

Explaining the Size Distribution of Cities: X-treme Economies*

Marcus Berliant[†] and Hiroki Watanabe[‡]

March 2008

Abstract

The methodology used by theories to explain the size distribution of cities takes an empirical fact and works backward to first obtain a reduced form of a model, then pushes this reduced form back to assumptions on primitives. The induced assumptions on consumer behavior, particularly about their inability to insure against the city-level productivity shocks in the model, are untenable. With either self insurance or insurance markets, and either an arbitrarily small cost of moving or the assumption that consumers do not perfectly observe the shocks to firms' technologies, the agents will never move. Even without these frictions, our analysis yields another equilibrium with insurance where consumers never move. Thus, insurance is a substitute for movement. Even aggregate shocks are insufficient to generate consumer movement, since consumers can borrow and save. We propose an alternative class of models, involving extreme risk against which consumers will not insure. Instead, they will move. JEL number: R12 Keywords: Zipf's Law, Gibrat's Law, Size Distribution of Cities, Extreme Value Theory

*The first author thanks the Kyoto Institute of Economic Research at Kyoto University for their kind hospitality during the writing of the first draft of this paper. The second author thanks the Center for Research in Economics and Strategy (CRES) at the Olin School of Business of Washington University in St. Louis for financial support. Both authors thank Matthias Wrede and the participants at the North American Regional Science Association International meetings in Savannah for helpful comments and especially Jan Eeckhout for incisive comments and for sharing his data with us, but retain responsibility for any errors.

[†]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, Washington University, Campus Box 1208, 1 Brookings Drive, St. Louis, MO 63130-4899 USA. e-mail: watanabe@wustl.edu

1 Introduction

1.1 Motivation

A small industry has developed that seeks to provide a theory to explain a singular but robust stylized fact in urban growth: the size distribution of cities. Zipf's law or the rank-size rule, as applied to the size distribution of cities, states that for any country, the rank of a city according to population (for example, New York is ranked number one in the US) multiplied by its population is constant. Thus, Los Angeles has half the population of New York, whereas Chicago has one third the population of New York. This stylized fact or factoid holds across many countries and time periods, but it is only one fact. In general, it is connected to Gibrat's law, stating that stochastic proportional growth tends to a lognormal distribution. The most compelling empirical work in this area shows that the size distribution of cities is lognormal (Eeckhout, 2004) when the data is not cut off at an arbitrary rank or population.¹ For those unfamiliar with the empirics associated with this literature, we display in Figure 1 a graph of Eeckhout's data, consisting of more than 25,000 places from U.S. Census 2000. Since population on the horizontal axis and rank on the vertical axis are both plotted in log scales, the rank-size rule, taken literally, would say that the plot should be linear with slope -1 . Deviations from the rule or law at the top and bottom of the size distribution are documented and discussed in the literature. See Gabaix and Ioannides (2004) for a fine survey of the entire area of research.

¹It is well known that the lognormal distribution applies not only to the size distribution of cities, but to many phenomena in the physical sciences as well; see the survey by Limpert *et al* (2001) with applications as broad as the length of sentences in the writings of George Bernard Shaw to geology and plant physiology. Of particular interest is the literature showing that rainfall (see Meneghini *et al*, 2001), soil moisture (see Janowicz *et al*, 2003), and crop abundance (see Halloy, 1999) are distributed according to lognormal distributions. In fact, Halloy (1999) displays graphs (see for example his figure 4) for crop abundance in New Zealand from 1842 to 1990 that look very much like the graphs used to explain the size distribution of cities. He also explains dynamics as in Duranton (2007).

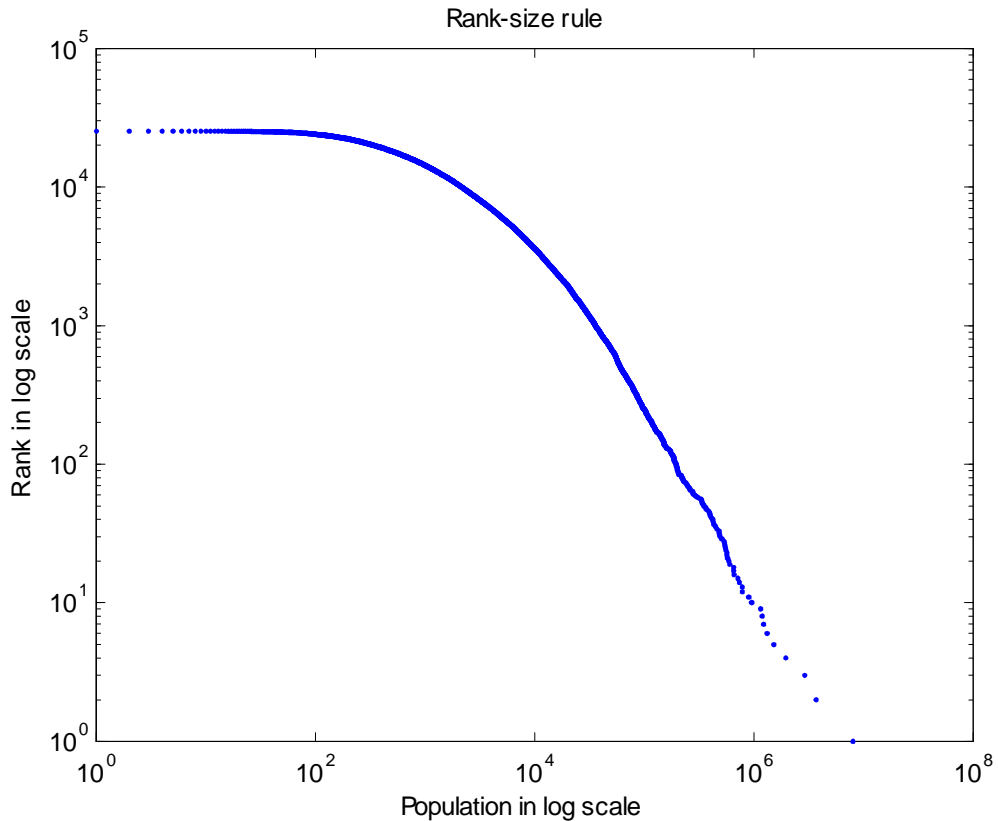


Figure 1: The rank-size rule. *Data Source:* Census 2000.

Explanation of this stylized fact or factoid by a theory has long been an objective of urban economists; it is quite robust, but also very difficult to theorize about. Three recent articles, Eeckhout (2004), Duranton (2007), and Rossi-Hansberg and Wright (2007), have made valiant attempts to tackle this issue head on.

The models seeking to explain this factoid are developed working backwards from empirics. That is, they take the empirical fact and attempt first to explain it using a reduced form, generally stochastic in nature, and then push this reduced form back to assumptions on primitives. As we shall see, these assumptions on primitives generally do not look natural, in the sense that if one were formulating a model of cities from scratch, it would not be obvious that one would want to begin with these assumptions. In fact, the literature of urban economics prior to the introduction of these models did not. Moreover, since these models are constructed for a single purpose, namely to explain an empirical factoid, they seem incapable of explaining other empirical regularities, though they seem to be judged exclusively on the basis of how well they explain the one empirical factoid. Finally, it is worthwhile to note

that there is likely an infinite number of models capable of explaining the size distribution of cities, and some of these might not even be stochastic; see, for example, Fujita and Mori (1997).

We shall focus primarily on the behavior of consumers, in particular the degree to which they can hedge against risk. Our findings are as follows.

