in average distance commuted. Average residential density falls, reflecting an increase in city area of approximately 6%. CBD rents and hence CBD employment density vary little as people live on average further out but commute longer distances.

Real income effects are shown on panel (b). A reduction in commuting costs brings a small fall in maximum productivity as the CBD spreads slightly. There is a significant fall in rents per capita, of around 6% as the area of the city has increased and residential density fallen. The net effect of cheaper commuting, lower rents and slightly lower productivity, raises real wages by around 3%.

---

[10] The experiment varies the parameter C in the travel cost factor, $\tau_{ij} = \exp(-Cd_{ij})$, so reduces effective distance.

The relationships in Figure 4.1 give the effects of better transport, allowing adjustments within the city, but holding city population constant. However, since one of the effects is an increase in real wages, a full equilibrium response is likely to include in-migration and an increase in the urban population. To capture this in the simplest possible way we assume a supply curve of labor to the city which is increasing in urban real wages relative to outside wages (which are assumed constant). Figure 4.2 gives effects when the elasticity of this supply curve is set at 2, in which case a 20% reduction in travel costs leads, via an increase in the real wage, to a population increase of slightly less than 6%. This somewhat reduces the real wage increase (as expected by the combination of figures 4.1(b) and 3.1(b)) and emphasises the ambiguous impact of a transport improvement on the CBD. Transport improvement now tends to reduce CBD rents as it allows dispersion, but raise them as it attracts population, with the ambiguous effects on illustrated on Figure 4.2.

Figure 4.2: Commuting costs in an open city (endogenous population)

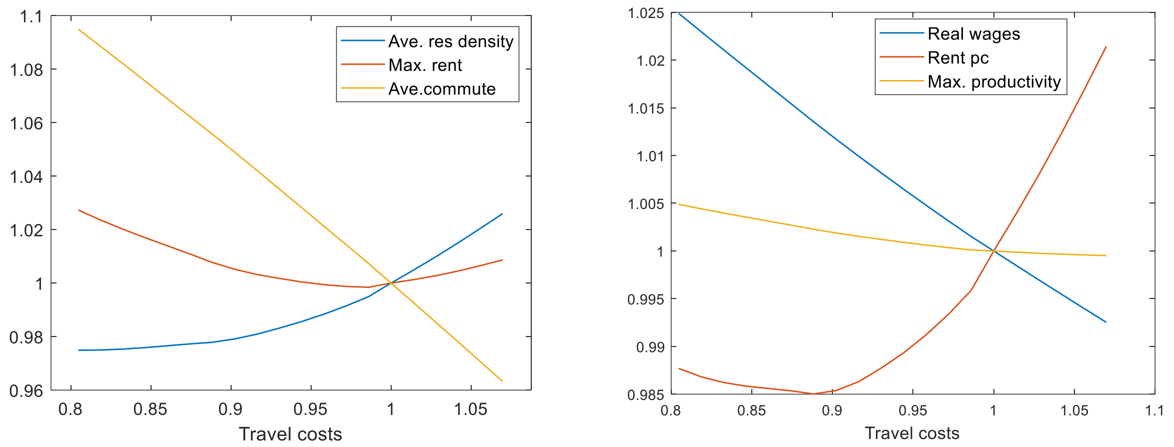

(a) City shape:

(b) Wages, rent, and productivity

Each of these transport quality experiments is assumed to affect transport quality everywhere in the city. Particular place specific transport projects change the shape of the city in more complex ways, and the more localised the project, the greater is the share of the benefit that accrues to landlords. This is because residential choice means that the demand for housing in a particular place is quite elastic, so the immobile factor in the neighbourhood of the transport improvement captures most of the benefit. We develop this further in section 5 of the paper, looking at the effect of a radial road network.

Finally, the benchmark model assumes that commuting costs take exponential form, so each unit increase in distance brings an equi-proportional reduction in utility. Replacing this with a linear form – each unit increase in distance brings an equal absolute reduction in utility, set with the same proportion of utility left after travelling one-unit distance – gives an outcome very close to the benchmark case. Employment is slightly more concentrated in the CBD (maximum value 5% higher) with the average distance commuted increasing by just 0.75%.

4.4 Construction costs and the supply of floor-space

The concentration of employment towards the center of the city is a key part of the effects described above. This concentration is enabled principally by the construction of more floor-space per unit land, and also by reduced consumption of floor-space per household and worker. This construction could take the form of 'infill', or be achieved by height, as we have described in the present model.

Figure 5 captures the effect of variation in the cost of building tall (the parameters ν^h, ν^f). Variation in these parameters may be due to variations in building technology, but we think of these parameters as being a simple way of capturing a set of other factors that influence building decisions. One is the cost of capital; construction is a durable investment, so countries with a high cost of capital will find capital intensive investments relatively more expensive (recall that construction costs are fraction $1/\nu$ of combined construction and land rent costs). A further factor is security of tenure, clarity of land rights, and costs of formalization and making these rights secure; evidently, the attractiveness of sinking capital in formal sector buildings is reduced by the presence of these obstacles.

Sensitivity of outcomes with respect to these factors is illustrated in Figures 5, where the horizontal axis is the variation in ν^h, ν^f, relative to their benchmark values (and recall that the elasticity of supply of floor-space of each type with respect to land rent is $1/\nu^h, 1/\nu^f$). Evidently, higher costs of building tall (and lower elasticity of supply) have a large effect on inhibiting central densification, reducing CBD land-rents, and employment and residential density, and increasing commuting distances. Productivity and real wages are reduced, as are rents.

