A scale-free transportation network explains the city-size distribution

Berliant, Marcus and Watanabe, Hiroki

Washington University in St. Louis, Lamar University

23 October 2014

A Scale-Free Transportation Network Explains the City-Size Distribution *

Marcus Berliant[†] and Hiroki Watanabe[‡]

October 23, 2014

Abstract

Zipf's law is one of the best-known empirical regularities in urban economics. There is extensive research on the subject, where each city is treated symmetrically in terms of the cost of transactions with other cities. Recent developments in network theory facilitate the examination of an asymmetric transport network. In a scale-free network, the chance of observing extremes in network connections becomes higher than the Gaussian distribution predicts and therefore it explains the emergence of large clusters. The city-size distribution shares the same pattern. This paper decodes how accessibility of a city to other cities on the transportation network can boost its local economy and explains the city-size distribution as a result of its underlying transportation network structure.

Keywords: Zipf's law, city-size distribution, scale-free network
JEL classification: R12, R40

1 Introduction

Cities develop in relation to other cities rather than in a vacuum. What we consume in a city differs from what we produce in a city. The gap between the range and scale of production and consumption at city level is bridged by the transportation network, over which cities trade their products with others. The transportation network, in turn, does not coordinate cities uniformly. Some cities have only limited connections while others receive many links from cities across the country, both large and small, near and far away. The fate of city's economy, and by extension its population size, is more or less conditioned by how it is positioned (inadvertently or otherwise) in the overall interurban network of cities and how accessible it is from others. We will show that the city-size distribution is the result of a particular class of network that our economy installs on itself for interurban trading purposes, namely, a scale-free network.

The way we treat the transportation network has been rather naïve and simplistic. Most existing models of city-size distribution implicitly or explicitly assume a completely isolated graph (Figure 1(a)) or complete graph (Figure 2(a)). Each

*This project received a grant from the Center for Research in Economics and Strategy (CRES) at the Olin Business School, Washington University in St. Louis. The authors thank Professors Wen-Tai Hsu, Sukkoo Kim, John Nachbar, Jody O'Sullivan and Victor Wickerhauser for their helpful comments and advice.

†Department of Economics, Washington University, Campus Box 1208, 1 Brookings Drive, St. Louis, MO 63130-4899 USA. Phone: (314) 935-8486, Fax: (314) 935-4156, e-mail: berliant@artsci.wustl.edu

‡Department of Economics and Finance, Lamar University, P.O. Box 10045, Beaumont, TX 77710 USA. Phone: (409) 880-8657, e-mail: watanabe.wustl@gmail.com.

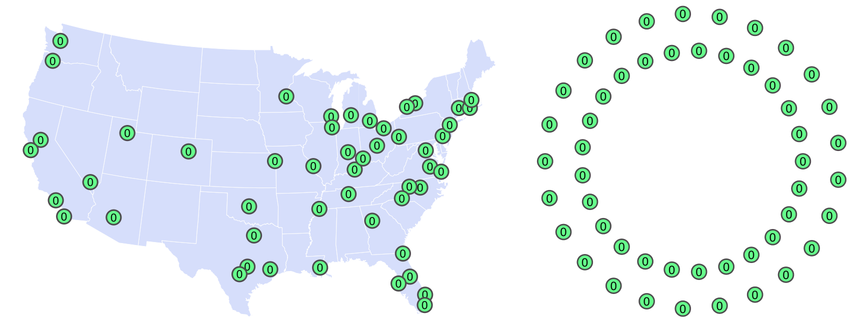

(a) The United States according to completely isolated graph with the 50 largest cities.

(b) Completely isolated graph with 50 nodes.

Figure 1. Completely isolated graph

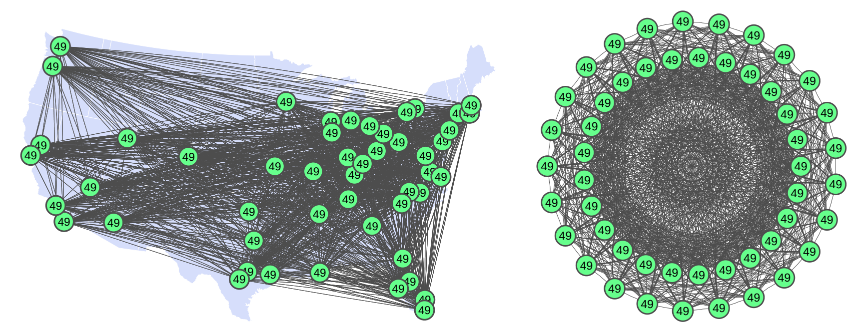

(a) The United States according to complete graph with the 50 largest cities.

(b) Complete graph with 50 nodes.

Figure 2. Complete graph

node represents a city and a link represents a route available for shipment. The number inside a node counts its degree, i.e., the number of edges or routes each node has. Commodities cannot be shipped at all on a completely isolated graph, but they can be shipped anywhere in a single step from any city on a complete graph. Either way, the resulting equilibrium will be an even split of population among the cities, which does not match the actual city-size distribution. To explain the city-size distribution, we have sought a source of variation other than what the nexus of interurban relationships has to offer. Some use a completely isolated graph (e.g., Eeckhout [Eec04]). Others such as Duranton [Dur06], Rossi-Hansburg and Wright [RHW07] or the New Economic Geography [FKV99] engage a complete graph as the transport structure, when in fact, transaction and/or communication between hub cities is much easier than between cities on peripheries. Behrens et al [BMMS13] introduce a more lifelike representation of transportation cost in that the delivered price depends on a particular city pair. The price differential reflects monopolistic pricing rather than the underlying transportation network structure, which is still an (ex-ante) complete graph. The literature usually introduces a tiebreaker in the form of externalities, random growth, economies of scale or scope to replicate the actual city-size distribution.

In practice, transportation cost differs greatly depending on where you are and where you are headed. We will drop the assumption that our economy operates on a complete or completely isolated graph and see how much explanatory power network structure exerts as the engine of local economies of various sizes.

The transaction pattern between any two cities affects both the way cities are populated and the overall city-size distribution. Cities are tied together in various ways both topologically and economically. Some cities function as an intersection of major transportation routes and they trade and process commodities frequently in large volume. Others are less active in the interurban exchange of commodities. Differences among cities in terms of exchange patterns reverberate in the city-size distribution. Cities heavily interrelated to many others are likely to grow due to increased economic activities, whereas cities with sparse connections to a limited number of cities are liable to remain small in size. Those small cities, however, will not be completely wiped off the map.

1.1 Cities on a Network

Intercity exchange patterns like Figures 1(a) and 2(a) are best described by a network with cities as a set of vertices and traffic by edges as in Figures 1(b) and 2(b). In this regard, network theory is indispensable when constructing a model of cities in the nationwide economy.

The recent seminal work by Barabási and Albert [BA99] has revitalized network theory. Classical network theory pioneered by Erdős and Rényi [ER59]'s model (ER network) cannot explain the emergence of a cluster or hub in a network, which we observe in most real social networks. In a classic random graph, each node is linked with an equal probability to any other and lacks distinctiveness, for the number of pre-existing links does not matter in forming a network. Barabási and Albert (BA) add a dynamic feature and preferential attachment to the classical random graph model so that the nodes are no longer ex-ante identical. Some nodes gather lots of links while others are wired to just a few. The model has been applied to many fields, including the emergence of web science, and has produced an improved description of the organization and development of networks. Most real-world networks have one thing in common: the resulting distributions of links are scale-invariant, that is, the distributions have fat tails. We can find nodes with an extremely large number of links rather easily with these networks compared to a classical random graph.

The city-size distribution shares the same pattern of scale invariance: the distribution of the 100 largest cities follows the same distribution as the one for the 1000 largest cities and so on, a property known as a power law, and in particular, Zipf's law in the city-size literature. We expect that the degree of a city is positively related to its population. And for that reason, we imagine that our economy is based on a BA network rather than an ER network. This turns out to be correct, but selection of the appropriate network structure depends on exactly how node degree is related to city size. We will decode their relationship in Section 3.8.

The urban economic application of network theory is in its very early stage of development and there is much room for advancement. Interaction between individual cities has not caught much attention so far. Our goal in this paper is to bring to the fore the interaction between transportation network structure and the city-size distribution. With this goal in mind we introduce (asymptotic) techniques from network theory and merge them with a tractable economic model in a new way. We do not intend this work to be the last word on this topic, but merely a suggestion of a first step into a bigger research program.

1.2 Some Transportation Networks Are Scale Free

Our economy operates on various modes of transportation and each mode comes with distinct network structures. Take a highway and airline network for example. Figures 3(a) and 4(a) are schematic representations of the Interstate System and a typical airline route map for the 50 largest US cities. Apparently, a network composed of the Interstates does not share its structure with that of airlines. The Interstate will remain relatively intact when we take away New York, Houston and Cleveland. On the other hand, it would prove devastating if we did the same

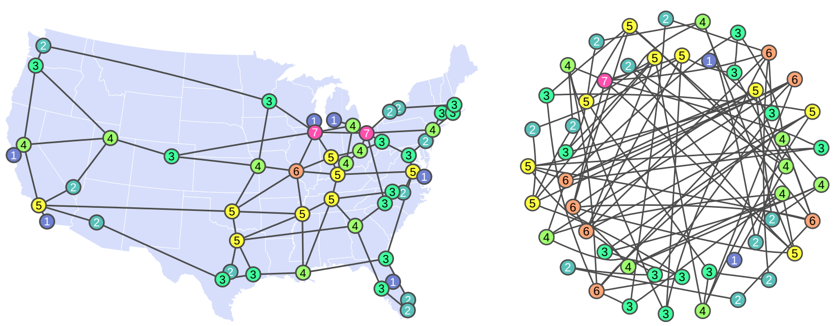

(a) The Interstate route map (abridged)

(b) ER network with 50 nodes.

Figure 3.

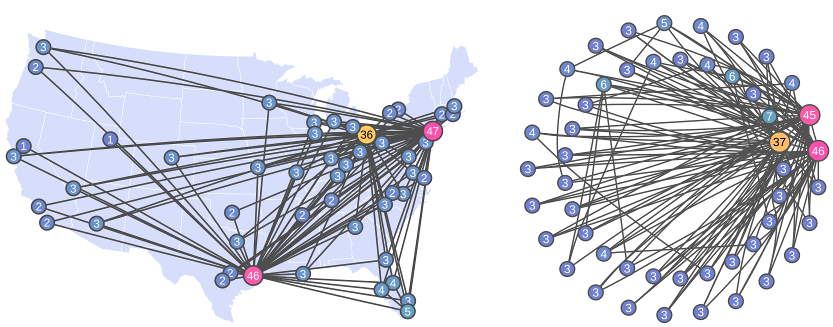

(a) A typical airline's route configuration (pre-merger Continental Airlines)

(b) BA network with 50 nodes.