First, the type of risk featured in the literature is city-level risk in the form of a shock each period to the city-wide production function. The shock is known to all before they make their decisions, in particular consumer decisions about location and consumption bundle. There is no aggregate risk, since the number of cities are large and shocks are i.i.d., and this is often stated explicitly in the papers. In this context, *insurance is a perfect substitute for consumer movement*. We consider either self insurance, where a consumer insures over time by borrowing or saving to smooth consumption, or insurance markets, where a consumer insures using the fact that shocks are independent over space at a given time. With either type of insurance (or a combination), we find an equilibrium that yields the same period by period utility for each consumer as the one presented in the literature, where consumers move and generate Zipf's law or Gibrat's law. Our equilibrium features no consumer movement. Moreover, with even arbitrarily small moving costs or arbitrarily small uncertainty about shocks on the part of consumers, only our equilibrium survives. The existing literature finds that initial conditions don't matter, in that the size distribution of cities eventually tends toward lognormal. For our equilibrium, initial conditions matter in that consumers never move.

Second, even aggregate risk is insufficient to generate consumer movement. For example, if there were a single aggregate shock common to all cities at each time, consumers could still insure by smoothing consumption through saving and borrowing over time and never moving.

Third, our proposed alternative model has aggregate risk of a specific kind. In the context of perfect competition, each city receives shocks to its productivity at each time. Only the city with the best technology in an industry produces at that time, driving out others. Our equilibrium has consumers moving to the cities producing with the best technology for some industry at that time. Insurance against shocks is too costly, as it is almost the total wage in a productive city. Our framework leads not to the central limit theorem or Gibrat's law, but rather to extreme value theory (the analog of the central limit theorem for maximal values instead of averages) and the Fisher-Tippett (1928) theorem. The implied functional form for the size distribution of cities

is different from the predictions in the literature. Independence of shocks across time was a requirement of the original Fisher-Tippett theorem, but is not required for modern versions.

The paper is organized as follows. First, in the balance of this section, we shall discuss the literature that attempts to refine the stylized fact, namely the rank-size rule, and explain it. Then in section 2 we shall raise specific objections, involving insurance against city-level risk, to these models. In section 3 we introduce Eeckhout's model and modify it to make the objections formal. In section 4, we propose an alternative type of model to explain the size distribution of cities, and implement it empirically. Finally, in section 5, we shall discuss our conclusions and directions for future work.

1.2 The Older Literature

The innovative work of Gabaix (1999a, 1999b) is the source from which the modern literature on the size distribution of cities flows. This work uses an overlapping generations structure where consumers live for two periods. It is assumed that moving costs are so high that consumers can only choose their location (city) when they are young. This location decision is made after shocks to production and amenities are realized for that period, and known to all. The consumer/workers cannot move again when old. The wages or income for the old in a city are never even specified, and it is simply assumed that the young make their decisions in a myopic manner. Moreover, the availability of insurance or capital markets is never discussed, so it is unknown whether the young can hedge against uncertainty about their wage when they are old in the city they choose.

If the old people are immobile, why is this important? It is important because when the young make their decisions, they can anticipate what happens when they are old, and might change their mind about their location decision when young. In other words, they won't behave myopically. Without myopia, insurance becomes important.

1.3 Recent Literature

Chief among recent work are Rossi-Hansberg and Wright (2007), Duranton (2006), Eeckhout (2004) and Duranton (2007). We focus on the latter two.

Eeckhout's model has consumers who are infinitely lived with foresight and who can move each period. There are technological shocks to production

in each city in each time period. It is movement of the consumer/worker population in response to these shocks that generates Gibrat’s law. The shocks generate changes in equilibrium wages, rents, and congestion across time and space that correspond to the consumer movements that equalize utility levels across space at each time. On p. 1445, the following statement is made: “Moreover, because there is no aggregate uncertainty over different locations, and because capital markets are perfect, the location decision in each period depends only on the current period utility. The problem is therefore a static problem of maximizing current utility for a given population distribution, and the population distribution must be such that in all cities, the population $S_{i,t}$ equates utilities across cities.”

Here we wish to make an important distinction between transfers of consumption across time, namely perfect capital markets, and across states, namely complete and perfect futures markets.

The actual consumer optimization problem does not involve state-dependent assets nor does it allow state-contingent transfers of income. If it were to allow this, as in a standard model of complete futures or insurance markets, then agents would never move. They would simply buy assets at the start of time that would pay them under a bad state in their city at a particular time, and such that they would pay under a good realization in their city. In other words, they would insure against the state of nature in their city.²

The basic model of Duranton (2007) has consumers maximizing an intertemporal utility function subject to an intertemporal budget constraint, without facing uncertainty. However, once the detailed urban features are added (in Section V and Duranton, 2006), the model looks similar to Eeckhout’s at least in terms of the urban features. One simply needs some dependence of local prices (land rents or wages) on the state of nature. Then utility equalization implies that people will move depending on the state realization, but this movement disappears if one allows insurance.

There isn’t enough detail about the urban market in Duranton (2006, 2007) to make specific statements about how insurance would work, but the consumers in a city face uncertainty about employment due to the uncertainty about innovations in various industries, so similar insurance arguments should

²It is important to recognize that in this model there are two factors determining a worker/consumer’s productivity, namely the city-specific shock, and the externality in production induced by total population in the city. Even if capital markets are perfect, the production externality is not internalized (even with a land market), so the equilibrium allocation is not necessarily first best.

work if the details of the model are filled in.

The bottom line here is not that complete and perfect futures markets are needed to upset the purpose of these rather fragile models. Rather, it is that any insurance at all will do the trick. Then the question becomes whether moving or buying insurance is cheaper for the consumers. If insurance is incomplete, it can still hinder mobility. Typically in these models and the real world, if moving costs are positive, it makes sense for consumers to stay put.

2 The Issue

2.1 How Insurance Reduces Population Movement

So how might this insurance occur in practice? Let's assume either that consumers cannot perfectly observe the technology shocks to cities, or moving has a small cost, or both.

- Self insurance. Since consumers can transfer consumption across time, and they know that shocks are i.i.d., then they can borrow or use their savings in bad times and save (or pay off their loans) in good, staying in the same city. In the literature, the intertemporal uncertainty faced by consumers does not show up in their objective function, whereas the possibility of self insurance does not show up in the budget constraint. The earlier quote from Eeckhout seems to imply that this is allowed, but the formal statement of the consumer budget constraint makes it clear that this is not allowed. This type of insurance exploits the fact that for any given city, the shocks are i.i.d. over time. Empirically, the place to look for self insurance is in the savings response to local employment shocks.
- Insurance markets. In all of these models, at each time the state of nature (the random shock to each production function for each city) is known to all and verifiable before consumers make their decisions about consumption bundles and location. So this is a perfect setting for a viable insurance market. An insurance firm can step in or the continuum of consumers can simply pool resources in each period, smoothing their consumption without changing location so it is independent of the state

in their city.³ This type of insurance exploits the fact that at any given time, the shocks are i.i.d. across cities. Empirically, one place to look for insurance is a cross-country comparison of how varying benefits of unemployment insurance affect mobility in response to local employment shocks.

- Futures markets. Consumers formulate plans to sell labor and buy consumption commodity and housing contingent on every possible state in every time period. There is no empirical complement. We mention this for completeness.

For our criticism to apply, insurance via either of the vehicles mentioned above need not be perfect. It only need be enough so that it throws off the mobility result, which requires that the response to shocks is *only* in consumer movement, rents and congestion.

Given that for Gibrat’s law to hold, the shocks to each city in each period must be “small” (see Eeckhout, 2004, p. 1447), it seems reasonable to think that insurance would yield higher consumer utility than movement, if moving costs are at all significant or if consumers cannot observe shocks to firms perfectly, and thus face even a small amount of uncertainty in their optimization problems.

We will show below that in general, for models in the literature, consumers will choose to insure instead of move when insurance is available. A new model in which people will choose not to insure but move even when insurance is available is proposed below. A common feature of both the models in the literature and the model presented below is the prediction that people will move and not insure. Thus, if we were to look for empirical evidence to distinguish among the models, evidence of substitution between movement and insurance is insufficient.⁴

³Although landlords (and the destination of land or housing rent) are not made explicit in these models, they might wish to participate in the insurance market as well, since their incomes fluctuate with the state of nature in their city. Of course, risk aversion on the part of landlords simply requires that their utility as a function of rental income be concave.