Figure 5: Construction costs

(a) City shape:

(b) Wages, rent, and productivity

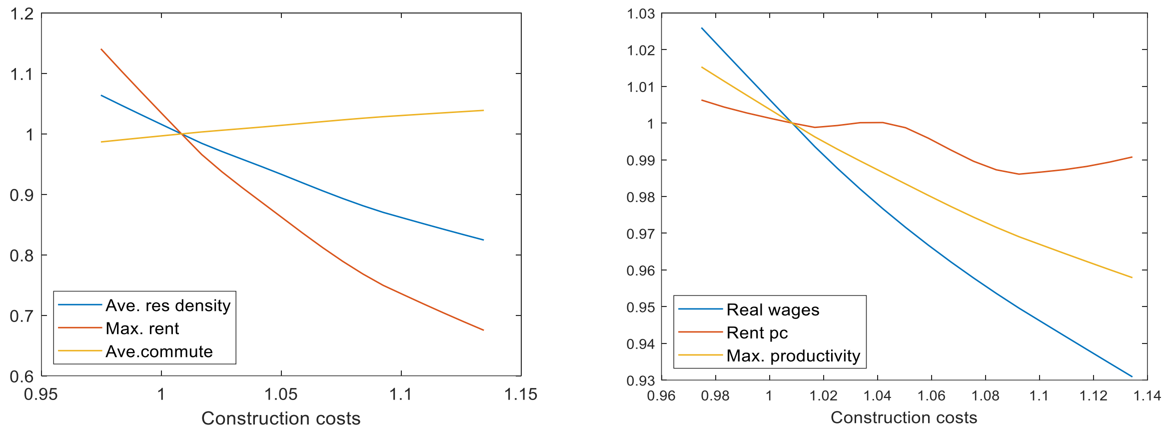

5. Policy experiments

In this section we move from broad city-wide comparative statics to more targeted policy interventions, looking at restrictions on building heights (or, equivalently, maximum floor-area ratios), zoning, and construction of a road network.

Height restrictions: Density is achieved by building tall. What is the effect of regulations that restrict height? To explore this we suppose that the regulation applies city-wide, but binds only for commercial building. We set the ceiling at just less than half the height of the tallest buildings in our base case (a value of 5 in table 2 compared to maximum height of 10.5 in figure 1.2.b).

As expected, employment is spread through a wider area of the city, and land rent at the center of the city is just two-thirds of that in the base case. Productivity at the CBD is somewhat lower, although it is slightly higher further out. The net effect is to reduce aggregate welfare by 3.9%, composed of an increase in total land-rent of 10% and a reduction in household real wages of 5.3% (table 2). It may seem surprising that a restriction on the use of land should increase total urban land rent, but there are two forces at work. In places where the constraint binds (i.e. near the CBD) rent is lower, but the height restriction's overall effect is equivalent to a reduction in the amount of land available to the city. Given total population, total land-rent rises, and the negative impact on household real income is correspondingly greater.

Table 2: Height restriction: real wages and excess rent relative to benchmark

	Real wages	Excess rent pc	Real wages + excess rent pc
Maximum height: 5	0.947	1.101	0.961
Maximum height: 4	0.921	1.208	0.948

Zoning: We explore four different zoning experiments in the first three of which land-use is restricted in a 60^o wedge of the city (i.e. a wedge with vertex at the benchmark CBD). Cases are tabulated in table 3, and in the first case land in this wedge is restricted for residential use only. Welfare increases by 2.4%, with rents 12.8% higher and real wages 1.3% higher.[11] The second case restricts the area to commercial use only; welfare is 2.7% higher than in the benchmark, with rents 24.6% higher and real wages 0.5% higher. In the final case the land is unavailable for either use. This raises welfare by 2.74%, rents going up by 24.9% and real wages increasing 0.5%. In each of these cases it is unsurprising that total rent goes up, as land supply is restricted. The overall welfare gain is surprising, and arises from the productivity externality. Essentially, land restriction forces firms closer together, this increasing productivity. (Unlike the case of height restrictions, where workers were forced further apart).

There is an artificiality about this experiment, in so far as the city remains very monocentric, simply with the CBD shifted somewhat away from the zoned 'wedge' of the city. An alternative experiment, that shown in the final row of table 3 is to take out a disk centered at the benchmark CBD, this impacting the maximum productivity that the cluster can attain. In this case (with zoned disk taking 5.6% of city area) there is a real income loss, substantial at 3.5% of welfare. Real wages are reduced by 3.7% and excess rents somewhat less, by 1.4%. Comparison of these cases indicates just how context specific effects can be, depending in particular on how activity reorganises in response to zoning.

---

[11] This calculation does not include any direct or amenity benefit from the zoning.

Table 3: Zoning: real wages and excess rent relative to benchmark

	Real wages	Excess rent pc	Real wages + excess rent pc
Residential only (60° wedge)	1.014	1.128	1.024
Commercial only (60° wedge)	1.005	1.246	1.027
Residential + commercial (60° wedge)	1.005	1.249	1.027
Residential + commercial (disk around CBD = 5.6% city area)	0.963	0.986	0.965

Road systems: What is the effect of placing a road network on the benchmark city? To explore this we add a radial network, with six roads leading out from the CBD on which it is assumed that journeys are 50% faster (as if distance was halved) than elsewhere. The commuting cost network between cells *i* and *j* is then based on the shortest path between these cells, following nodes at the center of each hexagon.[12] Productivity spillovers remain based on unchanged Euclidean distance, and none of the other aspects of the benchmark are changed. The road network removes the (near) perfect circularity of the city and replaces it with a six-fold symmetry. This is illustrated most clearly by the pattern of residence, shown on figure 6 (a vertical view of a figure similar to figure 1.1b). Cells beyond the city edge are dark blue, and the blue-green area in the center is where there is little residential population as land is occupied by production. The cells with the densest residential population are those in bright yellow, on the road system close to the center