Figure 4.

to the airline network (cf. [BB03]). More broadly, there is not much variance in the degree of nodes in the Interstate network, whereas the airline network has a limited number of heavily wired cities. The BA network (Figure 4(b)) explains the latter network better, as it follows a power law.

It should be noted, however, that what is geographically visible may *not* represent the real network that our economy relies on in effect. The Interstate network exhibits an ER-type topology as in Figure 3. Nonetheless, the economy may operate a transportation network of a scale-free class on it. Shipment from Memphis has to go through St. Louis even if its final destination is Chicago. In this case Memphis is connected to Chicago in a single step rather than in two steps via St. Louis. For a carrier making Chicago-bound shipment from Memphis, St. Louis (a seeming layover node) is no different from the cornfield they pass through along the way (just a part of the edge), in that neither one of them add anything to the shipment. An economically relevant network is buried beneath the easily noticeable surface network and we do not want to confuse one with the other.

It is very important to note here a difference between the literature on dynamic social network formation and transportation networks. In the standard economics literature on social networks, for example Mele [Mel11] or Christakis et al [CFIK10], it is the individual agents, represented by nodes, who make decisions about forming links among themselves. In contrast, the nodes of a transport network are cities. Typically, it is not the cities or their agents who make decisions about forming links. Rather, it is another agent who controls an entire networks, for example the federal government in the case of highways or airlines in the case of an airline system.

1.3 The City-Size Distribution Is Scale Free Too

The city-size distribution has a distinct feature. Figure 5 plots the frequency of the city-size distribution from US Census 2000. It is only when we take the log of population (Figure 5(b)) that the distribution exhibits resemblance to a familiar Gaussian distribution. Black and Henderson [BH03] and Soo [Soo05] explain how widespread scale-free distributions are in urban economics[1]. Under the scale-free distribution, the arithmetic mean (Hillsboro, TX in Figure 5) becomes less interpretive and the geometric mean (Sutton, NE) takes over the role of the average in the conventional sense.

The fat-tailed distribution also makes its appearance on a map. Figure 6 illustrates the population density of each metropolitan and micropolitan statistical area (MSA and μSA, collectively referred to as Core Based Statistical Area, CBSA) in the United States in 2000. Most of the cities have a low density and are painted in blue; there are only few cities that are green and only two cities are colored in red. If the city-size distribution followed a Gaussian distribution or Poisson distribution with a large mean[2], most of the cities should be green and only a few should be in blue or red. Just as for the airline network in Figure 4(a), if we take away the ten largest US cities, we will leave more than a quarter of urban population unaccounted for.

[1] Scale-free distributions are commonplace in the socioeconomic realm. It seems that something of an additive nature presides over natural phenomena, leading to a Gaussian distribution, and something of multiplicative nature (cf. [LSA01]) is at work among socioeconomic phenomena, leading to a scale-free domain. We study the latter.

[2] As in the degree distribution of an ER network.

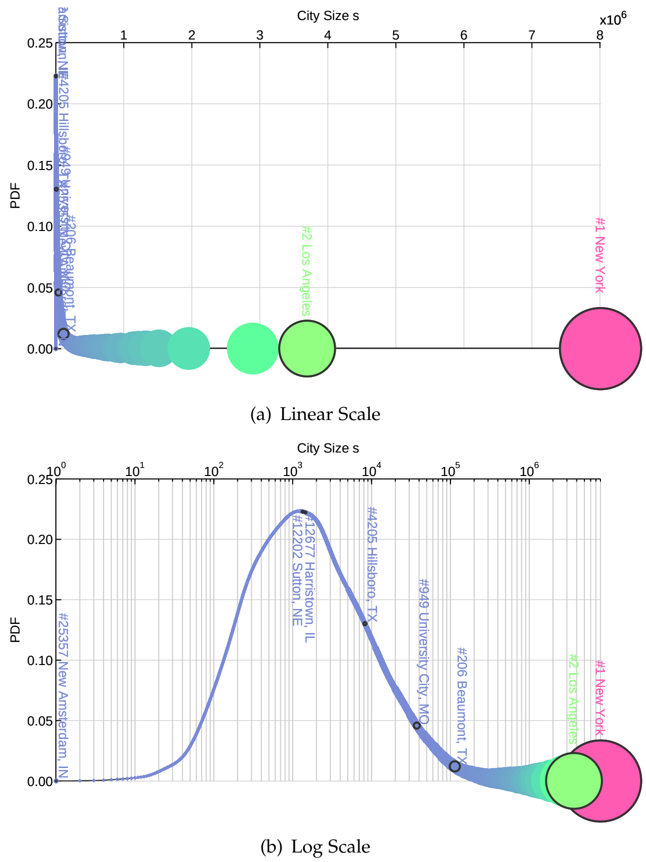

Figure 5. Frequency plot of the city-size distribution. Dots are size proportionate. See Table 1 for explanation of the cities selected in the figure. *Data source*: US Census 2000.

Our main findings are as follows. City sizes are positively related to their degree. A city with a high degree has good accessibility to other cities. Reduced transportation cost makes the city's product inexpensive and stimulates a large demand. As a consequence, the city creates large-scale employment. However, a marginal increase in degree contributes less to the city size as the degree increases. If a city is well-connected, then adding a new link to the city will not increase accessibility much because the city is already readily accessible from other cities through the existing grid.

We test implications of our model with Belgian and US data. The BA network leads to a result comparable to existing models, whereas the ER network fails to replicate the empirical city-size distribution. This confirms that the BA transport network is more consistent with reality.

The rest of the paper is organized as follows. In Section 2, we will go over the two types of network structures mentioned above as a preamble to the next section, where we introduce and develop a model of spatial equilibrium with a transporta-

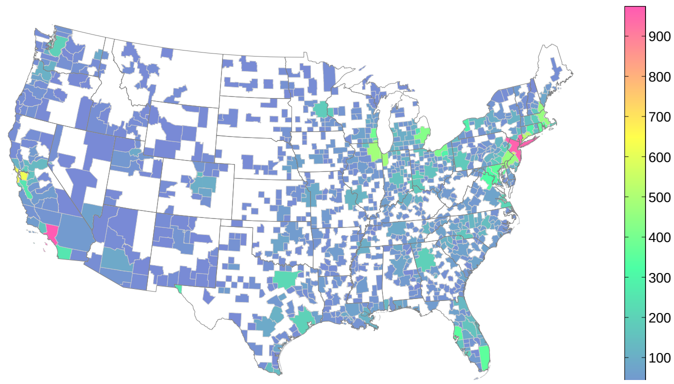

Figure 6. Population density by CBSA (persons/km^2). *Data source*: Census 2000.

tion network woven into it. Particularly, in Section 3.8, we will connect the network structure to the city-size distribution. In Section 4, we verify the prediction of our model with data before we draw conclusions from our project in Section 5.

2 Preliminaries

We will briefly review how ER and BA networks are built and examine the qualitative differences in terms of their degree distributions before we apply them to transportation networks.

2.1 ER Networks

The ER network is the simplest random graph of all. A pair of nodes are connected with a fixed connection probability. A completely isolated graph illustrated in Figure 1 and complete graph illustrated in Figure 2 are the special cases of the ER network where connection probability is zero and one, respectively.

The degree distribution of an ER network follows a Poisson distribution. The important feature is that the degree distribution is concentrated around its arithmetic mean[3] and we rarely observe a city with an exceedingly large degree. All pairs of nodes share the *same* ex-ante connection probability, which leads to a small variance, and the network is *egalitarian* in that sense.

2.2 BA Networks

The degree distribution of most real network structures does not follow a Poisson distribution. Rather, it follows a power law. This class of networks is called scale free. There are a number of proposed generative models that lead to power-law degree distributions (see Section VII of Albert and Barabási [AB02] for a review). To get a sense of how power-law type behavior emerges, consider the BA model

[3]Recall that arithmetic mean does not mean much for scale-free distributions like the city-size distribution or a BA degree distribution.

[BA99] for example. Two major characteristics of BA model are growth and preferential attachment. The model sets off with a complete graph of a fixed number of nodes as a starting grid. New nodes with edges will be added sequentially to the existing network (growth).

As we can see from this mechanism, in general, older nodes are likely to gain an excessively large number of edges. The rich get richer because they are already rich (known as the Matthew effect). The rest of the nodes are merely mediocre in terms of degree. They are poor because they are already poor. This type of variance in degree hardly arises with an ER network. That is, New York City will not happen if the links are formed uniformly at random. Compare BA network Figure 4(b) to ER network Figure 3(b). BA network is *not* egalitarian, as connection probability depends on the number of acquired edges, which is path dependent. We shall also emmploy the network structure of Jackson and Rogers [JR07] that contains both the ER and BA types of networks as special cases. Details shall be provided in Section 3.8.

3 Model

We propose a model where the trading costs of commodities among cities are explicitly specified. The city-size distribution is derived as a result of gains from trade and the underlying transport network configuration.

3.1 Location-Specific Commodities

There are J cities in the economy, with index j. A city is defined as a geographic entity within which it produces the same commodity and from within which the geodesic paths (the shortest path on the network) to any other city in the country have the same length. If Adam and Beth both live in St. Louis, then they have the same shipping cost schedule to everywhere in the nation. We know they are in different cities if Adam pays a 10% shipping charge to San Francisco and a 5% charge to Minneapolis, whereas Beth pays a 10% charge to San Francisco but an 8% charge to Minneapolis. The endogenous population of city j is given by s_j and in total, there are

$$\sum_{j=1}^{J} s_j = S \tag{1}$$

households in the economy. Each household supplies a unit of labor inelastically. City j produces consumption commodity c^j in a competitive environment. We assume that technology exhibits constant returns to scale and that one unit of labor produces one unit of commodity. In what follows a superscript denotes a city of production or origin, whereas a subscript denotes a city of consumption or destination.

The delivered price of commodity j in city i is denoted by p_i^j. The value of marginal product $p_j^j \cdot 1$ coincides with the local wage w^j in equilibrium:[4]

$$p_j^j = w^j \tag{2}$$

[4]Note that p_j^j denotes the mill price.

Consumer preferences are represented by a Cobb-Douglas utility function of the form $u(c_i) = \frac{1}{J}\sum_{j=1}^{J}\log(c_i^j)$. The set of consumption bundles is constrained by the budget $w^i \geq \sum_{j=1}^{J} p_i^j c_i^j$.