⁴The empirical investigation of the use of insurance as a substitute for migration, especially when consumer heterogeneity is taken into account, seems quite interesting as a topic for future research.

2.2 Possible Objections to the Criticism

We emphasize that the criticism we make is a purely theoretical point concerning models in the literature. Whether or not agents in the real world actually insure against city-wide risk is not relevant to the question at hand. Our point is that in the theoretical worlds of these models, insurance of the sort discussed in the previous subsection is implicitly excluded. The reasons are not given or, more importantly, included in the model. If these factors, such as asymmetric information, are included in the model to explain insurance market breakdown, other competing forces driving agglomeration can be important; see for example Berliant and Kung (2008).⁵ In other words, this criticism of the internal structure of the models, for example when there is a non-zero moving cost, is that the consumers are not behaving rationally if they don't insure.

Next we present a discussion of why insurance market breakdown is not natural in the context of the models. Again, this is not meant to be a statement about the real world, but rather about whether the exclusion of an insurance option for consumers in the models makes sense.

The usual cause of a breakdown of insurance markets is adverse selection, represented for example by cream-skimming on the part of insurance companies. In the models discussed here, the state is assumed realized and observable to all before decisions are made in a given time period. So there is no issue of adverse selection. But one can easily imagine variations of these models that incorporate some form of information asymmetry. It would not be natural for, say, only consumers to know the shock to the local economy, since the technology shock really affects firms. If only firms knew the realization of the shock before making their decisions, then consumers could draw inferences from firm behavior, or the consumers could self insure or insure. It is not clear what hidden information or hidden action on the part of consumers would cause an insurance market breakdown in this context. It is natural to assume that amenities are observable.

One can imagine moral hazard at the city level with insurance markets, in that a city might try to claim a productivity level lower than the actual one so the residents can collect more insurance money. However, there are no local governments in the models in the literature to coordinate this, and the assumption is that local productivity is observable to all, including non-residents of the city, when they make their location decisions.

⁵This work shows that adverse selection alone can generate agglomeration.

Another objection that could be raised is the commitment required on the part of consumers. In fact, commitment to a plan or contract is a requirement of models that feature self insurance, insurance or futures markets generally. For example, a consumer might experience regret over the purchase of a long-term health insurance contract after the state of the world that tells them that they are healthy is realized. Or the insurance company might experience regret if the consumer turns out to be unhealthy. But they are committed to their contracts. In the models of the size distribution of cities, for example, one could begin the random process of technological change and at any point in time, allow insurance and commitment to begin. Then the population distribution will not change from that point on.

One might easily object to even small moving costs or even a small amount of noise in consumer observations of shocks. Then what we present is another equilibrium, that yields exactly the same period by period utility as the equilibrium studied in this literature. This alternative equilibrium features a uniform distribution of consumers, and does not generate Zipf's law.

Finally, there are costs associated with insurance contracts that, from the point of view of consumers, must be balanced against the cost of moving. Such costs involve lawyers and potentially complex transactions. Moreover, unemployment insurance might fulfill the role of explicit contracts. Self insurance does not suffer from these problems. But credit constraints could limit self insurance. In any case, insurance does not need to be perfect. If there is substitution between insurance and mobility, the type of mobility needed to generate the various empirical distributions of city size can be upset.

But we emphasize again that although these various insurance market imperfections can cause insurance market breakdown, their inclusion in a formal model is necessary to ensure that consumers behave rationally when they don't insure, and the consequences of their inclusion are far from obvious.

3 A Model from the Literature Modified to Include Insurance

3.1 Notation

We use the model of Eeckhout (2004) as the basis for the analysis because it is explicit about consumer behavior, in the form of an optimization problem, as well as endogenous urban variables, namely local wages and land rents. We

conjecture that the other models in the literature can be modified in a similar fashion.

The original model is specified as follows. For complete detail, see Eeckhout (2004, pp. 1445-1446).

Time is discrete and indexed by t . The set of cities is indexed by $i, j \in I$. Production is constant returns to scale. The measure of population in city i at time t is $S_{i,t}$, and let $A_{i,t}$ be the technological productivity parameter of city i at time t . This parameter follows the law of motion:

$$A_{i,t} = A_{i,t-1}(1 + \sigma_{i,t}) \quad (1)$$

where $\sigma_{i,t}$ is the exogenous technological shock to city i at time t . It is assumed that $\sigma_{i,t}$ is i.i.d. with mean 0, symmetric, and satisfies $1 + \sigma_{i,t} > 0$. The positive local externality (spillover) function is given by $a_+(S_{i,t}) > 0$, where $a'_+(S_{i,t}) > 0$. The marginal product of a worker in city i at time t is given by

$$y_{i,t} = A_{i,t}a_+(S_{i,t})$$

Consumers are infinitely lived and identical. In city i at time t , consumption good is $c_{i,t}$, housing or land consumption is $h_{i,t}$ whereas leisure is $1 - l_{i,t}$ for labor supply $l_{i,t} \in [0, 1]$. Utility for a consumer in city i at time t is Cobb-Douglas:

$$u(c_{i,t}, h_{i,t}, l_{i,t}) = c_{i,t}^\alpha h_{i,t}^\beta (1 - l_{i,t})^{1-\alpha-\beta}$$

with $\alpha, \beta, \alpha + \beta \in (0, 1)$. For prices, let the consumption good be numéraire, the price of housing or land in city i at time t be $p_{i,t}$, and let the wage in city i at time t be $w_{i,t}$. The local negative externality or congestion function is given by $a_-(S_{i,t}) \in [0, 1]$, where $a'_-(S_{i,t}) < 0$. The optimization problem of a consumer in city i at time t is:

$$\max_{\{c_{i,t}, h_{i,t}, l_{i,t}\}} c_{i,t}^\alpha h_{i,t}^\beta (1 - l_{i,t})^{1-\alpha-\beta}$$

subject to

$$c_{i,t} + p_{i,t}h_{i,t} \leq w_{i,t}L_{i,t}$$

where $w_{i,t} = A_{i,t}a_+(S_{i,t})$ and $L_{i,t} = a_-(S_{i,t})l_{i,t}$. Total land or housing in a city is H . Using the first order conditions from this optimization problem and market clearance, equilibrium (denoted by asterisks) in city i at time t as a

function of population $S_{i,t}$ can be found:

$$\begin{aligned}
p_{i,t}^* &= \frac{\beta A_{i,t} a_+(S_{i,t}) a_-(S_{i,t}) S_{i,t}}{H} \\
w_{i,t}^* &= A_{i,t} a_+(S_{i,t}) \\
c_{i,t}^* &= \alpha A_{i,t} a_+(S_{i,t}) a_-(S_{i,t}) \\
h_{i,t}^* &= \frac{H}{S_{i,t}} \\
l_{i,t}^* &= \alpha + \beta
\end{aligned}$$

The last equation in particular, indicating that labor supply is independent of population, is an artifact of the Cobb-Douglas specification.