Figure 6: Residential population with a radial road network

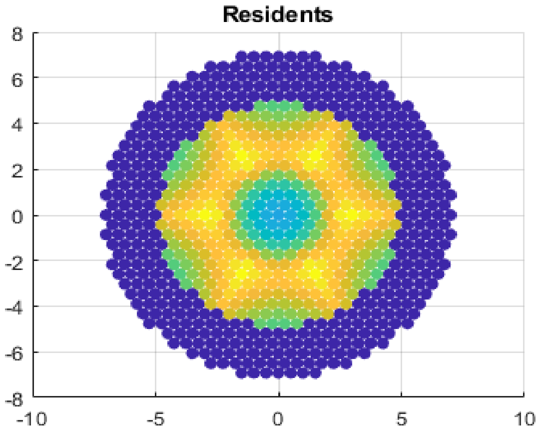

The presence of the network has a powerful centripetal force, nearly doubling employment density and land rent at the city center. The mechanism is essentially that the road system enables residents to move away from the CBD, this releasing land for employment near the center. The further concentration of employment brings a productivity gain of around 5% at the CBD.

The real income effects of the road network are substantial, and we break them down into the contributions of different mechanisms (table 4). If total city population is fixed, overall welfare is 6.5% higher, composed of an 8.5% increase in rents and 6.3% increase in real wages (row 3). This is comprised of the direct effect of reducing travel times on the network, holding constant places of residence, work, and the level of productivity, giving a 2.7% increase in welfare (row 1). Allowing adjustment in the location of jobs and residence but holding productivity constant raises the welfare gain slightly to 2.8%. The remaining welfare

---

[12] This uses the Dijkstra algorithm to compute the shortest path between all node pairs, implemented through the matlab command 'distances'.

gain, of 3.8% of initial welfare, is due to the productivity increase associated with agglomeration and greater concentration of production in the center of the city.

The numbers in the first three rows of table 4 are computed with city population constant. Making the city 'open' (as in section 4.2, with an elasticity of population with respect to real wages of 2), results in an increase in population of 12.4%, this further raising the per capita welfare of city residents and giving a total welfare increase of 8.7% (row 4).

The key message from this example is the differential between the direct effect of the transport improvement (essentially equal to 'user benefit' of faster travel times) is less than one-third of the total benefit once the relocation of firms and workers and consequent productivity gains have been achieved.

Table 4: Road network: real wages and excess rent relative to benchmark

	Real wages	Excess rent pc	Real wages + excess rent pc
1: Direct effect	1.03	1.00	1.027
2: Relocation, given productivity	1.03	1.016	1.028
3: Relocation + productivity change	1.063	1.085	1.065
4: Relocation + productivity change + population growth	1.068	1.263	1.087

6. Modeling the choice of residence and workplace.

An important element of the model is the form and parameterisation of the discrete choice function that guides households' choices of where to live and work. In the model as developed so far this depends on utility, and hence the economic variables of wages, commuting costs, and the price of floor-space. The responsive of choices to spatial variation in utility is captured by the parameter ε (where, in equation 12, $F(u_{ij}) = u_{ij}^{\varepsilon}$), a low value of ε meaning that households are not very sensitive to these variables.[13] If $\varepsilon = 0$ then all resident/employment location pairs are equi-probable, and jobs and employment are uniformly distributed within the city (but with massive amounts of commuting, as a worker living in one place has equal probability of working in any other). As ε becomes very large so choice equalises values of utility u_{ij} across all occupied cells, this giving a sharp partition of between work and residential areas, as in the standard Alonso-Muth-Mills model (see Duranton and Puga 2015).

The sensitivity of results with respect to this parameter around its benchmark value, $\varepsilon = 9$, are illustrated in figure 7.1, and shown to be large. A low value of ε means that there is less concentration of activity near the CBD, a larger city area per unit population, and longer commutes. As a consequence of his dispersion productivity, real wages, and rents are lower (the former two moving together).

---

[13] Micro-foundations can be given to this relationship if $\varepsilon > 1$ by assuming that idiosyncratic preferences over places are drawn from a Frechet distribution with shape parameter ε, see Alhfeldt et al. (2015).

Figure 7.1: Sensitivity of location choice:

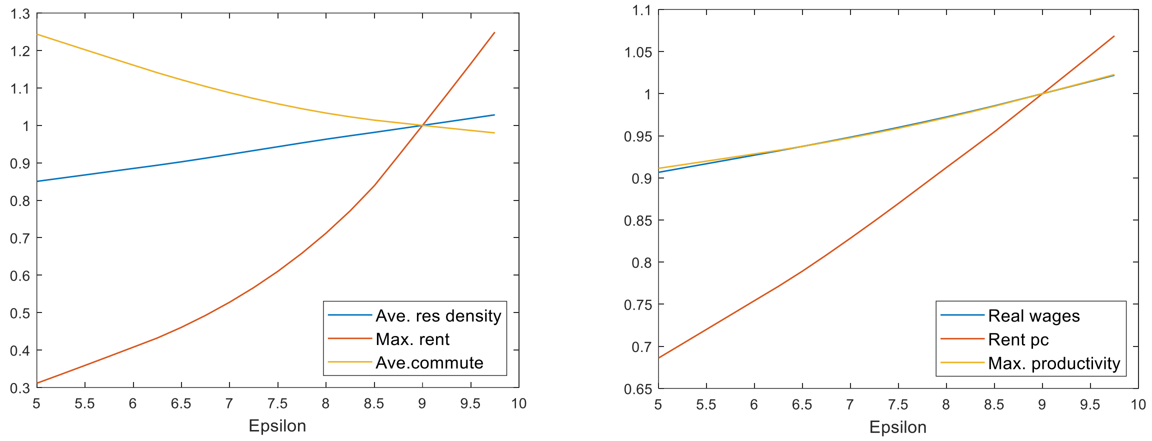

(a) City shape:

(b) Wages, rent, and productivity

A further issue arises as the standard model's specification of the discrete choice function captures returns to employment only through the wage on offer in each cell, w_j. The number of jobs in each cell is immaterial – indeed, the number might be zero, and the wage a notional one that would be offered if a firm were located there. This is conventional in so far as consumer choices are generally modeled as depending on prices, rather than quantities available, but it seems quite unsatisfactory in the present context. It seems likely that choice is influenced by the number of nearby jobs, as well as by the wage they offer.