3.2 Network Infrastructure and Delivered Price

The economy has a network infrastructure $\Gamma = (V, E)$, where $V = \{1, \cdots, J\}$ denotes the set of vertices representing each city and E denotes a set of edges. For example a completely isolated graph in Figure 1 is given by $\Gamma = (\{1, \cdots, 50\}, \emptyset)$ and a complete graph in Figure 2 by $\Gamma = (\{1, \cdots, 50\}, \{\{i, j\} : 1 \leq i < j \leq 50\})$. All the traffic flow will follow Γ. We assume that the network is unipartite (i.e., there is a path between any pair of nodes) to avoid multiple equilibria. Whereas consumers in city j can consume any commodity in the economy, they have to incur an extra iceberg transport cost to consume commodities brought in from other cities. Transportation cost piles up as a commodity travels from city to city along the path. To describe the exact transport cost structure, we define a metric $l_j^i : V \times V \to \mathbb{R}_+$ to measure a geodesic length between node i and j given Γ. The delivered price of commodity j shipped to city i is given by

$$p_i^j = \tau^{l_i^j} p_j^j, \tag{3}$$

where $\tau(\geq 1)$ marks the iceberg transportation parameter. We use the iceberg transport technology, standard in urban economics, for tractability reasons.[5] If you dispatch τ units of commodity to your neighboring city, one unit of it will be delivered and the rest melts en route. The delivered price snowballs as the package travels from one city to another and the initial mill price is inflated by $\tau^{l_j^i}$ by the time the package reaches its final destination l_j^i steps over. We assume that all the links share the same value of τ. The large fraction of transportation cost is a location-invariant fixed cost. Having τ dependent on each link will not add much to our analysis but will make our equilibrium analytically insolvable.

3.3 Equilibrium

Simple calculations yield the Marshallian demand for commodity c_i^j:

$$\varphi_i^j(p_i^1, \cdots, p_i^J, w^i) = w^i(\tau^{l_i^j} p_j^j)^{-1} J^{-1}.$$

The aggregate demand for commodity j is the sum of demand from all the cities in the country: $C^j(p, w) := \sum_{i \in V} s_i \varphi_i^j(\cdot)$.[6] Recalling that each household supplies one unit of labor inelastically and one unit of labor produces one unit of output, the commodity market j clears when

$$s_j = C^j(p, w) = \left(p_j^j\right)^{-1} J^{-1} \sum_{i \in V} s_i w^i \tag{4}$$

[5]For detailed discussion, see McCann [McC05].

[6] This expression may seem incredulous at first, for it does not include τ. A large τ discourages demand but it also means that firms have to ship more commodities. A large portion of shipment will melt on its way. They cancel each other in equilibrium. This propitious cancellation may not occur with other preference specifications.

The indirect utility function is given by

$$\begin{aligned} v(p_i^1, \cdots, p_i^J, w^i) &= \tfrac{1}{J} \textstyle\sum_{j=1}^J \log \varphi_i^j(\cdot) \\ &= \log w^i - \log J - \tfrac{1}{J} \textstyle\sum_{j \in V} \log p_j^j - a_i \log \tau, \end{aligned}$$

where

$$a_i := \frac{1}{J} \sum_{j=1}^{J} l_i^j = \langle l_i \rangle \tag{5}$$

is a remoteness parameter, or an average geodesic length from city i, where $l_i : j \mapsto l_i^j$. In what follows $\langle x \rangle$ denotes the average value of x. The parameter measures how hard it is to reach city j from other cities in the economy. The higher the value is, the more remote the city is because we have to go through many links to get there. We will explore the role of accessibility later.

Free mobility of consumers implies

$$v(p_i^1, \cdots, p_i^J, w^i) = v(p_j^1, \cdots, p_j^J, w^j) \tag{6}$$

for all $i, j \in V$ in equilibrium.

The equilibrium $(s_1, \cdots, s_J; p_1^1, \cdots, p_J^J; w^1, \cdots, w^J)$ satisfies (1), (2), (4) and (6). Utility equalization (6) leads to

$$\log p_i^i - \log p_j^j = (a_i - a_j) \log \tau. \tag{7}$$

Equation (7), together with (4), implies $s_j = \tau^{a_i - a_j} s_i$. With the population condition (1), we obtain the city-size distribution

$$s_i = \frac{S}{\tau^{a_i} \sum_{j \in V} \tau^{-a_j}}. \tag{8}$$

3.4 How Does a Network Break Symmetry?

An obvious implication of (8) is that cities with better accessibility have larger equilibrium population. Naturally, we are tempted to conclude that the entire population will collapse into the city with the best accessibility and the rest of the cities will be completely vacated. As it turns out, this is not the case. The city-size distribution will not become degenerate. Let us break down (8) both mathematically and economically to see why.

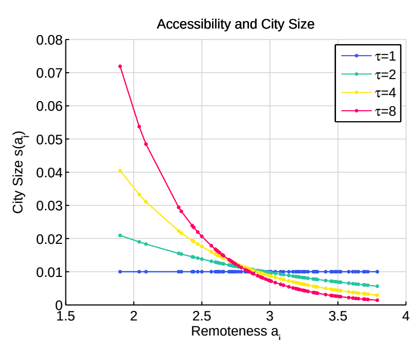

Figure 7.

First, let us recast the relationship (8) to explore how accessibility translates to the population of a city. We can rewrite (8) as $s(a_i) = \langle s \rangle \tau^{-a_i} / \langle \tau^{-a} \rangle$, where $\langle s \rangle := S/J$ is a base city size and $\langle \tau^{-a} \rangle := \sum_j \tau^{-a_j}/J$ gives the average of τ^{-a_j}. The city size spreads around the canonical size $\langle s \rangle$. A better accessibility (i.e., small remoteness value a_i) contributes to the city by augmenting the baseline size $\langle s \rangle$ by a factor of $\tau^{-a_i} / \langle \tau^{-a} \rangle$. The multiplier is large when τ^{-a_i} is greater than the national average $\langle \tau^{-a} \rangle$ and vice versa. Furthermore, the multiplier grows *more than proportionally* as the city's accessibility improves as can be seen in Figure 7. The multiplier τ^{-a_i} is monotone decreasing and convex in a_i. Does this mean New York City sweeps away all the population off the rest of the cities? — Not really. And it calls for an economic exposition of (8) to see why.

Although restricted accessibility of a city raises its delivered prices, demand for its produced commodity does not cease to exist. Eliminating a commodity from the basket will punish consumers a lot. They appreciate variety and missing a single variety will push the utility level down to negative infinity. Workers in a poorly connected city will have to pay a high price for imported commodities due to a poor network infrastructure, but they are compensated with a high nominal wage, as indicated by the wage (2) and utility equalization (7). These two equations imply that the mill price (and ultimately, the nominal wage) is positively related to the average geodesic length $\langle l_i \rangle$ from city i in equilibrium, i.e., a sparsely connected city has a high mill price. The prices adjust to make it worth living in small cities in equilibrium. The scale of local production is small, but each commodity is sold high to make up for an increased cost of living due to remoteness and the resulting costly transport.

Variance in city sizes is solely due to the structure of the network. The abovementioned trade-off entails two counteracting forces. The agglomerative force is heterogenous accessibility, which tends to spread out the city-size distribution. The dispersion force is preference for variety, which tends to push the distribution back to a collection of equal-sized cities.

There are alternative ways to derive city size with a tractable economic model, particularly for the dispersion force. In this model, location-specific commodity production drives dispersion, as a bundle of all goods is desired by consumers. An alternative model would use another natural dispersive force, say housing or land markets. If we had just a few produced commodities (say one for illustration), then Starrett's Spatial Impossibility Theorem (Fujita and Thisse [FT02], Ch.2) applies, and we would have an autarkic equilibrium where no commodity is transported.[7] Yet another alternative is to introduce a congestion externality, but then the model begins to look more complicated and, at the same time, arbitrary.

Obviously, this trade-off disappears and there will be no variance in city sizes if the agglomerative force is removed. This can happen when shipment becomes costless (to be discussed in Proposition 3.1) or network structure becomes redundant, that is, if it turns into a complete graph. Although we introduced a location-specific technology, commodities are symmetric. Technology is linear everywhere. Consumer preferences are identical and they put the same weight on each commodity. If we take the network structure out of the equation, the resulting equilibrium is such that all the cities share the same size $\langle s \rangle$ and every household consumes an equal portion of all the commodities available.

3.5 Transportation Cost Skews the City-Size Distribution

Along with remoteness a_i, transportation cost τ plays a leading role in the determination of the city-size distribution. Depending on its magnitude, τ can nullify or amplify the influence of a network structure over the economy. Figure 7 compares

[7] Starrett's Theorem makes no assumption about the transport network or transport cost.

the relationship between accessibility and the city-size distribution under different transportation costs.

In the extreme situation where shipment is free ($\tau = 1$), all the cities will be of an equal size regardless of the network structure. The city size $s(a_i)$ becomes constant against a_i (see the blue line in Figure 7). The network becomes a complete graph in effect, because the delivered price will be the same no matter how long the geodesic length is. For $\tau > 1$, city size (8) becomes a strictly convex function of remoteness.

The transportation network Γ starts to sink in as τ grows. A large τ implies that the geodesic length exerts a more dominant influence on the size of a city. With a small value of τ, a city with good accessibility does not distinguish itself well from other cities because the effect of path length is limited due to low transportation cost. On the other hand, if shipping is costly, a city with a good accessibility benefits from a low a_i value because high transportation cost amplifies the effect of accessibility. In other words, a high transportation cost reveals the network structure and projects the network Γ onto the city-size distribution in a more pronounced, clear-cut manner than with a low transportation cost. As a result, holding the remoteness distribution constant, large τ skews the city-size distribution and makes the emergence of disproportionately large hubs more likely. To measure how the cost of transportation τ bends the city-size distribution, consider a measure

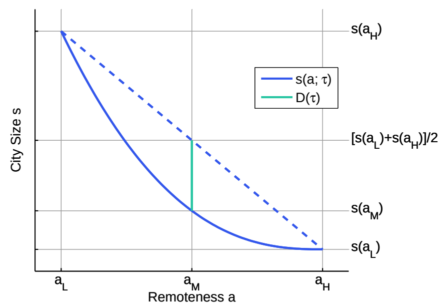

Figure 8. $D(\tau)$ measures the convexity of $s(a_i)$. The midpoint $(a_H + a_L)/2$ is given by a_M above.

$$D(\tau) = \frac{s(a_H) + s(a_L)}{2} - s\left(\frac{a_H + a_L}{2}\right),$$

where a_H and a_L are the highest and lowest remoteness of a given network. The first term is the average of the smallest and the largest city whereas the second term is the city size of average remoteness. For a given distribution of remoteness a_i, $D(\tau)$ measures the convexity of $s(a_i)$, which gauges how spread out the distribution of city size $s(a_i)$ is for each τ. See Figure 8. When $\tau = 1$, $s(\cdot)$ lays flat and $D(\tau) = 0$. As τ grows, $s(\cdot)$ bends more and $D(\tau)$ grows accordingly as can be seen in Figure 7.