Substituting back into the utility function, indirect equilibrium utility as a function of population $u^*(S_{i,t})$ can be written as

$$u^*(S_{i,t}) = [\alpha A_{i,t} \cdot a_+(S_{i,t}) a_-(S_{i,t})]^\alpha S_{i,t}^{-\beta} H^\beta [1 - \alpha - \beta]^{1-\alpha-\beta} \quad (2)$$

Under free mobility of consumers, indirect utility is equated across cities in each time period, determining their populations as a function of their productivity and their realized history of shocks, summarized by $A_{i,t}$. Instantaneous utility is constant over both time and location in equilibrium. Again using Eeckhout's notation, call this instantaneous utility level U . Denote the local size effect as

$$\Lambda(S) = a_+(S) \cdot a_-(S) \cdot S^{-\beta/\alpha}$$

3.2 Insurance

Let the discount factor be denoted by $\rho \in (0, 1]$. In correspondence with the assumption of complete capital markets, it is assumed that all consumers can borrow or lend at rate $\frac{1}{\rho} - 1$. The consumer optimization problem (at time 0) becomes:

$$\begin{aligned}
&\max_{\{c_{i,t}, h_{i,t}, l_{i,t}\}} \sum_{t=1}^{\infty} \rho^t \cdot c_{i,t}^\alpha h_{i,t}^\beta (1 - l_{i,t})^{1-\alpha-\beta} \\
&\text{subject to} \\
&\sum_{t=1}^{\infty} \rho^t \cdot (c_{i,t} + p_{i,t} h_{i,t}) \leq \sum_{t=1}^{\infty} \rho^t \cdot w_{i,t} L_{i,t}
\end{aligned}$$

As stated by Eeckhout, the problem reduces to the one period optimization problem if there is no insurance or futures markets. Formally, there should

be an expectation in the objective function and a requirement that the budget constraint hold for every state of nature. However, this is omitted in the literature since the problem is reduced to a static optimization problem where the state of nature is observed before consumers make their choices.

There are several important points to be made at this juncture. First, it is useful to imagine the consumers stepping back at $t = 0$ and making decisions about their cities of residence and their consumption bundles for the entire time stream of their infinite lives, contingent on state realizations at each time. Second, and more important, *it does not matter which interpretation of the model one employs*. Specifically, resources can be transferred across states of the world (at any given time) in one or more of several ways. In the end, what a consumer is choosing is their residence and consumption bundle for every time and for every possible state of the world, optimizing utility subject to the budget constraint. The state of the world at time t affects the optimization problem through the prices, $p_{i,t}$ and $w_{i,t}$, and income (through $a_-(S_{i,t})$ and $L_{i,t}$) only. These variables depend on $A_{i,t}$ both directly and indirectly, the latter because $S_{i,t}$ depends on $A_{i,t}$ in equilibrium. The state of the world at time t does not enter into the consumer optimization problem otherwise. For example, it does not enter into the utility function. We could index these prices and incomes by the state of the world, but that would only serve to complicate notation.

As already mentioned, what will matter are only the lifetime choices of residence and consumption bundles, contingent on the state of the world in each period. The method used to actually implement them, via *transfers across states* in a time period as opposed to across time periods, does not matter; there are many possibilities. With complete futures markets, at time $t = 0$ the consumers can sell their labor in every future time period and state, buying consumption good and housing in every future time period and state. With insurance markets, at $t = 0$ the consumers can buy actuarially fair insurance against price and income changes. With self insurance, they can commit to a plan of borrowing and saving under all possible scenarios, namely realizations of states in each time period.

To get the basic idea across, in the next subsection we show how insurance would work from the beginning when all cities have the same initial state (productivity) and population. This yields no movement at any time in equilibrium. In the next subsection, we discuss how to extend this so that insurance can begin from equilibrium of the model at any time t . From that

time on, there is no consumer movement unless the insurance is switched off.

3.2.1 Insurance when the initial state is the same for all cities

To illustrate the ideas behind insurance, we begin with an example where all cities begin with the same state at time 0 and consumer insure from then on.

For notational purposes, let \bar{S} be the mean population of cities, that is $\bar{S} = \frac{\sum_{i \in I} S_{i,t}}{|I|}$, where $|I|$ is the cardinality of the set I .⁶ Let $A_0 = A_{i,0}$ denote the common initial technology level for all the identical cities before the process begins. Let $S_{i,0} = \bar{S}$ for all cities i , so they all have the same initial population. We assume that

$$U = u^*(\bar{S}) = [\alpha A_0 \cdot a_+(\bar{S}) a_-(\bar{S})]^\alpha \bar{S}^{-\beta} H^\beta [1 - \alpha - \beta]^{1-\alpha-\beta}$$

Thus, we assume for illustrative purposes that the initial configuration of shock A_0 and uniform population distribution \bar{S} generate the instantaneous equilibrium utility. This is to get the idea across; in the next section, we will show how to start insurance from equilibrium at an arbitrary given time. In either case, no consumer movement will occur once insurance begins.

With insurance, self insurance, or a futures market (or some combination of all 3), we propose the following equilibrium solution for all cities i and times t :

$$\begin{aligned} \bar{p}_{i,t} &= \frac{\beta A_0 a_+(\bar{S}) a_-(\bar{S}) \bar{S}}{H} \\ \bar{w}_{i,t} &= A_0 a_+(\bar{S}) \\ \bar{c}_{i,t} &= \alpha A_0 a_+(\bar{S}) a_-(\bar{S}) \\ \bar{h}_{i,t} &= \frac{H}{\bar{S}} \\ \bar{l}_{i,t} &= \alpha + \beta \end{aligned}$$

In other words, this is the allocation generated by a constant, over both time and state, allocation with a uniform distribution of consumers. By construction, it generates the same instantaneous utility stream for all consumers in all cities and in all times as both the initial distribution and the equilibrium studied by Eeckhout.

But how does this work in a pragmatic sense? Regarding futures markets, each consumer works the same hours, independent of state. If the state realization is good, i.e. if the consumer is in city i at time 0 and $A_{i,t} > A_0$, income

⁶There are technical issues concerning the cardinality of I , but we shall ignore them here.

in excess of $A_0 a_+ (\bar{S}) \cdot (\alpha + \beta)$ is paid to the market. If the state realization is bad, then the consumer receives income from the market, smoothing consumption. Under self-insurance, the consumer commits to a plan of saving income in a good state, and withdrawing from savings or borrowing in a bad state, thus smoothing consumption. The banks know that $E(A_{i,t}) = A_0$, so they are willing to lend. Under mutual insurance, the same type of idea, with commitment, has consumers who are in cities with good states at time t contributing to an insurance pool, and those in cities with bad states receiving payments from an insurance pool. If the number of cities is large, the law of large numbers implies that the mutual insurance pool is solvent.

It is interesting to note that the phenomenon we describe is something like another manifestation of Starrett's spatial impossibility theorem (see Starrett, 1978; Fujita, 1986; and Fujita and Thisse, 2002 chapter 2.3), though here markets are incomplete due to the presence of unpriced local externalities, both positive and negative.

In summary, the equilibrium time path of utility for every consumer is the same, and constant, under insurance and under the equilibrium that generates movement and eventually becomes lognormal. At the very least, a discussion of why the latter equilibrium is selected should be offered in the literature.

With any moving cost, the insurance or futures market equilibrium (the one denoted with bars) clearly dominates the path with asterisks, the one put forth in the literature. Given a choice between moving along the equilibrium path or insuring at $t = 0$, each consumer will individually choose to insure.

A second, and perhaps more reasonable possibility, is that consumers observe A_{it} imperfectly when they make their location decisions each period. In that case as well, the consumers will insure rather than move, since they are risk averse. This can be seen in equation (2). When consumers cannot perfectly observe $A_{i,t}$, equilibrium expected utility will vary in proportion to $E(A_{i,t})^\alpha$.

3.2.2 Insurance starting when the state is an equilibrium at a given time

The preceding subsection was provided to give intuition. However, it has drawbacks in terms of commitment on the part of consumers if they use mutual insurance at each given time, and on the part of banks and consumers at time 0 if the consumers use self-insurance. Moreover, there is a strong assumption that at time 0, A_0 is the same across cities, each city has the same population

\bar{S} , and this combination produces the instantaneous equilibrium utility level. Here we discuss how to dispense with some of these assumptions.

Suppose that we start running the model without insurance, so that consumers are generally moving around, and stop it at some arbitrary time t . At this time, the instantaneous utility level of each consumer is, of course, U . Consider a consumer in city i and the possibility of self insurance. At that point, the productivity parameter in the city is $A_{i,t}$, and everyone knows from equation (1) that for $t' > t$, $E(A_{i,t'}) = A_{i,t}$. So if the consumers in that city freeze their consumption bundle at whatever it is at that time, and commit to staying in that city and consuming that consumption bundle forever through a plan of borrowing and saving, they will obtain utility level U in each period. This exploits the law of large numbers over time.