This can be modeled if households choose not a place of work and place of residence, but a job (which is place specific) and a place of residence. In this case the discrete choice function of equation (2) becomes

$$\pi_{ij} = L_j F(u_{ij}) \Big/ \sum_i \sum_j L_j F(u_{ij}), \quad i, j \in I. \tag{12}$$

As before, π_{ij} gives the probability of living in cell *i* and working in cell *j*, but this is now the sum over jobs in *j* of probabilities of choosing particular job and residence combinations. Since all jobs in a particular cell are identical this is simply multiplication by the number of jobs, L_j, this number also entering the denominator as household choices are now over this expanded set.[14]

What difference does this make? Households now look at the number of jobs in each place and are, other things equal, more likely to locate close to a place with many jobs. This is a powerful additional force for agglomeration, strengthening the central tendency in the city. Using the model to explore the consequences of specification (12) (with benchmark $\varepsilon = 9$) we find that, if there are no productivity spillovers at all, then CBD employment density becomes 50% greater than CBD resident density, as compared to near equality of the two in the standard case (figure 2.1.c). A similar degree of centrality to that found in the benchmark case arises if specification (12) is combined with agglomeration elasticity reduced from $\chi = 0.1$ to $\chi = 0.02$. This is illustrated in figure 7.2. The qualitative shape of the city and extent of concentration in the CBD is similar to

---

[14] This argument could be pushed further, if household choice were over pairs of job *and particular house*, in which case $\pi_{ij} = R_i L_j F(u_{ij}) / \sum_i \sum_j R_i L_j F(u_{ij})$. Simulation with $\chi = 0.02$ gives the city center entirely residential, beyond which there is a ring of employment, and then a further ring of residency.

that in figure 1.1, but the city as a whole occupies slightly less land (there is higher residential density as workers seek to locate close to places with many jobs) and, since the agglomeration elasticity is cut by a factor of five, much lower productivity. The sensitivity of the model with respect to this aspect of model specification is an important area of future research.

Figure 7.2: Urban density when location of jobs matters; ($\chi = 0.02$).

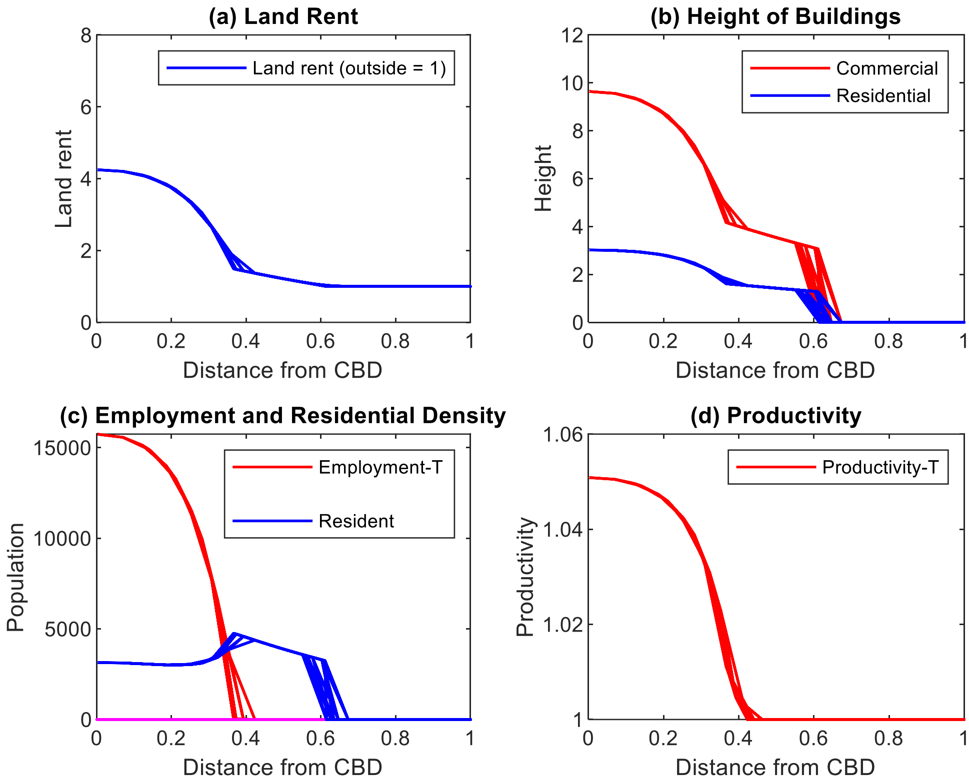

7. Multiple sectors: Tradable goods and consumer services

The benchmark model has a single productive sector the output of which is tradable – within the city and with the rest of the world – at zero cost. Adding multiple sectors increases the realism of the model, while adding complication and reducing transparency. It is worth moving some way along this trade-off by adding a 'consumer services' sector which, depending on context, includes retail, hospitality, maintenance and other residential and household services, as well as the informal sector. A majority of employment in most cities is in such sectors. We assume that the output of this sector is produced using labor, land, and, as an intermediate good, the output of the other (perfectly tradable) sector; recall that this can be traded into goods supplied from outside the city, some of which – such as food – are an input to the retail sector.