We confirm the observation above as follows:

Proposition 3.1 Transportation Cost Skews the City-Size Distribution

Suppose that the economy has a unipartite network Γ. The city-size distribution s_i is a convex function of remoteness a_i for $\tau \geq 1$. Moreover, the degree of convexity measured by the size difference $D(\tau)$ between the city of average size and the city of average remoteness increases with τ.

Proof. See Appendix A.1. □

3.6 Geodesic-Length Distribution

The city-size distribution (8) depends on the distribution of remoteness (5), which, in turn, rests on the distribution of geodesic length. While most of the research on network topology is focused on *mean* intervertex distance ([NSW01], [FFH04], [ZLG+09]), what we need here is the geodesic length between *individual* nodes. Mean intervertex distance comes in handy when we gauge how efficient a network is, but we are not here to see if the transportation network that our economy relies on is optimally configured (that would be another paper). We would like to derive the city-size distribution, not the average size of cities or the remoteness thereof.

There is not much research that looks into the geodesic length between each pair of nodes. At the time of writing, the analytical form of geodesic length between individual nodes is yet to be discovered[8]. There is an attempt to track down the geodesic length by guessing the analytical form from sequentially generated, fractal-like networks reverse-engineered from a Pareto degree distribution ([DMO06]), which we cannot use because our distribution (14) is not a Pareto distribution.

Hołyst et al [HSF+05] take a different approach to derive an intuitive solution for a wide range of network types. They measure the expected geodesic length between any pair of nodes i and j as follows:

$$l^i_j = A - B\log(k_i k_j), \tag{9}$$

where $A := 1 + \log(J\langle k\rangle)/\log\kappa$ and $B := (\log\kappa)^{-1}$. The number k_i denotes the degree of node i. Rearrange the nodes so that we have a tree with node i as its root. The average number of children is called an average branching factor and denoted by κ. For more details see Appendix A.2.

Although [HSF+05] does not provide a formal proof of (9), but rather is based on a heuristic, it appears to be the best we can do given the current state of network theory. Zhang et al [ZLG+09] provide an analytical background for the mean intervertex distance for a special case. We hope that its extension to individual distances will become available in the near future.

Meanwhile, (9) proves to be quite useful in translating a network structure into economic context without loss of generality. A path length is a *global* property whereas a degree is a *local* property. We cannot compute the individual geodesic path unless we compare all the possible pathes between a city pair of interest and pick the shortest one, which calls for a systemic search all across the board. The geodesic path thus obtained is too specific to the particular network in question and does not have wide implications beyond the specific network itself. Degree is much easier to compute because we do not have to launch a nationwide search for it, and the degree distribution is readily available for a wide range of networks. Equation (9) succinctly writes a global property (a path length) in terms of the analytically manageable local property (a degree). It implies that the path length will be short if your city and/or your destination city have many edges to choose from to begin with and/or to end with. This abundance in selection should save you from being thrown to circuitous paths, and vice versa when your degree is small. Absent this conversion of the global property into the local property, we would not be able to describe a general relationship between degree and city size,

[8] The one for the average intervertex separation has already been brought out into the open. Cf. [NW99], [NMW00], [ZLG+09].

when in fact, there is an obvious symbiotic interaction between them waiting to be investigated.

3.7 City-Size Distribution

From (9), remoteness (5) is written as

$$a_i (= \langle l_i \rangle) = A - B \log k_i - B \langle \log k \rangle. \tag{10}$$

We observe that accessibility improves as a city acquires more edges, but only on the logarithmic order. Taking the log of (8), we have

$$\log s_i = \log S - (A - B \log k_i - B \langle \log k \rangle) \log \tau - \log \left(\sum_j \tau^{-a_j} \right).$$

The last term is approximated by $\log J - \langle a \rangle \log \tau$[9] so that

$$\log s_i = \log \langle s \rangle + B \log \tau \left(\log k_i - \langle \log k \rangle \right). \tag{11}$$

A couple of observations are in order. The equation above answers two questions concerning the relationship between a network structure and a system of cities. The first one is "Does construction of an edge boost the local economy?" The answer is "Apparently." The second, and more interesting question is "How so?" The answer is twofold.

In terms of a linear scale, (11) can be rewritten as $s_i = \langle s \rangle \left(\frac{k_i}{\gamma} \right)^{B \log \tau}$, where $\gamma := \prod_{i=1}^{J} k_i^{1/J}$ is the geometric mean of the degree. It indicates that city size is anchored around the base city size $\langle s \rangle$ multiplied by the deviation $(k_i/\gamma)^{B \log \tau}$. If a city has a large degree, then its size becomes larger than the standard city size by a factor of $(k_i/\gamma)^{B \log \tau}$ and vice versa for a city with a small degree. The city size coincides with the cornerstone size of $\langle s \rangle$ exactly when its degree matches the national (geometric) average.[10] The deviation is amplified as shipment becomes costly, which, in turn, confirms our observation made in Proposition 3.1.

We also note that adding an edge to a city increases its size, but the change in size is inversely proportional to the current degree provided $B \log \tau < 1$. If city i is highly wired already, then the introduction of a new edge to city j does not add much to city i. The geodesic length to city j is already short before the establishment of the new edge. You can go to many cities in a single step and city j is likely to be linked to at least one of those many neighboring cities already, making the geodesic length to city j just two. The added edge will only reduce the geodesic length by one. On the other hand, if the current degree of city i

[9] Let $\vec{a} := (a_1, a_2, \cdots, a_J)$ and $\langle \vec{a} \rangle := (\langle a \rangle, \langle a \rangle, \cdots, \langle a \rangle)$. The Taylor series expansion about $\vec{a} = \langle \vec{a} \rangle$ tends to

$$\begin{aligned} \log\left(\sum_j \tau^{-a_j}\right) = & \quad \log\left(\sum_j \tau^{-\langle a \rangle}\right) + (\vec{a} - \langle \vec{a} \rangle) \cdot D \log\left(\sum_j \tau^{-a_j}\right)\Big|_{\vec{a}=\langle\vec{a}\rangle} + O\left[(\vec{a} - \langle \vec{a} \rangle) \cdot (\vec{a} - \langle \vec{a} \rangle)\right] \\ \rightarrow & \quad \log J - \langle a \rangle \log \tau, \end{aligned}$$

by the law of large numbers.

[10] This examination begs one question: If my city has the average number of edges, is my city larger or smaller than the national average in size? The answer is "larger". Since transportation cost and the branching factor are both greater than one, $\frac{\log \tau}{\log \kappa}$ is positive. Plus, the geometrical mean is smaller than the arithmetic mean. To score a national average $\langle s \rangle$ you only need γ edges. It should be noted, however, that in a scale-free world, arithmetic mean does not carry much information. The lognormal is the new normal (or any heavy-tailed distribution is for that matter) and the geometric average is the new average in this world as we saw in Figure 5(b).

is low, then the link to city j will not only reduce the geodesic length to city j greatly but also reduce the geodesic lengths to the cities in city j's neighborhood. Consequently, city i will see significant reduction in its average geodesic length.

Based on the degree-size relationship (11) the city-size distribution is given as follows:

PROPOSITION 3.2 CITY-SIZE DISTRIBUTION

Suppose that the economy has a unipartite network Γ with the associated degree distribution $G(k)$. The city-size distribution of this economy follows the distribution function $F(s)$, defined by

$$F(s) = G(k(s)), \tag{12}$$

where $k(s) := \gamma(s/\langle s\rangle)^{\frac{\log\kappa}{\log\tau}}$. Its probability density function (PDF) is

$$f(s) = k'(s)g[k(s)] = \frac{\log\kappa}{\log\tau}k(s)s^{-1}g\left[k(s)\right], \tag{13}$$

where $g(\cdot)$ denotes the PDF of degree k.

Since the transport cost and average branching factor only come into the equation in the form of a quotient of their logarithmic values, $\frac{\log\kappa}{\log\tau}$, we will denote this by δ for estimation purposes, in which case, (13) becomes $f(s) = \gamma\delta\langle s\rangle^{-\delta}s^{\delta-1}g\left[k(s)\right]$. As we have already seen a small δ stretches out the distribution and a large δ does the opposite.

3.8 City-Size Distribution under Different Network Systems

Now that we have the city-size distribution based on the city's degree, we can make our predictions based on different transport network structures. There are two network models of particular interest: ER and BA networks.

Note that empirical determination of the transport network relevant to the formation of a system of cities is a tough job. The task at hand is to find a network that is consistent with the real city-size distribution (and we have already discarded complete and completely isolated networks in Section 3.4). The most consistent network structure will give us a clue as to the shape of a network that is germane to the formation of cities.

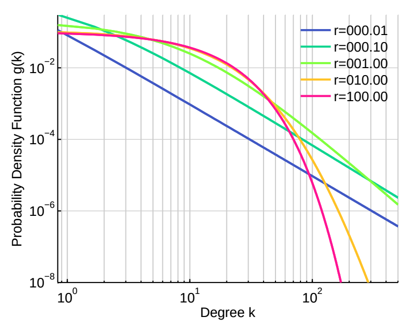

Figure 9. Probability density function of degree with $k_0 = 0$ and $m = 10$.