Mutual insurance, exploiting the law of large numbers over space at a given time, is more interesting. Pick an arbitrary time t and freeze all the consumers in their locations as well as their consumption bundles. All consumers obtain utility U in this situation at time t . Now consider what would happen if they maintain the same location and consumption bundle in time $t + 1$. Given equation (1), the surplus or deficit in total wage payments for city i is

$$\sigma_{i,t+1} \cdot A_{i,t} \cdot a_+(S_{i,t}) \cdot (\alpha + \beta) \cdot S_{i,t}$$

Thus, to ensure that this system of mutual insurance across cities is solvent at time $t + 1$, it is necessary that

$$\sum_{i \in I} \frac{1}{|I|} \cdot \sigma_{i,t+1} \cdot A_{i,t} \cdot a_+(S_{i,t}) \cdot (\alpha + \beta) \cdot S_{i,t} = 0$$

Although this cannot be assured for finite $|I|$, we can see that as the number of cities $|I|$ tends to infinity, the limiting result is a consequence of a law of large numbers with weights given by $\frac{1}{|I|} \cdot A_{i,t} \cdot a_+(S_{i,t}) \cdot (\alpha + \beta) \cdot S_{i,t}$. In fact, Eeckhout (or a little algebra) shows that there is a constant K such that

$$A_{i,t} \cdot \Lambda(S_{i,t}) = K$$

so the weight attached to $\sigma_{i,t+1}$ can be written as

$$\frac{1}{|I|} \cdot \frac{K}{\Lambda(S_{i,t})} \cdot a_+(S_{i,t}) \cdot (\alpha + \beta) \cdot S_{i,t} = \frac{1}{|I|} \cdot K \cdot a_-(S_{i,t}) \cdot (\alpha + \beta) \cdot S_{i,t}^{1+\beta/\alpha}$$

The size of $S_{i,t}$ can be bounded by constraining the support of the random variable σ , since $A_{i,t}$ and $S_{i,t}$ are positively related. There is an extensive

literature on law of large numbers for sums of weighted random variables. Our framework would fit, for example, in Cabrera and Volodin (2005, Corollary 1).

Notice that there is no commitment required under mutual insurance beyond the next period. So it can be switched on and off as desired, with no consumer movement when it is on, and movement when it is off.

4 Modeling the Size Distribution of Cities

4.1 A Model

This model is loosely based on Duranton (2007), but in the context of perfect competition⁷ instead of monopolistic competition. In contrast with the other models in the literature, there is *economy-wide* risk in addition to *city-level* risk. But this in itself is not sufficient to generate consumer movement. For example, if all cities faced the same shock at each time, namely A_{jt} is independent of j , consumers could still insure against this risk by smoothing their consumption through borrowing and saving. Thus, we employ a more extreme form of aggregate risk.

Time is discrete and all consumers are infinitely lived. Assume that there are many cities (indexed by $j = 1, \dots, m$) and many industries, each producing one consumption commodity (indexed by $i = 1, \dots, n$). All commodities are freely mobile. The production function for commodity i in city j at time t is given by

$$y_{ijt} = A_{ijt} \cdot l_{ijt}$$

where y_{ijt} is the output of commodity i in city j at time t , and l_{ijt} is labor input. The random variable $A_{ijt} \in \mathbb{R}_{++}$ will be discussed in detail shortly. Suppose that each consumer supplies 1 unit of labor inelastically and that the total number of consumers as well as total labor supply is given by N .⁸

In each time period t , each city j receives a random draw for its productivity in producing commodity i , namely A_{ijt} . Since we will be using the Fisher-Tippett limit theorem from extreme value theory rather than the central limit theorem, there is no requirement that these random variables be independent. It is assumed that with probability 1, the random draws for 2 industries at time t for city i are not both maximal among all cities for these given industries.

⁷Since there is no market failure built into our model, equilibrium allocations will be Pareto optimal.

⁸The assumption of Starrett's spatial impossibility theorem that is violated by this model is the assumption of location-independent production sets.

In equilibrium, only the cities with the highest draw of the random variable for some industry will have employees and population. (Alternatively, we could simply classify cities exogenously by industry, and assume that a city in an industry receives only a draw for that industry.) Extensions that imply several cities produce in equilibrium will be discussed shortly, but first we must explain the basic model.

The wage rate for the (freely mobile) population of consumers is given by $w(t)$. In equilibrium, it will be the same across industries.

As is standard in this literature, the utility function of a consumer at time t is given by

$$u(t) = \sum_{i=1}^n \frac{1}{n} c_i(t)^\gamma$$

where $c_i(t)$ is the consumption of commodity i by a consumer at time t and $\gamma \in (0, 1)$. Let $p_i(t)$ be the price of commodity i at time t . Then a consumer's budget constraint is

$$\sum_{i=1}^n p_i(t) \cdot c_i(t) = w(t)$$

Let $\lambda(t)$ be the Lagrange multiplier associated with the budget constraint in the consumer optimization problem. Standard calculations yield demand for commodity i at time t for a single consumer $d_i(t)$:

$$d_i(t) = \left(\frac{\gamma}{-\lambda(t)n \cdot p_i(t)} \right)^{\frac{1}{1-\gamma}}$$

Aggregate demand is given by

$$N \cdot d_i(t) = N \left(\frac{\gamma}{-\lambda(t)n \cdot p_i(t)} \right)^{\frac{1}{1-\gamma}}$$

Profit optimization yields, for each t :

$$\text{For } i = 1, \dots, n, \text{ for } j^* \text{ with } A_{ij^*t} = \max_{1 \leq j \leq m, 0 \leq t' \leq t} A_{ijt'}$$

$$p_i(t) \cdot A_{ij^*t} = w(t)$$

Here we are assuming total recall, in that the best technology from the past is remembered, so new technologies are not used unless they are better than all the old ones. Also, only the best technology in industry i survives, where the best is across all cities and previous time periods. This assumption is made for convenience. We discuss it more below.

Hence

$$\begin{aligned} \text{For } i &= 1, \dots, n, \text{ for } j^* \text{ with } A_{ij^*t} = \max_{1 \leq j \leq m, 0 \leq t' \leq t} A_{ij't'} \\ p_i(t) &= \frac{w(t)}{A_{ij^*t}} \end{aligned} \quad (3)$$

In other words, even though wage is constant across occupied cities, output price varies inversely with the production shock. Consumption commodity market clearance requires, for each t :

$$\begin{aligned} \text{For } i &= 1, \dots, n, \text{ for } j^* \text{ with } A_{ij^*t} = \max_{1 \leq j \leq m, 0 \leq t' \leq t} A_{ij't'} \\ l_{ij^*t} \cdot A_{ij^*t} &= N \cdot d_i(t) = N \left(\frac{\gamma}{-\lambda(t)n \cdot p_i(t)} \right)^{\frac{1}{1-\gamma}} \end{aligned} \quad (4)$$

This is the key equation for our analysis.⁹

Setting the constant $\kappa(t)$ to be

$$\kappa(t) = N \left(\frac{\gamma}{-\lambda(t) \cdot n \cdot w(t)} \right)^{\frac{1}{1-\gamma}}$$

and using (3) and (4), we obtain

$$\begin{aligned} \text{For } i &= 1, \dots, n, \text{ for } j^* \text{ with } A_{ij^*t} = \max_{1 \leq j \leq m, 0 \leq t' \leq t} A_{ij't'} \\ l_{ij^*t} \cdot (A_{ij^*t})^{\frac{\gamma}{\gamma-1}} &= \kappa(t) \end{aligned}$$

Hence

$$\begin{aligned} \text{For } i &= 1, \dots, n, \text{ for } j^* \text{ with } A_{ij^*t} = \max_{1 \leq j \leq m, 0 \leq t' \leq t} A_{ij't'} \\ l_{ij^*t} &= \kappa(t) \cdot (A_{ij^*t})^{\frac{\gamma}{1-\gamma}} \end{aligned} \quad (5)$$