Deliveries around the city of the output of this consumer services sector are costly, and none are sold outside the city. We represent the sector by the standard Dixit-Stiglitz model of product differentiation and monopolistic competition in which firms each offer a distinct variety of good or service. The number of such firms is endogenous, and they now also make discrete choices of where in the city to operate, this depending on costs of land and labor, and also on ease of access to the city's residential population (market access).

Making this extension to the model, the indirect utility function becomes

$$u_{ij} = \{w_j t_{ij} B_i / (G_i^{\mu} q_{hi}^{1-\alpha-\mu})\} \tag{13}$$

where G_i is the price index of consumer services consumed in cell *i*, and μ is the share of expenditure devoted to these goods. The price index is a CES function of the prices and numbers of varieties supplied, taking the form

$$G_i = \left[\sum_j n_j \left(p_j T_{ji}\right)^{1-\sigma}\right]^{1/(1-\sigma)}. \tag{14}$$

The number of varieties produced in cell *j* is n_j, their producer price is p_j and their price delivered to cell *i* is $p_j T_{ji}$, where T_{ji} is an iceberg transport cost factor, with $T_{ii} = 1$. This demand system implies that the total sales of a single variety produced in cell *i* is

$$x_i = \sum_j p_i^{-\sigma} T_{ij}^{1-\sigma} E_j G_j^{\sigma-1} \tag{15}$$

where E_j is cell *j* expenditure on consumer services (share μ of wage income in that cell).

In each cell *i* firms maximise profits, $\Omega_i = p_i x_i - [x_i + F]c_i$, where c_i is unit cost and Fc_i is fixed cost, this giving price

$$p_i = \sigma c_i / (\sigma - 1). \tag{16}$$

Unit cost is a Cobb-Douglas function of inputs of labor, land, and the tradable (numeraire) good, so

$$c_i = w_{2i}^{\beta_2} q_{2i}^{\gamma_2} / a_{2i}, \tag{17}$$

where additional subscript 2 distinguishes this sector from the perfectly competitive one, $1 - \beta_2 - \gamma_2$ is the share of intermediates (at price unity), and a_{2i} is a sector and cell specific productivity, which we set equal to unity for all *i*.

Firms make zero profits if they sell output level $\bar{x} = (\sigma - 1)F$. We set this to unity (without loss of generality). Free entry and exit of firms is then the condition $\sum_j p_i^{-\sigma} T_{ij}^{1-\sigma} E_j G_j^{\sigma-1} = 1$ (from 15) and this, together with (16) and (17) gives equation

$$w_{2i}^{\beta_2} = q_{2i}^{-\gamma_2} a_{2i} \left(\frac{\sigma - 1}{\sigma}\right) \left(\sum_j T_{ij}^{1-\sigma} E_j G_j^{\sigma-1}\right)^{1/\sigma}. \tag{18}$$

This is the 'wage equation' giving the wage at which firms in cell *i* make zero profits. An analogous equation for the perfectly competitive sector comes from the unit cost function, equation (4), from which firms in this sector offer wage $w_{1i}^{\beta_1} = A_i q_{fi}^{-\gamma_1}$ (using sub-script 1 to distinguish this sector, where necessary). If both sectors are operating in cell *i* we assume that employment in that cell is divided between sectors according to

$$L_{2i}/L_i = H(w_{2i}) / \left(H(w_{1i}) + H(w_{2i})\right) \tag{19}$$

Where $H(.)$ is an increasing function, set to be iso-elastic with elasticity ε_w.[15] The level of employment determines the number of firms in the sector since (from labor demand), $n_i = L_{2i}w_{2i}/(\gamma_2 p_i x_i)$.

Implementation of the two-sector model: Parameters are given in table 5 below. We assume that there are no productivity spillovers in sector 2, and that commercial building in this sector uses the residential construction technology.

Table 5: Benchmark parameters:				**Two-sector parameters**	
Preferences	$\alpha = 0.75$;	$\varepsilon = 9$;	Preferences	$\alpha = 0.25$; $\mu = 0.5$;	$\varepsilon = 9$; $\varepsilon_w = 40.0$;
Firms	$\beta = 0.6$;	$\gamma = 0.2$;	Firms	$\beta_1 = 0.6$; $\beta_2 = 0.6$;	$\gamma_1 = 0.2$; $\gamma_2 = 0.2$;
Commuting:	utility left after 1 radius c = 0.7:		Shipping sector 2 goods/ services:	Iceberg left after 1 radius T = 0.7:	

Adding this sector to the benchmark model – with agglomeration, $\chi = 0.10$ – gives the city configuration illustrated in figure 8a. The main feature is the addition of 'employment-N', i.e. employment in this non-tradable consumer services sector. As shown in panel (c) the tradable (T-) sector – as modeled in the benchmark case – remains closest to the center, while the N-sector locates next to consumers, except where central land rents crowd it out. The T-sector is now small, but somewhat more crowded into the center and it is this that gives marginally higher central productivity and land rents compared to figure. 1.2.

While the N-sector is assumed not to have agglomeration economies, it does have costs of shipping goods to consumers (or equivalently, of consumers visiting the firms), and this gives a centripetal force to the location of this industry. This becomes the dominant force if T-sector agglomeration economies are small enough, in which case the locations of the N and T sectors are reversed as illustrated in Figure 8.2. The N-sector and a large mass of consumers now locate near the center, while the T-sector attains greatest employment density near the edge of the city.