Jackson and Rogers [JR07] constructed a degree distribution of a directed[11] dynamic network as follows:

$$G(k) = 1 - \left(\frac{k_0 + rm}{k + rm}\right)^{1+r} \quad \text{for} \quad k \geq k_0, \tag{14}$$

[11]Commodities can flow either way on an edge. We take an arrowhead on a directed edge just as a decorative memorabilia indicating from which end the edge was constructed, but nothing more. We represent degree distribution by an in-degree distribution. It is impossible to tell different networks apart with an *out*-degree distribution due to the way a network is constructed in [JR07]. Any network comes with a degenerate out-degree distribution.

where k_0 denotes an in-degree with which an entering node is endowed. This value is shared across all the nodes. The ratio of the number of links formed by an ER-like random connection and a BA-like network-based connection is given by r, and m is the average out-degree of a node. Five PDF's of (14) are depicted in Figure 9 as a visual cue. In the figure parameter r ranges from .01 (over 99% network-based and less than 1% random links) to 100 (the other way around). A predominantly random PDF (with large r) tapers off quickly whereas a mostly network-based PDF (with small r) only gradually dissipates with degree. We expect that our economy operates with a small r. In what follows we refer to in-degree as the degree unless otherwise stated. BA network's degree distribution is (14) with $r = 0$, in which case, (14) turns into a Pareto distribution. ER network calls for $r \to \infty$, in which case (14) is no longer well defined and the degree distribution turns into an exponential distribution.[12]

What is left to do is write the mean branching factor κ in terms of other parameters in (14) before we can fully identify the city-size distribution.[13] The actual mean branching factor cannot be computed until after the network is formed. Hołyst et al [HSF+05] provide a good approximate to κ:

$$\kappa = \sum_{k=1}^{J} k \frac{kg(k)}{\sum_{x=1}^{J} xg(x)} - 1 = \frac{\sum_k (2k-1)G(k)}{\sum_x G(x)} - 1 = \frac{\mu_k^2 + \sigma_k^2}{\mu_k} - 1, \tag{15}$$

where μ_k and σ_k^2 denote the mean and variance of k, respectively. For details, see Appendix A.3.

While [JR07] is microfounded and sufficient to generate a fat-tailed degree distribution, it is not necessarily the only degree distribution which a BA network gives rise to. There is a chance that our economy's transportation network may have come around from a different mechanism than [JR07]. In this regard we experimented with other fat-tail distributions as a candidate degree distribution along with (14). In particular, we tested lognormal and generalized extreme value (GEV) distributions for use as a degree distribution. To our knowledge, these degree distributions are not yet microfounded.

4 Empirical Implementation

Now that the model with an explicit transport system is at the ready, we will pitch it against the actual city-size distributions to identify what class of network governs the city-size distribution. By and large the results are in full support of our initial inkling that a scale-free network explains the city-size distribution but ER or other network structures commonly adopted do not.

All told, we have four sets of data on our plate: Belgium, Metropolitan Area (MA), CBSA and Places.[14] Descriptive statistics for each data set are in Table 1. The Belgian data is included to see if our model's predictive value is subject to

[12] The original ER network [ER59] comes with a Poisson degree distribution rather than an exponential degree distribution. The differences in the distribution arise from the way the network is constructed: [JR07] is dynamic, whereas [ER59] is static.

[13] The branching factor is not a free parameter and it cannot be directly estimated from the data, because the estimation algorithm will either explode or create indeterminacy. It is dependent on the shape of the network, which, in turn, is characterized by the other parameters via (15).

[14] The Belgian data is provided courtesy of Soo [Soo05] and the remainder are from US Census 2000. For definitions of MA and CBSA, see http://www.census.gov/population/metro/about/ and for Places, see http://www.census.gov/geo/reference/gtc/gtc_place.html. We thank Jan Eeckhout for sharing his data used in [Eec04].

Data	Belgium	MA	CBSA	Places
Data size J	69	276	922	25,358
Total urban population S	4,344,222	225,981,679	261,534,991	208,735,266
Population covered	42.38%	80.30%	92.93%	74.17%
Largest city	Antwerp	New York CMSA	New York MSA	New York city
Largest size	446,525	21,199,865	18,323,002	8,008,278
City near arithmetic mean	Genk	Oklahoma, OK MSA	Green Bay, WI MSA	Hillsboro city, TX
Arithmetic mean	62,960	818,774	283,661	8,232
Median city	Beringen	Anchorage, AK MSA	Hinesville-Fort Stewart, GA MSA	Harristown village, IL
Median size	39,261	259,600	71,800	1,338
Smallest city	Arlon	Enid, OK MSA	Andrews, TX µSA	New Amsterdam, IN
Smallest size	24,791	57,813	13,004	1
Standard deviation	61,240	1,968,621	974,190	68,390
Skewness	4.183	6.682	10.98	75.53
City near geometric mean	Mouscron	Huntsville, AL MSA	Sunbury, PA µSA	Sutton city, NE
Geometric mean	50,809	342,844	94,373	1,447
Mean of $\log(s)$	10.84	12.75	11.46	7.278
Standard deviation of $\log(s)$	.5697	1.119	1.191	1.754
Skewness of $\log(s)$	1.498	1.048	1.187	.2091

Table 1. Descriptive Statistics. The statistics above the line (shaded in blue) are related to a linear scale and the below (shaded in green) are related to a log scale. Mean of $\log(s)$ is same as the log of geometric mean.

both the area and population size of a country under study. (It was not.) MA and CBSA are a popular choice in the literature. The smallest unit of measurement is a county and they suffer from data truncation ([Eec04]). Places have the finest unit of measurement and are free of truncation. We tested the following five distributions against them: ER/BA, BA, lognormal, GEV and the degenerate distribution. The first two distributions are estimated in three ways: maximum spacing estimation (MSE), minimum Kolomogorov-Smirnov estimation (minKS) and maximum likelihood estimation (MLE), and the remainder in MSE.

In what follows a hat on parameter x indicates its estimate, $\hat{x}$.

4.1 Estimation Methods Employed

The first choice is to go for MLE, which does not work with (14). The likelihood function is monotone increasing in k_0. As a workaround to MLE, we calculated the estimates by MSE. While its use is limited in the city-size literature so far especially when compared to MLE, it is more robust and easier to handle than MLE. The problem we have with MLE is exactly the one exemplified in Ranneby [Ran84] and we used his solution. The MSE estimator maximizes the geometric mean of the gap or step between two adjacent CDF's

$$F(s_i;\theta)-F(s_{i-1};\theta),$$

where θ is a vector of parameters to be estimated and data sequence s is rearranged in the ascending order $s_1 \le s_2 \le \cdots \le s_J$.[15] The idea here is to split the interval $[0,1]$,

[15] The first and last gap are defined by $F(s_1;\theta)-F(-\infty;\theta)$ and $F(\infty;\theta)-F(s_J;\theta)$ each.

the range of a CDF, in J steps in the way that none of the assigned $F(s_i;\theta)$ will create a disruptively large gap with its neighbors and the gaps should be evenly spaced as much as possible on the *logarithmic* scale. Maximizing the *arithmetic* mean does not work here because it will always be $1/J$ no matter what estimates we toss in. This actually works as a ap on our geometric mean in turn, by Jensen's inequality. Thus, we can safely rule out the possibility that the maximand tends to infinity, which is exactly the reason why we had to discard MLE. For more on MSE, see Appendix A.4.

4.2 A Scale-Free Transportation Network Explains the City-Size Distribution

Estimation with four different data sets unanimously chooses BA over ER as the underlying transport network in our economy. We report our results in Table 2 and Figures 10 to 13 along with other distributions.

ER/BA in the table corresponds to (14). We left the estimated distribution functions for ER/BA in Figures 14 to 17. As low values of $\hat{r}$ indicate, edges are formed predominantly through networking rather than by random selection. We cross-checked estimates with minKS and MLE[16] and we obtained a similar result. To be doubly sure of our findings, we ran estimation with $r \to \infty$. ER in Table 2 lists the statistics with $r \to \infty$. The statistics of ER seem to be comparable with other distributions except that the estimated transportation cost is unreasonably high. A one-dollar pen will cost more than the US GDP five towns over on the purely random network.[17] Thus, we conclude that a scale-free transportation network explains the city-size distribution but a scale-variant network does not.

Estimated $\hat{\delta}$ ranges from .9911 to 2.536.[18] As we discussed in reference to (11) we confirm that in most cases, the impact of adding an edge on city size wears off as degree itself becomes saturated (it cannot exceed $J-1$), or put differently, New York has more edges, size for size, than any other cities as it takes more edges to raise the city size as the city grows further.

We ran MSE with three other distributions representative of the existing city-size models to compare with our model. Eeckhout [Eec04]'s model leads to a lognormal distribution and Berliant and Watanabe [BW14] predict a GEV distribution as the city-size distribution. A complete graph will result in a degenerate probability distribution. The BA economy fits comfortably into the circle of existing testable models based on all the statistics we computed in Table 2 (usually coming in second on all fronts except for Places).

In addition we put two other fat-tailed degree distributions to the test. The results (the last two rows in Table 2) seem to indicate that the network formation does not necessarily have to be of [JR07] type. Regardless of how it came about, a network with a fat-tailed degree distribution results in the city-size distribution that closely resembles the actual distribution.

[16] With k_0 fixed at zero to prevent explosion. MSE and minKS point $\hat{k}_0$ towards zero.

[17] There is not enough variance in the ER degree distribution, certainly not power-law type behavior. To generate the empirical city-size distribution, the ER economy has to amplify and capitalize on what little variance its degree distribution has to offer (cf. Proposition 3.1). As a result τ has to be ludicrously large to make things work. On the other hand, if the transportation infrastructure is in its early stage of development without any hubs, then the country's transportation cost will probably be higher than more BA-like countries because Zipf's law is a universally observed phenomenon. There is a trade-off between τ and how close the transport network is to BA, provided that Zipf's law holds at all times.

[18] The estimate tends to decrease as data size J increases.