Since $\gamma < 1$, labor usage l_{ij^*t} and the shock A_{ij^*t} are positively correlated. Notice that cities that do not have an industry with the largest shock in that industry at time t are empty.¹⁰

⁹Labor market clearance requires, for each t :

$$\sum_{i=1}^n \sum_{\{j^* \text{ with } A_{ij^*t} = \max_{1 \leq j \leq m, 0 \leq t' \leq t} A_{ij't'}\}} l_{ij^*t} = N$$

¹⁰Existence of an equilibrium is not an issue here, since the equilibrium prices and quantities can be solved analytically. For example, at $t = 1$, setting $p_1(1) = 1$, then $w(1) = A_{1j^*1}$, $p_i(1) = A_{1j^*1}/A_{ij^*1}$, $\lambda(1) = -\frac{\gamma}{nA_{1j^*1}} (\sum_{i=1}^n A_{ij^*1}^{\frac{\gamma}{1-\gamma}})^{1-\gamma}$, $l_{ij^*1} = N \left(\frac{\gamma}{-\lambda(1)nA_{1j^*1}} \right)^{\frac{1}{1-\gamma}} A_{ij^*1}^{\frac{\gamma}{1-\gamma}}$, and so forth. Thus, equilibrium is also unique.

The original work on the asymptotic distribution of maxima drawn from a distribution is due to Fisher and Tippett (1928). Modern, more general treatments are given in Coles (2001) and Embrechts *et al* (1997). We shall return to a discussion of extreme value theory momentarily, but first we will draw the implications for our analysis.

The bottom line from this literature is that A_{ij^*t} has an asymptotic distribution of the following form, known as the generalized extreme value (GEV) distribution:

$$F_{GEV}(x) = \begin{cases} \exp\{-[1 + \xi b(x - u)]^{-\frac{1}{\xi}}\} & \text{when } \xi \neq 0 \\ \exp\{-\exp(-b(x - u))\} & \text{when } \xi = 0 \end{cases}$$

Notice that there are 3 free parameters to be estimated here, namely b , u , and ξ . Also notice that to use rank as the left hand side variable in the regression, one simply computes $1 - F_{GEV}(x)$. But from a pragmatic point of view, it is easier to simply use $\ln(F_{GEV}(x))$ as the left hand side variable.

If there are no upper or lower bounds on the distribution, then $\xi = 0$ and the distribution is Gumbel. If there is an upper bound on the distribution, then $\xi < 0$ and the distribution is reverse Weibull. If there is a lower bound on the distribution, for example 0 in our case, then $\xi > 0$ and the distribution is Fréchet.

Substituting (5),

$$\ln(F(l)) = \begin{cases} -[1 + \xi b((\frac{l}{\kappa(t)})^{\frac{1-\gamma}{\gamma}} - u)]^{-\frac{1}{\xi}} & \text{when } \xi \neq 0 \\ -\exp(-b((\frac{l}{\kappa(t)})^{\frac{1-\gamma}{\gamma}} - u)) & \text{when } \xi = 0 \end{cases} \quad (6)$$

Notice that if we use cross section data, then t and hence $\kappa(t)$ is constant. Thus, in addition to the 3 parameters for the distribution of A_{ij^*t} (namely b , u and ξ), for the distribution of l_{ij^*t} there are two additional parameters, namely κ and γ .

In conclusion, we note that consumers will not want to insure against this risk. If only a small percentage of cities produce at any time, then insurance would cost only slightly less than the wage, so they might as well move and receive the wage in each period. For example, to keep things simple suppose that there are 100 industries (or consumption commodities) and 100 cities in each industry (that is, each city is capable of producing only one commodity). Then there is only one city producing in each industry at each given time, and 100 cities out of 10,000 producing in each given time. As time plays out, as long as some consumers are willing to move, each of the cities producing at a given time will be replaced by another in the industry. The city using old

technology has zero wage and no production. So if some workers don't move, their average wage tends to one percent of the expected new wage with time. Under symmetry of cities in an industry, actuarially fair insurance would cost 99% of the expected new wage. In other words, if workers move they will receive the wage next period, but if they insure they will receive 1% of the wage next period. The only way workers won't move is if they all agree to use old, frozen technology in each industry, and collude so that none will move for a higher wage. In contrast, we assume competitive behavior.

The consumers still might want to insure against aggregate wage volatility (namely movement in $w(t)$ over time) by saving and borrowing to smooth consumption, but their spatial distribution is still as we have laid out.

Returning now to our assumptions and extreme value theory, the original theory of Fisher and Tippett presumed that, fixing i , the random variables, A_{ijt} in our case, were i.i.d. across j and t . Of course, in our context this makes little sense. *In general, the city with the best technology for some good i at a particular time t is more likely to innovate and produce a better technology for the next period than an arbitrary city.* Moreover, it is possible that cities nearby are more likely to innovate than an arbitrary city. Fortunately, much progress has been made in extreme value theory since 1928. The modern versions of the Fisher-Tippett theorem, as given by Coles (2001, Theorem 5.1) and Embrechts *et al* (1997, Theorem 4.4.1) allow some dependence. Specifically, what is required is that the sequence of random variables be stationary and that a form of asymptotic independence (as blocks of random variables become farther apart in time) hold.¹¹ It is also important to note that the model and results can be extended to the case where more than one city in an industry produces. This could happen, for example, if there is transportation cost for consumption goods between cities, so a city with a high realization of productivity for a commodity, but not the highest, might serve a local market. It turns out that extreme value theory applies not only to the maximum of a sequence of random variables, but also to the upper order statistics. A detailed discussion of the results can be found in Embrechts *et al* (1997, Section 4.2).

A few more remarks are in order. First, the role of having different industries i , as in the other models in the literature, is to generate a full distribution of limiting populations rather than just one realization of the asymptotic dis-

¹¹An easy way to fit our structure into the theory is to fix an industry i and imagine that at each time t , there are m subperiods. A city j draws its random variable A_{ijt} in subperiod j of time t .

tribution of city populations. Second, in contrast with other models in the literature, the cities without the best technology for some industry at a given time have zero population, so they don't show up in the data because they are rural. Third, the derivations above work fine if there is not complete recall of previous technologies. For instance, if there were no recall, then the realizations of random variables for all cities at one time are independent of those at another time, so we have a form of block independence that is commonly used in extreme value theory.

4.2 Empirical Implementation

Notice that we are not overly concerned with identification of the 5 parameters in equation (6). As in the balance of the literature on the size distribution of cities, these are just the parameters of a reduced form equilibrium distribution implied by a theory. In essence, the parameters are identified by the functional form itself.

We use the Census 2000 data set also used by Eeckhout.

As noted in the sources we cite for extreme value theory, the most common method of estimating extreme value distributions is to use maximum likelihood. The maximum likelihood estimator (MLE) does not yield the smallest Kolmogorov-Smirnov (KS) statistic in our data set. The KS statistic measures the maximum distance between a sample distribution and its estimate. As noted by Goldstein *et al* (2004) in the context of social networks and later by Eeckhout (2007) in the context of the size distribution of cities, using a simple log-log regression can lead to serious statistical problems. The use of MLE and the KS statistic is preferred. It is interesting to note that both the literature on estimation of the GEV distribution and the literature on Zipf's law seem to be (independently) converging on MLE as the preferred method of estimation.

For purposes of comparison with Eeckhout (2004), we produce estimates using each of the lognormal (his) distribution and the generalized extreme value (our) distribution using equation (6), for both maximum likelihood estimation and minimization of the KS statistic.¹² The tables below summarize the estimation results. The results of maximum likelihood estimation for the lognormal distribution are identical to Eeckhout's.

¹²As usual, with any numerical maximum likelihood estimate, one can never be sure that one has found a global maximum. The analog applies to minimization of the KS statistic.