The different patterns in figures 8.1 and 8.2 accord with perceptions of some urban areas with centers occupied by tradable sectors, while others have consumer services such as retail in the center, while tradable sectors – factories rather than tradable services -- are at the edge. Outcomes depend on building requirements, as well as the extent of agglomeration economies. Such outcomes may be associated with clusters of tradable activity at the city edge, this destroying the circular city structure, and possibly creating a multiplicity of equilibria, the subject to which we now turn.

---

[15] If $\varepsilon_w = \infty$ then there is, in general, specialization of cells. Any non-specialized cell would pay the same wage in both sectors.

Figure 8: Urban density with consumer services

i) Benchmark agglomeration: $\chi = 0.10$

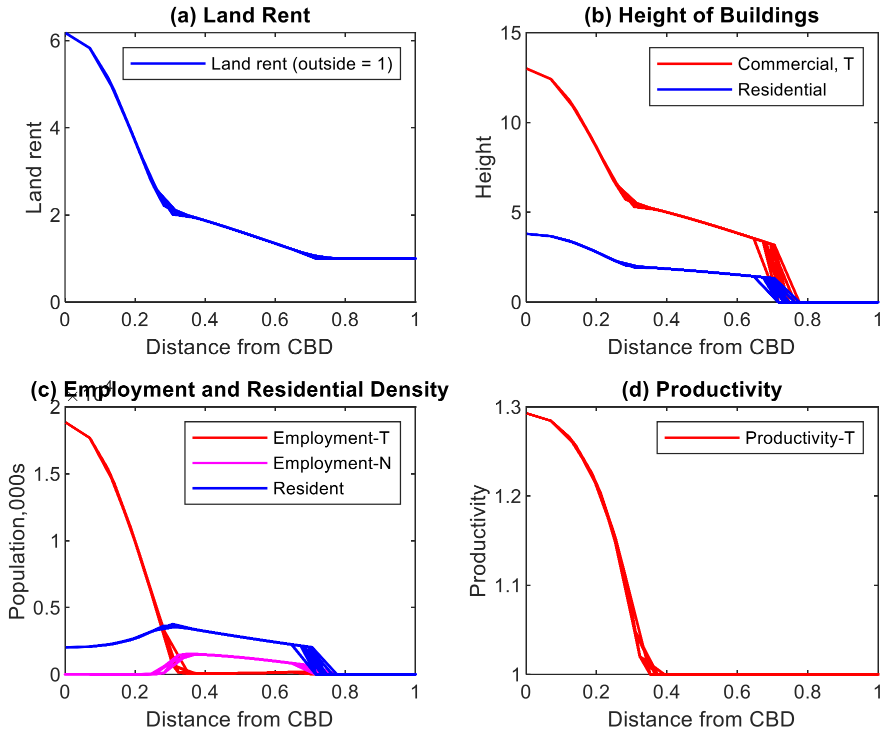

ii) Without agglomeration: $\chi = 0.0$

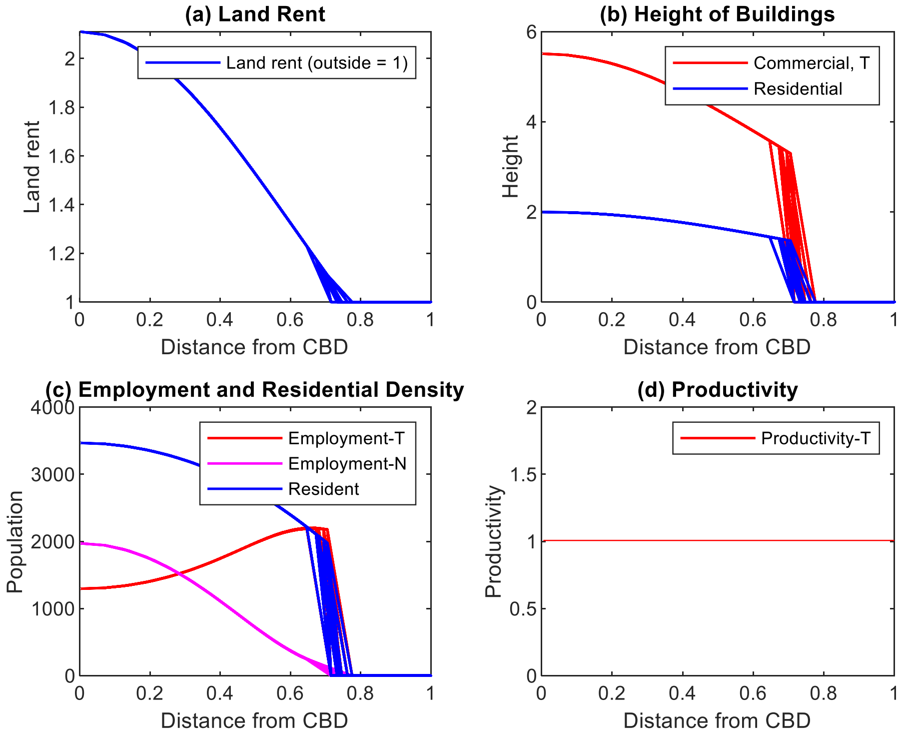

8. The emergence of polycentricity

The base case and the previous experiments were all in a monocentric city but, as the city expands, so polycentric structures can emerge. This is essentially because the relationship between productivity and spillover access is sharply concave, while commuting costs continue to increase with distance. Polycentricity can take different forms. Simulation indicates that in some cases the city is doughnut like, with employment concentrated in a ring around, but not including, the center. Agglomeration forces may cause the ring to break up, becoming crown like (employment lying in a ring, but of variable height and density), or form distinct sub-centers.[16] The example shown in figure 9 is for the two-sector model, in which all distances (travel costs for commuting and shipping goods) are doubled relative to figure 8; this experiment essentially makes the city 'larger', but we double transport costs rather than double population in order to keep the city within our geographical space.