Data	Distribution	⟨log LH⟩ ▲	KS ▼	⟨log step⟩ ▲	geo/arith ▲	$\lvert\theta\rvert$	BIC ▼	AIC ▼	r
Belgium	Lognormal (Eeckhout)	-11.69	.1986	-5.266	.005166/.01449	2	1621	1617	
Belgium	GEV (Berliant & Watanabe)	**-11.40**	**.1122**	**-4.981**	**.006870**/.01449	5	**1594**	**1583**	
Belgium	Complete Graph (de facto)	$-\infty$	.6812	$-\infty$	0/.01449	1	∞	∞	
Belgium	ER/BA (Jackson & Rogers)	**-11.47**	**.1348**	**-5.072**	**.006268**/.01449	5	1604	**1593**	.002745
Belgium	ER (Jackson & Rogers)	-11.49	.1766	-5.086	.006185/.01449	4	**1603**	1594	∞
MA	Lognormal (Eeckhout)	-14.28	.1036	-6.232	.001996/.003623	2	7891	7884	
MA	GEV (Berliant & Watanabe)	**-14.13**	**.04334**	**-6.089**	**.002267**/.003623	5	**7828**	**7810**	
MA	Complete Graph (de facto)	$-\infty$	.7935	$-\infty$	0/.003623	1	∞	∞	
MA	ER/BA (Jackson & Rogers)	**-14.17**	**.06102**	**-6.134**	**.002168**/.003623	5	**7852**	**7834**	.001154
MA	ER (Jackson & Rogers)	-14.21	.1057	-6.173	.002084/.003623	4	7860	7851	∞
CBSA	Lognormal (Eeckhout)	-13.05	.09402	-7.548	.0005270/.001085	2	2.407e+04	2.4063+04	
CBSA	GEV (Berliant & Watanabe)	**-12.91**	**.02606**	**-7.409**	**.0006056**/.001085	5	**2.384e+04**	**2.382e+04**	
CBSA	Complete Graph (de facto)	$-\infty$	.8362	$-\infty$	0/.001085	1	∞	∞	
CBSA	ER/BA (Jackson & Rogers)	**-12.95**	**.05922**	**-7.449**	**.0005819**/.001085	5	**2.391e+04**	**2.389e+04**	.0004526
CBSA	ER (Jackson & Rogers)	-13.29	.1762	-7.794	.0004121/.001085	4	2.454e+04	2.452e+04	∞
Places	Lognormal (Eeckhout)	**-9.258**	**.01895**	**-8.840**	0/3.944e-05	2	**4.696e+05**	**4.696e+05**	
Places	GEV (Berliant & Watanabe)	**-9.254**	**.008847**	**-8.836**	0/3.944e-05	5	**4.694e+05**	**4.693e+05**	
Places	Complete Graph (de facto)	$-\infty$	.8342	$-\infty$	0/3.944e-05	1	∞	∞	
Places	ER/BA (Jackson & Rogers)	-9.268	.02198	-8.849	0/3.944e-05	5	4.701e+05	4.700e+05	.0003171
Places	ER (Jackson & Rogers)	-9.392	.1134	-8.974	0/3.944e-05	4	4.764e+05	4.763e+05	∞
Places	Lognormal as Degree Dist.	-9.258	.01896	-8.840	0/3.944e-05	4	4.696e+05	4.696e+05	
Places	GEV as Degree Dist.	-9.255	.01159	-8.836	0/3.944e-05	5	4.694e+05	4.694e+05	

Table 2. Model Comparison

Row color corresponds to the line colors in Figures 10 to 13. ▲ denotes a statistic the higher value of which indicates a better fit and ▼, the other way around. ⟨log LH⟩ denotes the average of the log of likelihood values, KS denotes the Kolomogorov-Smirnov statistic, ⟨log step⟩ measures the geometric mean of the step $F(s_i;\theta)-F(s_{i-1};\theta)$ in logarithms. Geo/arith measures the ratio between geometric mean and arithmetic mean of the step. The closer the geometric mean is to the arithmetic mean, the better the fit is. It is zero for Places due to multiple cities having the same size. $|\theta|$ counts the number of parameters. BIC and AIC stand for Bayesian and Akaike Information Criteria for detecting overfitting. **Boldface with white foreground** marks the winner and **boldface with black foreground** denotes the runner-up among the five distributions tested.

5 Conclusion and Extensions

We examined how the network of cities affects the city-size distribution. We built a simple economic model with an explicit transport network. The bridge between network structure and city size is represented in (11), where we learned that *there is a log-linear relationship between city size and city degree.*

We put two commonly studied networks to the test. The classical ER random graph is too egalitarian to generate gravitationally large cities like New York City. The BA model explains the city-size distribution better than the ER model and bears very close comparison with other proposed city-size models in existence. The BA network has a scale-free degree distribution and the resulting city-size distribution behaves similarly via (11). In fact, it is would be odd if the city-size distribution was *not* scale free under a BA network. Large nodes with a high degree like Chicago attract a large mass of people because A) goods produced in Chicago are in high demand for its inexpensive delivered price owing to its high degree and B) goods available for consumption in Chicago are also inexpensive thanks to its high degree. The exact opposite applies to small cities. But there are still some people knowingly living in small cities because we cannot afford to wipe them off the map due to preference for variety. This gives rise to a few cities of an overwhelming size and a myriad of small cities. The actual city-size distributions (we tried Belgium and the United States in particular) unanimously opt for a BA network.

From this point on, it would be reasonable to combine GEV to determine firm productivity as in [BW14] and BA for transportation network structure by way of simulations, but we will not have an analytical solution due to the added complexity.

We argued that network structures motivate the population to form a specific distribution of city sizes. The structure of the network is pre-selected. Considering the fact that it is easier to relocate people than to build transport infrastructure, this is not an unreasonable assumption in the short run. New York City would have been much smaller had it not been the entrepôt to Europe. However, the degree-city relationship is not a one-way street. It may be the other way around: the relocation of people forces the transportation network to follow a specific pattern. It can also be the case that the network structure and its associated city-size distribution are in fact a product of some common underlying causes. The United States has seen a number of drastic changes in its network structure. Tracing the historical co-development of the network structure with the city-size distribution may reveal a clue to identifying the direction of causality. A problem with this methodology is that transportation networks are not unique, in that there are generally multiple modes of transport and multiple companies providing services in each mode.

For now, as a preview, consider a commodity transportation firm that is installing a transportation network that maximizes its profit by choosing r. It uses the mechanism described by [JR07] to add links to its network, namely random links and friends of friends, where r determines the relative proportions. As noted in the introduction, the difference between social networks and transportation networks is who makes the decisions about links, the nodes themselves or the owner of the network. If the revenues and cost of the network are additive across nodes, then the profit from a network is additive across nodes, so there is no distinction between maximizing the objective of a node and maximizing profit or utility of the entire network. In other words, profit of the network owner corresponds to efficiency of the network in [JR07]. Suppose that shipping industry is competitive and the shipping firm's indirect revenue function is additively separable across nodes and convex in node degree, and also assume that its cost is additive across nodes and proportional to node degree. Then we can follow the framework proposed in section IV of [JR07], in particular Corollary 1, to find r that maximizes its profit, independent of the city-size distribution. This allows us to take the network development process as exogenous, and leading to the BA network $r = 0$.

We finish our discussion with one last remark. It has been suggested that other networks be implemented in our framework, for example the optimal transport network for a given population distribution (assuming a cost function). This would require the geodesic length or degree distribution for the optimal network. We are not aware of any results addressing this issue.

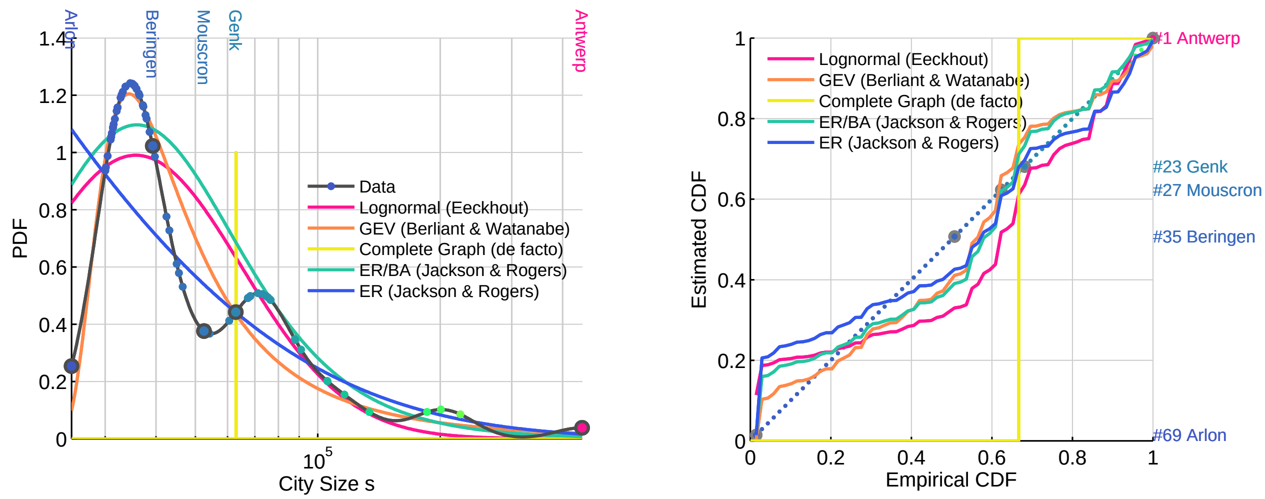

(a) PDF

(b) PP Plot

Figure 10. Model Comparison (Belgium)

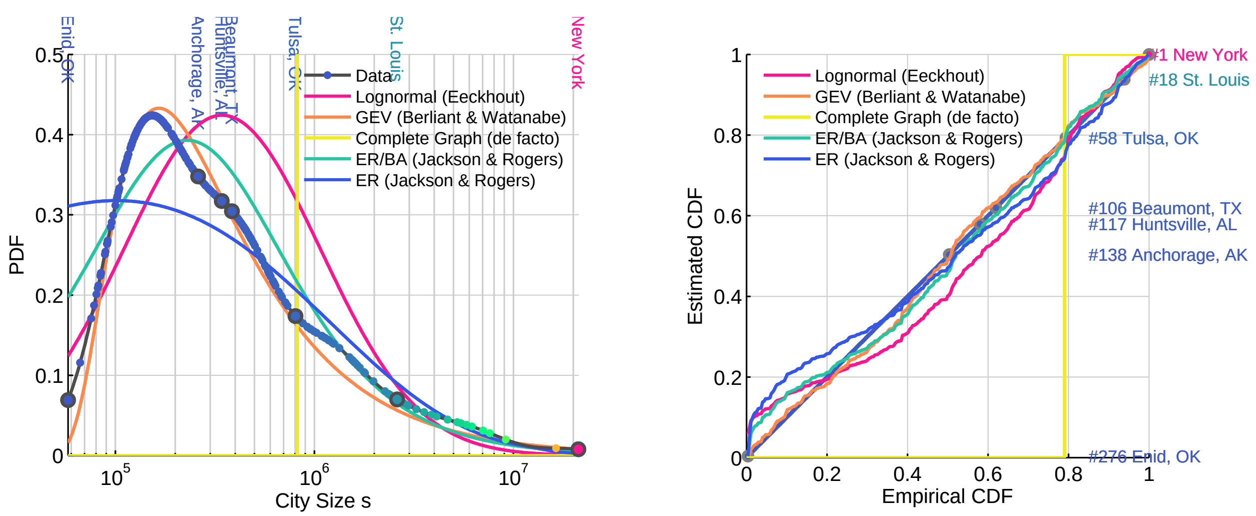

(a) PDF

(b) PP Plot

Figure 11. Model Comparison (MA)