Table 1: Maximum Likelihood Estimation

Distribution	Parameter Estimates	KS statistic
Lognormal	$\hat{\mu} = 7.2775, \hat{\sigma} = 1.7540$.018913
GEV	$\hat{b} = 1.624, \hat{u} = 1.5959, \hat{\kappa} = 102.00, \hat{\gamma} = .80990, \hat{\xi} = .13726$.006985

Table 2: Estimation by Minimizing Kolmogorov-Smirnov Statistic

Distribution	Parameter Estimates	KS statistic
Lognormal	$\hat{\mu} = 7.2496, \hat{\sigma} = 1.7375$.013378
GEV	$\hat{b} = 1.6274, \hat{u} = 1.5959, \hat{\kappa} = 102.00, \hat{\gamma} = .80990, \hat{\xi} = .1373$.006967

Notice that since $\hat{\xi} > 0$, we indeed obtain a Fréchet distribution.

Of course, the comparison between lognormal and GEV is not quite fair, since there are only two parameters in the lognormal distribution whereas there are five parameters in the generalized extreme value distribution, and these parameters do not contain the parameters used for the lognormal distribution.

Graphically (in color), the estimates and data plots follow.

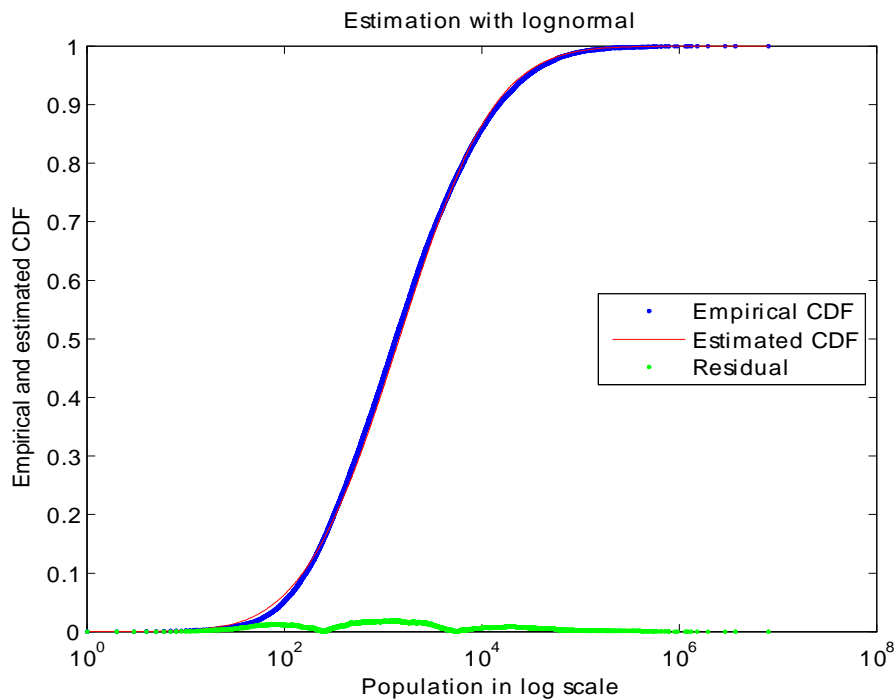


Figure 2: Maximum likelihood estimation of the lognormal distribution

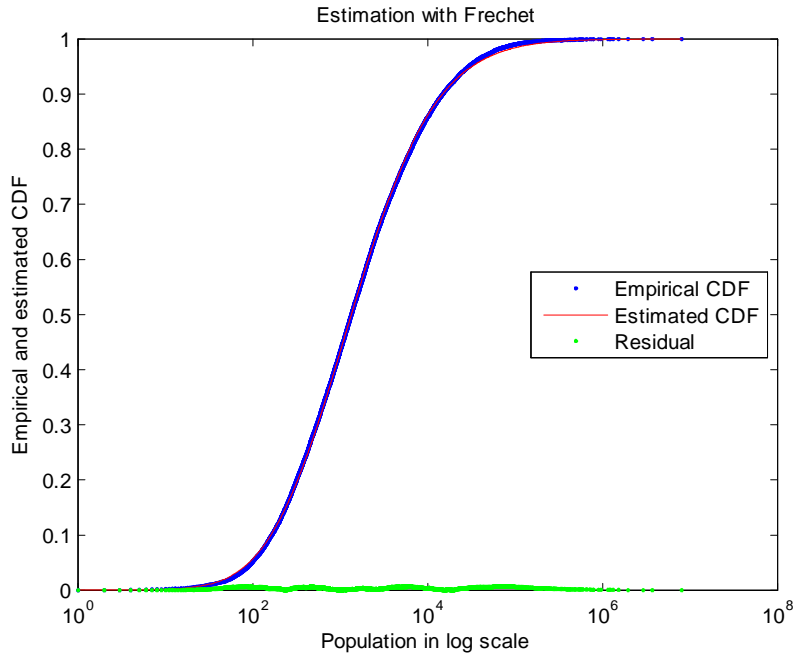


Figure 3: Maximum likelihood estimation of the generalized extreme value distribution

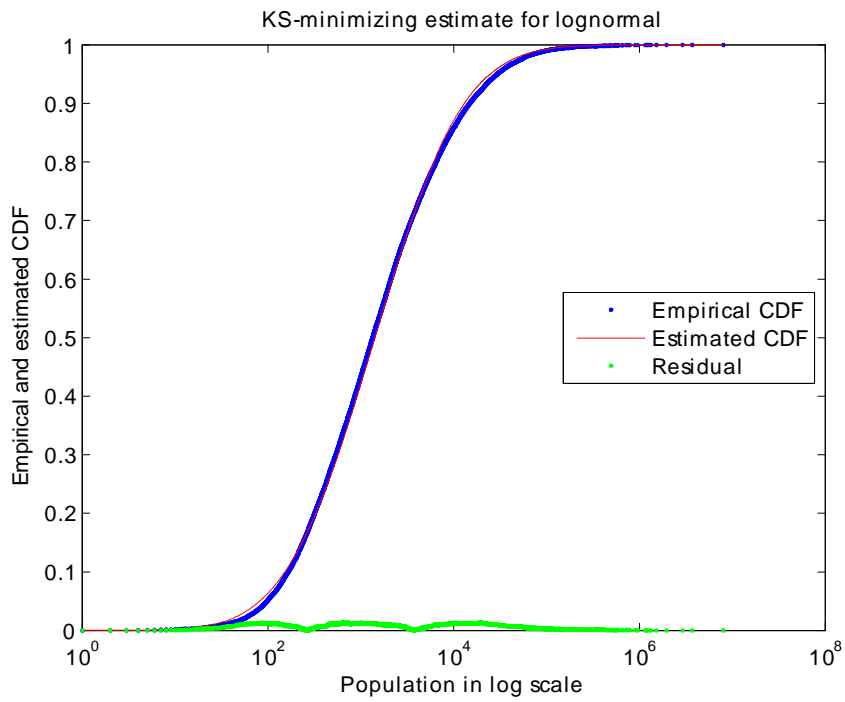


Figure 4: Estimation of the lognormal distribution by minimizing the Kolmogorov-Smirnov statistic

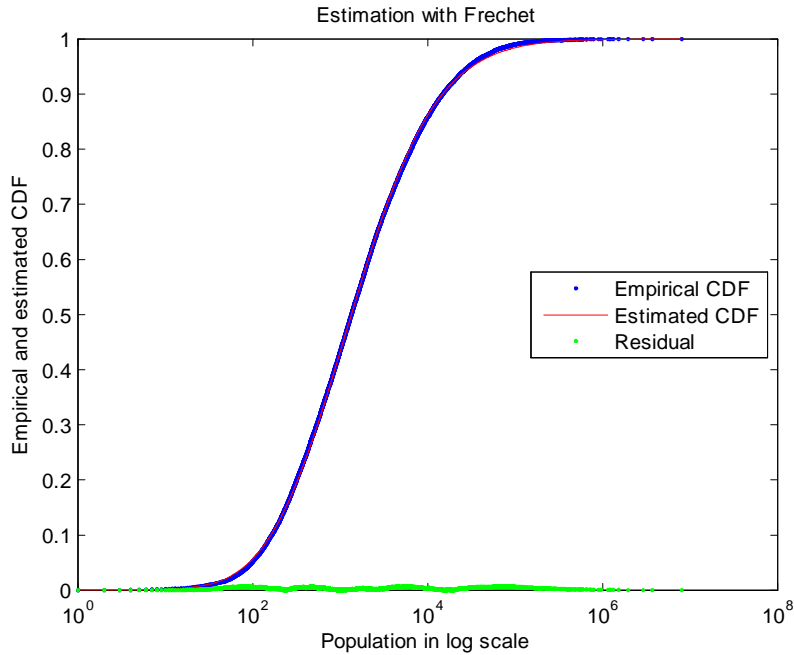


Figure 5: Estimation of the generalized extreme value distribution by minimizing the Kolmogorov-Smirnov statistic

In summary, estimates using the generalized extreme value distribution are quite competitive.