The equilibrium illustrated in figure 9 is polycentric, with rotational symmetry of order 3. Employment in the N-sector is shown in 9a; this peaks in the city center (dark red). The dark areas are where there are very few N-sector firms; the dark ring round the edge is outside the city, and the three separate dark areas are where there are clusters of T-sector employment, associated with high rents and tall buildings, and crowding out N-sector firms. Figure 9b shows the corresponding distribution of residential population, with highest density in the bright yellow areas close to the city center, and lower further out and where residents are crowded out by high T-sector employment.

Figure 9a: Employment – consumer services.

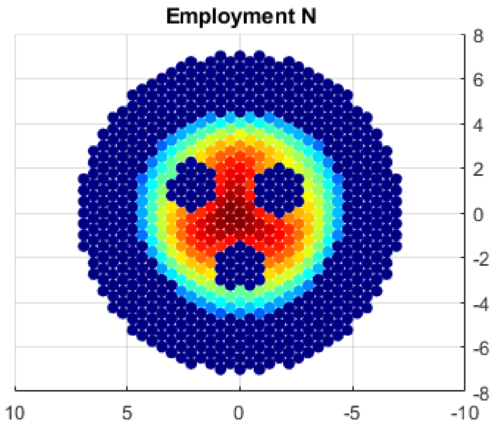

Figure 9b: Residential population

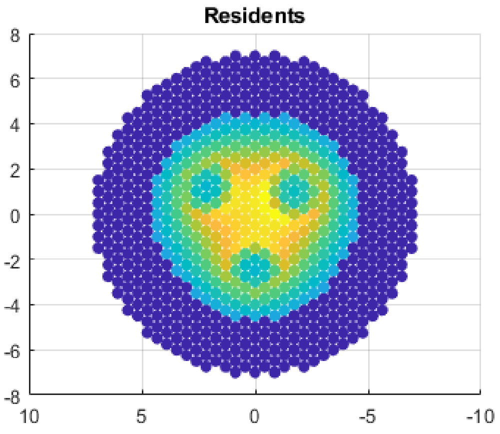

A number of further remarks are needed. First, this equilibrium is not unique. In a trivial sense, any rotation of the patterns shown in the figures is also an equilibrium. More fundamentally, it is possible that at the same parameters there could be an equilibrium with, two, or perhaps four, centers of T-sector activity. The particular configuration illustrated in figure 9 is computed by first finding the equilibrium with zero T-sector agglomeration economies, which is a flat-ish dome with rotational symmetry of infinite order, as in figure 8.2.

---

[16] The classic reference on poly-centricity is Fujita and Ogawa (1982). Lucas and Rossi-Hansberg (2002) derive the doughnut, although their 2-dimensional generalization of a 1-dimensional linear city rules out the possibility of clusters forming around the ring.

Agglomeration economies are then switched on and the pattern of employment is given a small random perturbation. The model then iterates to the new equilibrium, with workers' and firms' choices adjusting to new prices, wages, and rents, and land and product markets clearing at each step. The model runs until convergence is achieved at a new equilibrium, stable under this adjustment process. Importantly, while different runs of the model each start with its own random perturbation from the zero agglomeration starting point, the 3-fold pattern of centers is repeated (at these parameter values).[17] This approach enables some results to be obtained about the relationship between underlying parameters of the model and the form of polycentricity but the fact remains that, in many parameter ranges, the model is likely to produce multiple equilibria.[18]

Multiplicity of equilibria raises many issues, one of the most important of which is that comparative statics need not lead to continuous changes in outcomes. There is a possibility that some exogenous change may lead a new center to emerge – or cause an existing center to collapse. Such 'transformative' changes are often precisely what policy makers are seeking to achieve. The urban model provides a way of capturing these changes, and it is important that models do not rule out such changes by imposing conditions that ensure uniqueness of equilibrium. However, much work remains to be done before modeling techniques are able to capture the dynamics – of sunk costs and future expectations – that are a real world part of such changes.[19]

9. Concluding comments

The standard urban model, as illustrated by the benchmark case in this paper, provides a model of the city as a whole, with a full set of general equilibrium interactions in goods, land, and labor markets coordinating the actions of households, developers and firms. The model can isolate the forces that determine the shape and productivity of a city and is able to trace through the full impact of parameters, including those set by policy. Models of this type may well be able to assist policy makers in assessing outcomes of planned changes, particularly as the model can be calibrated to capture the circumstances of particular cities, as in Sturm et al. (2021).

If these models are to be used for practical and policy purposes they need to be thoroughly explored. Do they capture salient features of actual cities? Can the responses generated by the model be understood, and are they consistent with empirical observation of cities? How sensitive is the model to values of parameters, particularly those about which we have little information? What are the key areas where extension and refinement are needed?

To address these questions this paper sets out a 'pure' form of the model in which there are no idiosyncratic features of natural geography, enabling the basic model mechanisms to be seen relatively clearly. The paper contains two sorts of exercise. The first is to work through a series of comparative static exercises, asking how outcomes depend on a number of the exogenous inputs to the model. The second is to investigate several extensions to the basic model; many features could be added, and the paper looks at three that are of some importance.

---

[17] This is for reasons explored in Fujita et al. (1999).

[18] Explicit conditions for uniqueness of equilibrium can be found in simpler models, see for example Allen and Arkolakis (2014), Allen and Donaldson (2020).

[19] See Bird and Venables (2019) for a policy experiment of this type.