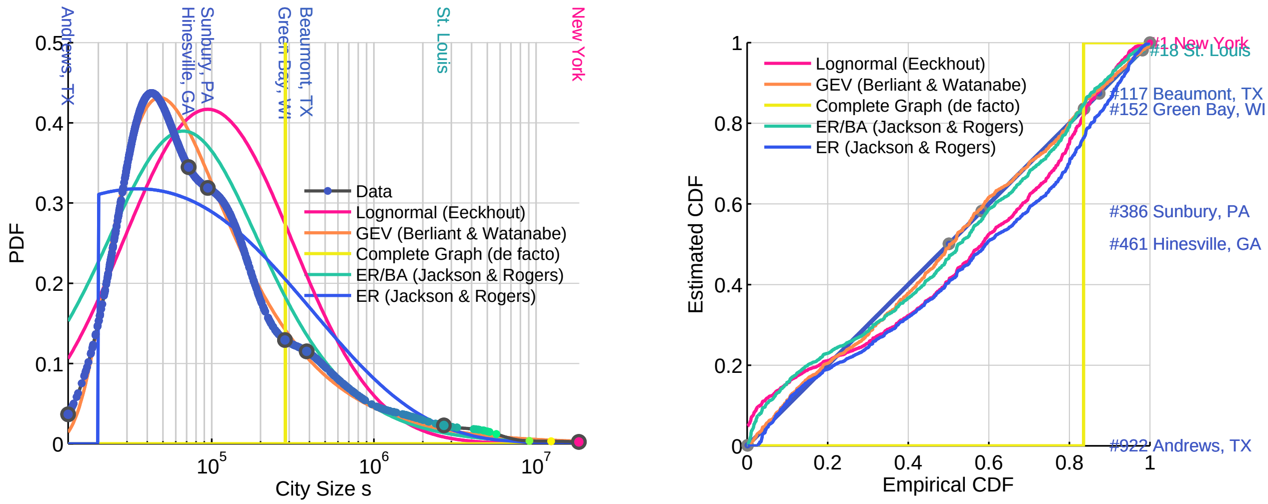

(a) PDF (b) PP Plot

Figure 12. Model Comparison (CBSA)

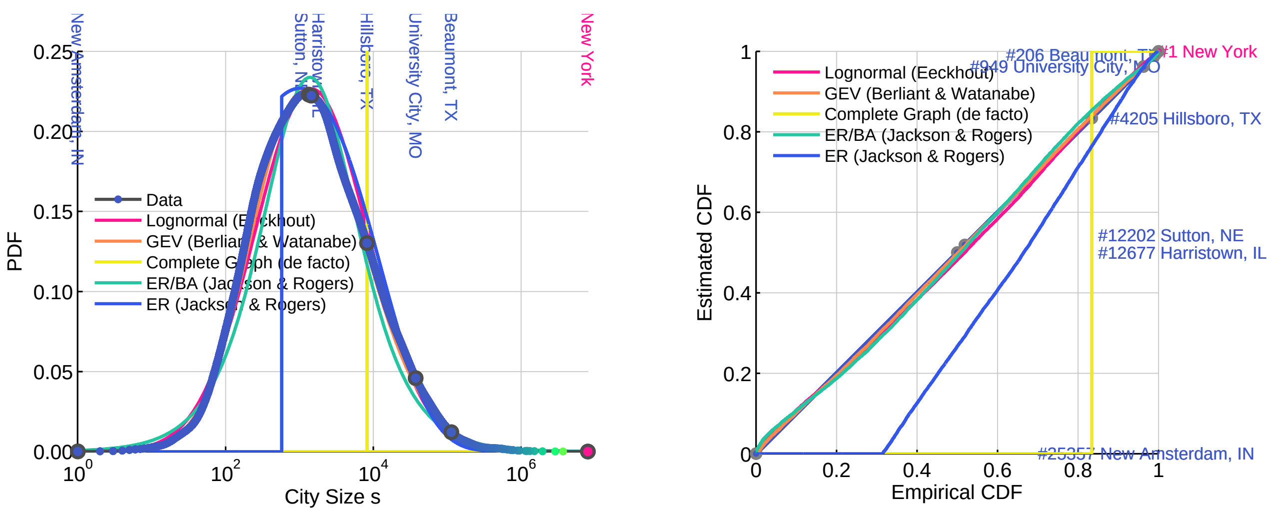

(a) PDF (b) PP Plot

Figure 13. Model Comparison (Places)

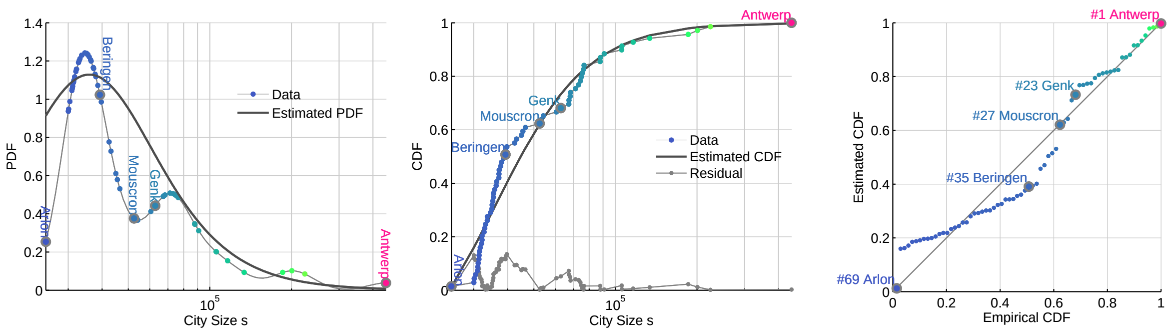

(a) PDF (b) CDF (c) PP Plot

Figure 14. ER/BA for Belgium, 2000 [S0005].

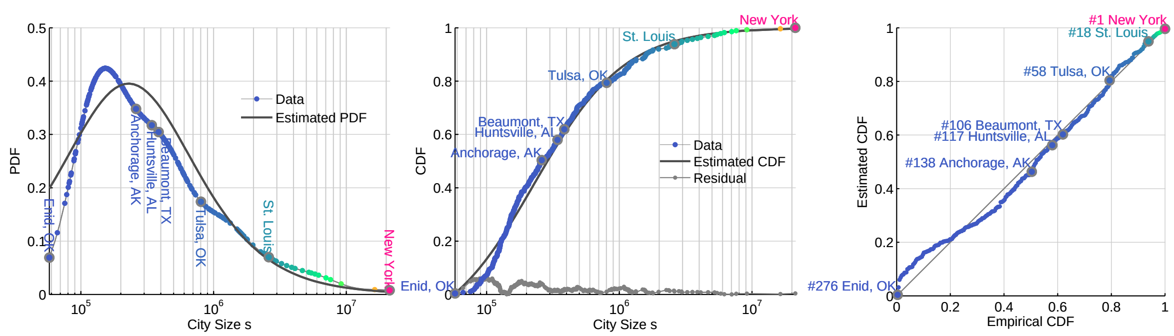

(a) PDF (b) CDF (c) PP Plot

Figure 15. ER/BA for MA, Census 2000.

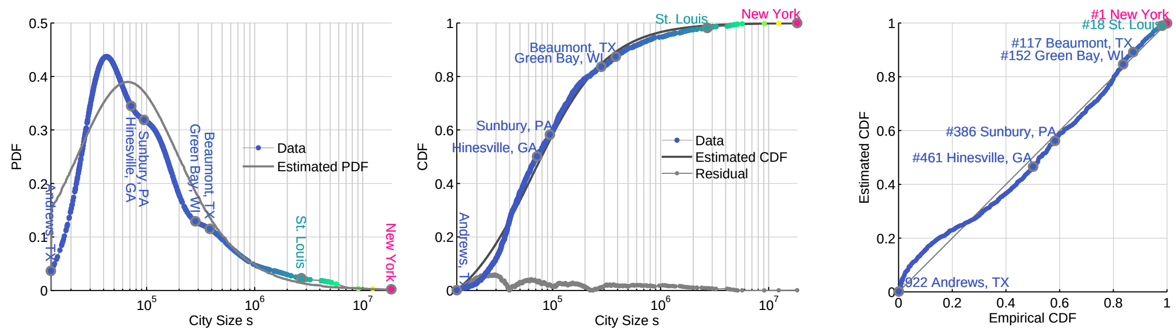

(a) PDF (b) CDF (c) PP Plot

Figure 16. ER/BA for CBSA, Census 2000.

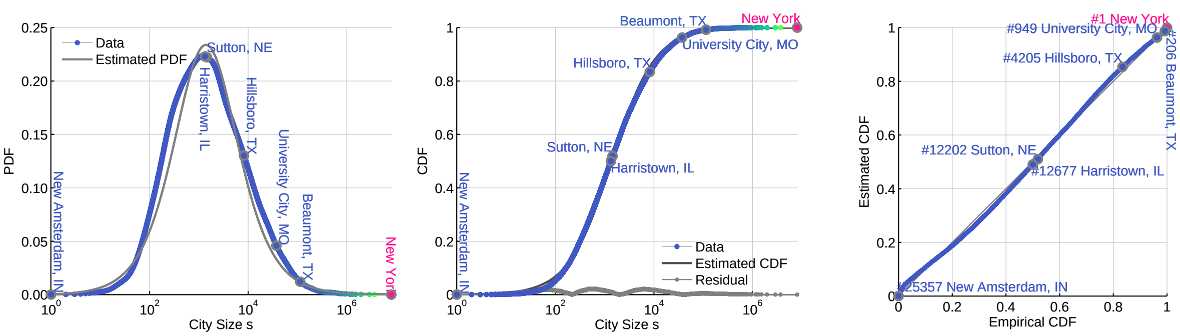

(a) PDF (b) CDF (c) PP Plot

Figure 17. ER/BA for Places, Census 2000.

A Appendix

A.1 Proof of Proposition 3.1

Proof. Note that $s(a_i)$ is monotone decreasing in a_i. Suppose $J > 2$ and the network is neither complete or completely isolated. We have

$$s'(a_i) := \frac{ds(a_i)}{da_i} = -(\log\tau)s(a_i)S^{-1}(S - s_i) \leq 0$$

with equality iff $\tau = 1$. The second derivative is, therefore,

$$\frac{d^2s(a_i)}{da_i^2} = [s'(a_i)]^2\frac{S - 2s_i}{s_i(S - s_i)} \geq 0,$$

with equality iff $\tau = 1$. Hence $s(a_i)$ is strictly convex in a_i.