5 Conclusions

So what's the point? Well, actually, there are several related points.

- First, when a model, markedly different from those found previously in the literature, is constructed to explain a specific empirical phenomenon, the microeconomic, structural assumptions about individual behavior and markets must make sense. Here, there is a rather obvious problem that self-insurance and insurance markets are assumed not to be functional. Models in the literature feature city-level risk, and it is generally possible to insure against such risk through many vehicles, barring asymmetric information. The latter does not arise naturally in these models, since consumers are assumed to know the state of nature before making their location and consumption decisions.
- With time in the model, it is even possible to insure against aggregate risk through borrowing and saving.

- However, it is much more difficult to insure against extreme aggregate risk, so we propose such a model. Our model begins with microfoundations and delivers a different functional form for the size distribution of cities than has been used in the literature.
- When one tests a model of the size distribution of cities, the real test is *not* whether it can explain the phenomenon it was specifically designed to explain, but rather whether it is also consistent with regularities that it was not explicitly designed to explain.

Regarding the last item, it is convenient to use an analogy with the agglomeration literature. Cities happen. But that does not provide evidence that any particular model of agglomeration and city formation is correct; there is now a large variety of such models. Further testable hypotheses and evidence, such as predictions about trade, land rent, and wages, are necessary to tease out the contributions of various forces and models. The literature on the size distributions of cities could learn from this example.

In summary, we began with a criticism of the literature based on the fact that a primitive assumption in previous work, that consumers cannot insure (either by borrowing and saving or by pooling resources) against the random productivity variable for each city that is observable to all. Taking Eeckhout's (2004) model as an example, we showed that if insurance is allowed, there is another equilibrium of the model with a uniform distribution of consumers where there is never any migration. Instead, consumers insure against the risk, and the utility stream they obtain in this manner is the same as that in the equilibrium used in the literature. If there is any moving cost or residual uncertainty, the equilibrium used in the literature disappears. Finally, we propose an alternative model based on primitive assumptions, not designed to match any particular empirical factoid. Insurance is allowed, but consumers will never use it, as it is very costly. Instead, they move. The new model is based on extreme value theory and yields a functional form for the size distribution of cities different from the other models, and this prediction is empirically competitive with the ones in the literature.

Future work includes testing further predictions of the model, for example the wage and rent distributions when transport costs for consumption commodities are introduced, and applying the model in new (but appropriate) contexts, such as finance (see Gabaix et al., 2003, for an application of Gibrat's law to finance).

References

- [1] Berliant, M., and F.-c. Kung, 2008, “Can Information Asymmetry Cause Agglomeration?” Working paper.
- [2] Cabrera, M.O., and A.I. Volodin, 2005. “Mean Convergence Theorems and Weak Laws of Large Numbers for Weighted Sums of Random Variables under a Condition of Weighted Integrability.” *Journal of Mathematical Analysis and Applications* 305, 644-658.
- [3] Coles, S., 2001. *An Introduction to Statistical Modelling of Extreme Values*. Springer-Verlag.
- [4] Duranton, Gilles, 2006. “Some Foundations for Zipf’s Law: Product Proliferation and Local Spillovers.” *Regional Science and Urban Economics* 36(4), 542-563.
- [5] Duranton, Gilles, 2007. “Urban Evolutions: The Fast, The Slow, and the Still.” *American Economic Review* 97(1), 197-221.
- [6] Eeckhout, J., 2004. “Gibrat’s Law for (All) Cities.” *American Economic Review* 94(5), 1429-1451.
- [7] Eeckhout, J., 2007. “Gibrat’s Law for (All) Cities: Reply.” Working paper.
- [8] Embrechts, P., Kluppelberg, C., and T. Mikosch, 1997. *Modelling Extremal Events*. Springer-Verlag.
- [9] Fisher, R.A., and L.H.C. Tippett, 1928. “Limiting Forms of the Frequency by Distribution of the Largest or Smallest Members of a Sample.” *Proceedings of the Cambridge Philosophical Society*, 180-190.
- [10] Fujita, M., 1986. Urban land use theory, in *Location Theory*. Edited by J. Lesourne and H. Sonnenschein. New York: Harwood Academic Publishers.
- [11] Fujita, M., and T. Mori, 1997. “Structural Stability and Evolution of Urban Systems.” *Regional Science and Urban Economics* 27, 399-442.
- [12] Fujita, M. and J.-F. Thisse, 2002, *Economics of Agglomeration*. Cambridge: Cambridge University Press.

- [13] Gabaix, X., 1999a. “Zipf’s Law for Cities: An Explanation.” *Quarterly Journal of Economics* 114(3), 739-767.
- [14] Gabaix, X., 1999b. “Zipf’s Law and the Growth of Cities.” *American Economic Review Papers and Proceedings* 89(2), 129-132.
- [15] Gabaix, X.; Gopikrishnan, P.; Plerou, V. and H. E. Stanley, 2003. “A Theory of Power Law Distributions in Financial Market Fluctuations.” *Nature* 423, 267-70.
- [16] Gabaix, Xavier and Y. Ioannides, 2004. “Evolution of City Size Distributions.” Chapter 53, 2341–2378, in J. Vernon Henderson and Jacques-Francois Thisse, eds., *Handbook of Regional and Urban Economics IV: Cities and Geography*, North-Holland, Amsterdam.
- [17] Goldstein, M.L.; Morris, S.A. and G.G. Yen, 2004. “Problems with Fitting to the Power-Law Distribution.” *The European Physical Journal B - Condensed Matter Physics* 41(2), 255-258.
- [18] Halloy, Stephan R.P., 1999. “The Dynamic Contribution of New Crops to the Agricultural Economy: Is it Predictable?” *Perspectives on New Crops and New Uses*, J. Janick (ed.), ASHS Press, Alexandria, VA.
- [19] Janowicz, J.R.; Gray, D.M. and J.W. Pomeroy, 2003. “Spatial Variability of Fall Soil Moisture and Spring Snow Water Equivalent Within a Mountainous Sub-Arctic Watershed.” *60th Eastern Snow Conference*, Sherbrooke, Québec, Canada.
- [20] Limpert, Eckhard; Staehl, Werner A. and Markus Abbt, 2001. “Log-normal Distributions across the Sciences: Keys and Clues.” *Bioscience* 51(5), 341-352.
- [21] Meneghini, R.; Jones, J.A.; Iguchi, T.; Okamoto, K.; and Kwiatkowski, J., 2001. “Statistical Methods of Estimating Average Rainfall over Large Space-Timescales Using Data from the TRMM Precipitation Radar.” *Journal of Applied Meteorology* 40(3), 568-585.
- [22] Rossi-Hansberg, E. and M.L.J. Wright, 2007. “Urban Structure and Growth.” *The Review of Economic Studies* 74(2), 597-624.
- [23] Starrett, D., 1978. Market allocations of location choice in a model with free mobility, *Journal of Economic Theory* 17, 21-37.