Some of the comparative static results are straightforward, for example that stronger agglomeration economies create higher rents and a denser cluster of employment in the CBD. Some involve trade-offs between forces pulling in opposite directions. Transport improvements enable concentration of employment and dispersal of residency, as raising income and thereby attracting population; their effects on CBD rent and density are ambiguous. Population growth raises productivity and CBD density, but also the average length and cost of commuting. Rents increase, but both real wages and average per capita income (the sum of wages and rents) fall with population size. Building technology and the ability to use land to produce floor-space are important for city productivity and income.

Policy measures have direct effects – for example costs saved in faster travel – and also determine the shape of the city. This can be quite fundamental as, for example a radial road network reinforces the urban center. The model indicates the direction and possible long-run magnitude of such effects and also gives estimates of their implications for income and welfare. These income changes can be quite large, both because the change in location of activities in the city are non-marginal, and because these changes interact with market failures, above all the externalities created by agglomeration. Thus, appropriate zoning that concentrates employment can raise productivity and welfare. The full equilibrium effects of a transport improvement may be several times larger than the direct 'user-benefits'.

Sections 6 – 8 of the paper explore extensions of the model. The most important is adding a consumer services sector, representing activities including retail, domestic services, hospitality, maintenance services, and the informal sector. There are intra-city transport costs in accessing and supplying consumer services and these costs create their own agglomeration force, favouring central locations and hence possibly forcing employment in tradable sectors to the edge of the city. This is one factor – but not the only one – that can create poly-centricity, as demonstrated in section 8.

The extended model contains multiple drivers of city shape. There are the spillovers that are captured in the benchmark model, the shape of transport systems (e.g. the centripetal force created by radial roads), and the effect of transport costs in consumer services. All of these operate through the choices made by households of where to live and where to work. We show (section 6) both the sensitivity of results with respect to parameters of the household choice function and, more fundamentally, the powerful additional clustering forces that are created if employment considerations enter households' location choices not just by wages offered in each place, but also by the number of jobs in each place. These are all forces for clustering, and the contribution of each needs to be better understood.

Finally, this paper is an 'exploration'; numerous further questions need to be addressed (for example the interactions between different policy measures) and extensions made. The most challenging are to do with multiple equilibria and the possibility of securing 'transformative change' through policy measures, particularly so in a dynamic environment where investments are long-lived and expectations matter.

References

Ahlfeldt, Gabriel M., Stephen J. Redding, Daniel M. Sturm, and Nikolaus Wolf (2015), 'The economics of density: Evidence from the Berlin Wall', *Econometric*a, 83(6), 2127-2189

Allen, T., and C. Arkolakis (2014): 'Trade and the Topography of the Spatial Economy,' *The Quarterly Journal of Economics*, 129(3), 1085-1140.

Allen, T., and D. Donaldson (2020): 'Persistence and Path Dependence on the Spatial Economy' https://sites.google.com/site/treballen/research

Bird, J. and A.J. Venables (2019), 'Growing a Developing City : A Computable Spatial General Equilibrium Model Applied to Dhaka. Policy Research Working Paper. 8762. World Bank, Washington, DC.

Bird, J. and A.J. Venables (2020), 'Land tenure and land-use in a developing city; a quantitative spatial model applied to Kampala Uganda' with J. Bird, *Journal of Urban Economics*, 119, 103268.

Duranton, G. and D. Puga, (2004). 'Micro-foundations of urban agglomeration economies', in: J. V. Henderson & J. F. Thisse (eds.), *Handbook of Regional and Urban Economics*, volume 4, chapter 48, pages 2063-2117, Elsevier.

Duranton, G. and D. Puga, (2015). 'Urban Land-use', in: G. Duranton, J. V. Henderson and W. Strange (eds.), *Handbook of Regional and Urban Economics*, volume 5, chapter 8, pages 467-560, Elsevier.

Fujita, Masahisa, Paul Krugman, and Anthony J. Venables. (1999) 'The spatial economy: Cities, regions, and international trade' Cambridge MA: MIT Press.

Fujita, Masahisa and Hideaki Ogawa, (1982) "Multiple equilibria and structural transition of non-monocentric urban configurations", *Regional Science and Urban Economics*, 12.2, 161-196.

Heblich, S., S. Redding and D. Sturm, (2020) "The Making of the Modern Metropolis: Evidence from London" *Quarterly Journal of Economics*, 135(4) 2059 – 2133

Henderson, J. V., Regan, T., and A.J. Venables, (2021), 'Building the city: from slums to a modern metropolis'', forthcoming, *Review of Economic Studies*

Lucas, R.E. and E. Rossi-Hansberg (2003), 'On the internal structure of cities', *Econometrica*, 70, 4, 1445-1476.

Monte, F., S. J. Redding, and E. Rossi-Hansberg, (2018) "Commuting, migration, and local employment elasticities." *American Economic Review* 108.12: 3855-90.

Redding, S. J., and M.A. Turner, (2015). 'Transportation Costs and the Spatial Organization of Economic Activity', in: J. V. Henderson & J. F. Thisse (eds.), *Handbook of Regional and Urban Economics*, volume 4, chapter 20, pages 1339-1398, Elsevier.

Redding, S.J. and Rossi-Hansberg, E.A. (2017) 'Quantitative spatial economics', *Annual Review of Economics*, 9(1), 21-58.

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

Sturm, D.S, K. Takeda and A.J. Venables (2021) 'How useful are quantitative urban models to analyse cities in developing countries? Evidence from Dhaka'

Tsivanidis, N. (2019), 'Evaluating the Impact of Urban Transit Infrastructure: Evidence from Bogotá's TransMilenio', unpublished.