To show that $s(a_i)$ bulges as τ grows, first note $\frac{\partial s(a_i)}{\partial\tau} = -\tau^{-1}s(a_i)(a_i - AB^{-1})$, where $A := \sum_j a_j\tau^{-a_j}$ and $B := \sum_j \tau^{-a_j}$. Then

$$\frac{dD(\tau)}{d\tau} = \frac{1}{2\tau}\left\{[s(a_M) - s(a_H)](a_H - AB^{-1}) + [s(a_M) - s(a_L)](a_L - AB^{-1})\right\},$$

where $a_M := (a_H + a_L)/2$. The first term in the curly braces is positive because $s(a_M) - s(a_H) > 0$ and $a_H - AB^{-1} = B^{-1}\sum_{j\neq H}(a_H - a_j)\tau^{-a_j} > 0$. Likewise, the second term is positive because $s(a_M) - s(a_L) < 0$ and $a_L - AB^{-1} < 0$. Therefore $\frac{dD(\tau)}{d\tau} > 0$, which establishes the claim. ◻

A.2 Idea behind Geodesic Length (9)

We briefly repeat [HSF+05]'s arguments to obtain (9) in our context. Consider a geodesic between nodes v_i and v_j. We ignore loops. The probability that a child node traces back to its ancestors via some circumvention is proportional to $1/J$. It becomes negligible as the system size J grows (our system size ranges from 69 to 25,358 in Section 4). As shown in [HSF+05], the resulting error is minimal. A tree is a sequence of nodes where each node except for the root node has exactly one parent (or ancestor) node. Each node may or may not be followed by (a) child node(s). There are no cycles on a tree. If we pick a random tree starting from v_i, we will wind up at v_j somewhere along the tree $k_j/\sum_{r\in V}k_r$ of the time and we will not reach v_i the remaining $1 - k_j/\sum_r k_r$ of the time. On average, we will reach v_j within $\sum_r k_r/k_j$ trials. Suppose that the depth (the number of parent nodes that you have to go through before reaching your root node) of v_j is l. There are $k_i\kappa^{l-1}$ nodes whose depth is l. Therefore, on average, we arrive at v_j in l steps if

$$\frac{\sum_r k_r}{k_j} = k_i\kappa^{l-1}, \tag{16}$$

from which we obtain (9). In other words, if, on average, it takes more than $k_i\kappa^{l-1}$ trials to reach city j, i.e., $\frac{\sum_r k_r}{k_j} > k_i\kappa^{l-1}$, then it is likely that city j is more than l steps away from your city i. You would try $k_i\kappa^{l-1}$ times to find city j, when in fact you would need additional $\frac{\sum_r k_r}{k_j} - k_i\kappa^{l-1}$ trials to reach city j, meaning that city j is not in the group of cities l steps away from you but actually located somewhere farther down. On the contrary if it takes less than $k_i \kappa^{l-1}$ trials to reach city j, then city j should be less than l steps away from you. You would not need that many trials to find a city j, the implication being that, once again, you are looking at a wrong group of cities. Thus, city i and j are l steps apart from each other exactly when (16) is satisfied with equality.

A.3 Branching Factor

Take a random edge and walk towards one arbitrarily selected end. Call where you arrived at a neighboring node. The average degree of neighboring nodes thus reached approximates the mean branching factor κ. In effect, we will take one degree off the average degree found above because the edge we just walked on cannot be used to reach the destination city. We are climbing up a tree, not down (recall how goods find their destination city in Section 3.6). Also note that the mean branching factor is not just a mean degree $\langle k \rangle$. We are not hopping from one city to another but climbing a tree from one neighbor to next to reach the destination city. Thus, a city charged with lots of links is more likely to be a neighbor of some city than a poorly connected city, and cities are duly weighted when fed into the mean branching factor. In other words, Houston is rare while there are quite a few mid-sized cities but that does not mean Houston is hard to reach at random for its rarity. Houston has far more edges than mid-sized cities and we are likely to travel through Houston at some point or another (cf. Figure 4(a)). In particular a node of degree k has a chance proportional to $kg(k)$ of being at one end of an arbitrary direction on a randomly chosen edge, where $g(k)$ is a probability density function of (14). Or put differently, if we parachute into a random edge and then flip a coin to decide which direction to go in, we will arrive at a k-th degree city $kg(k)$ out of $\sum_{x=1}^{J} xg(x)$ times. Thus, the mean branching factor is given by (15).

A.4 Maximum Spacing Estimation

It might be easier to make sense of the use of geometric mean in MSE if we recast it as an analogue of a more familiar, linear regression. The geometric mean of steps here corresponds to ordinary least squares and the arithmetic mean corresponds to a plain sum of residuals. Say we are trying to regress $y = (-1, 0, 1)$ on $x = (-1, 0, 1)$. If we aim to minimize the sum of residuals, *any* real estimate that makes the regression line run through the origin $(0, 0)$ will work, just as much as any estimate will make the arithmetic mean of gaps $1/J$. We will end up with infinitely many estimates because residual at $x = 1$ always offsets the one at $x = -1$. To ward off this cancellation problem, we usually try to minimize the sum of *squared* residuals, which leads to a unique estimate, a 45-degree line. Similarly, the use of *geometric* mean will solve the indeterminacy problem that comes with arithmetic mean and will promise us sensible estimates.

The geometric mean also comes in handy here. The gap tends to get tighter near the top and/or the bottom of most distributions as the CDF creeps up to one and/or bears down on zero. However, this does not mean New York or New Amsterdam, IN counts less than other cities as a sample. The geometric mean offsets this general tendency and duly stretches small gaps so that these extremities will receive no less attention than the ones in the middle. There is no particular reason to let the mid-sized cities punch above their weight.

On a related matter, we report Kolomogorov-Smirnov (KS) statistic. MSE is similar to KS in that both KS and the maximand of MSE are a power mean. KS statistic is a power mean of the form

$$\left\{ \frac{1}{J} \sum_i |\text{Empirical } F(s_i) - F(s_i)|^p \right\}^{\frac{1}{p}} \tag{17}$$

with $p \to \infty$ (i.e., the maximum of the residuals, the L^∞ norm), whereas the maximand of MSE is a power mean of the form

$$\left\{ \frac{1}{J} \sum_i (F(s_i) - F(s_{i-1}))^p \right\}^{\frac{1}{p}} \tag{18}$$

with $p \to 0$ (i.e., the geometric mean of the gaps). The way they aggregate the data is where their difference comes in. KS statistic only picks up a single city where the predicted value deviates from the actual value the most. It does not tell us anything about the selected model's performance over the remainder of cities other than the fact that their gap is tighter than the KS value (but *not* by how far). On the other hand, the maximand of MSE is determined by the step gap log-averaged over the entire range of the cities, and probably a better measuring tool to gauge the model's performance in that respect.

To get a sense of what MSE hunts for, consider what happens if we pull out the estimate that *minimizes* the geometric mean instead. Minimum spacing estimator would dump the entire interval $[0, 1]$ on one particular city i (any city will do) so that $F(s_j; \theta) = 0$ for all $j < i$ and $F(s_j; \theta) = 1$ for all $j \geq i$, in which case, the geometric mean would be zero, the smallest value possible (practically the same result when you try to maximize the arithmetic mean as we mentioned above, in the sense that *any* estimate will be as good as any other). This would make such a pointless estimator. MSE does the exact opposite.

References

[AB02] Réka Albert and Albert-László Barabási. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 2002.

[BA99] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *Science*, 1999.

[BB03] Albert-László Barabási and Eric Bonabeau. Scale-free networks. *Scientific American*, 2003.

[BH03] Duncan Black and Vernon Henderson. Urban evolution in the USA. *Journal of Economic Geography*, 2003.

[BMMS13] Kristian Behrens, Giordano Mion, Yasusada Murata, and Jens Südekum. Spatial frictions. *Working Paper*, 2013.

[BW14] Marcus Berliant and Hiroki Watanabe. Explaining the size distribution of cities: X-treme economies. *Quantitative Economics*, 2014.

[CFIK10] Nicholas A. Christakis, James H. Fowler, Guido W. Imbens, and Karthik Kalyanaraman. An empirical model for strategic network formation. *mimeo*, 2010.

[DMO06] S. N. Dorogovtsev, J. F. F. Mendes, and J. G. Oliveira. Degree-dependent intervertex separation in complex networks. *Phyisical Review E*, 2006.

[Dur06] Gilles Duranton. Some foundations for zipf's law: Product proliferation and local spillovers. *Regional Science & Urban Economics*, 2006.

[Eec04] Jan Eeckhout. Gibrat's law for (all) cities. *American Economic Review*, 2004.

[ER59] Paul Erdős and Alfréd Rényi. On random graphs. *Publicationes Mathematicae*, 1959.

[FFH04] Agata Fronczak, Piotr Fronczak, and Janusz A. Hołyst. Average path length in random networks. *Physical Review E*, 2004.

[FKV99] Masahisa Fujita, Paul Krugman, and Anthony J. Venables. *The Spatial Economy*. The MIT Press, 1999.

[FT02] Masahisa Fujita and Jacques-François Thisse. *Economics of Agglomeration*. Cambridge University Press, 2002.

[HSF+05] Janusz A. Hołyst, Julian Sienkiewicz, Agata Fronczak, Piotr Fronczak, and Krzysztof Suchecki. Universal scaling of distances in complex networks. *Physical Review E*, 2005.

[JR07] Matthew O. Jackson and Brian W. Rogers. Meeting strangers and friends of friends: How random are social networks? *The American Economic Review*, 2007.

[LSA01] Eckhard Limpert, Werner A. Stahel, and Markus Abbt. Log-normal distributions across the sciences: Keys and clues. *BioScience*, 2001.

[McC05] Philip McCann. Transport costs and new economic geography. *Journal of Economic Geography*, 2005.

[Mel11] Angelo Mele. A structural model of segregation in social networks. *mimeo*, 2011.

[NMW00] M. E. J. Newman, C. Moore, and D. J. Watts. Mean-field solution of the small-world network model. *Physical Review Letters*, 2000.

[NSW01] M. E. J. Newman, S. H. Strogatz, and D. J. Watts. Random graphs with arbitrary degree distributions and their applications. *Physical Review E*, 2001.

[NW99] M. E. J. Newman and D. J. Watts. Renormalization group analysis of the small-world network model. *Physics Letters A*, 1999.

[Ran84] Bo Ranneby. The maximum spacing method. an estimation method related to the maximum likelihood method. *Scandinavian Journal of Statistics*, 1984.

[RHW07] Esteban Rossi-Hansberg and Mark L. J. Wright. Urban structure and growth. *The Review of Economic Studies*, 2007.

[Soo05] Kwok Tong Soo. Zipf's law for cities: a cross-country investigation. *Regional Science and Urban Economics*, 2005.

[ZLG+09] Zhongzhi Zhang, Yuan Lin, Shuyang Gao, Shuigeng Zhou, and Jihong Guan. Average distance in a hierarchical scale-free network: an exact solution. *Journal of Statistical Mechanics: Theory and Experiment*, 2009.