

Yannis Menelaos Ioannides, [Henry G. Overman](#), Esteban Rossi-Hansberg and Kurt Schmidheiny

The effect of information and communication technologies on urban structure

Article (Accepted version)
(Refereed)

Original citation:

Ioannides, Yannis Menelaos and Overman, Henry G. and Rossi-Hansberg, Esteban and Schmidheiny, Kurt (2008) The effect of information and communication technologies on urban structure. *Economic Policy*, 54 . pp. 201-242. ISSN 0266-4658

DOI: [10.1111/j.1468-0327.2008.00200.x](https://doi.org/10.1111/j.1468-0327.2008.00200.x)

© 2008 [CEPR](#), [CES](#), [MSH](#). Published by [Wiley-Blackwell](#)

This version available at: <http://eprints.lse.ac.uk/30815/>

Available in LSE Research Online: February 2011

LSE has developed LSE Research Online so that users may access research output of the School. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LSE Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain. You may freely distribute the URL (<http://eprints.lse.ac.uk>) of the LSE Research Online website.

This document is the author's final manuscript accepted version of the journal article, incorporating any revisions agreed during the peer review process. Some differences between this version and the published version may remain. You are advised to consult the publisher's version if you wish to cite from it.

The Effect of Information and Communication Technologies on Urban Structure

Yannis M. Ioannides, Henry G. Overman,

Esteban Rossi-Hansberg, and Kurt Schmidheiny

Tufts University, London School of Economics, Princeton University, and Universitat Pompeu Fabra

December 7, 2007 Paper presented at the 46th Panel Meeting of *Economic Policy* in Lisbon, October 2007.

1. Introduction

The geographic concentration of economic activity occurs because transport costs for goods, people and ideas give individuals and organisations incentives to locate close to each other. If such costs did not exist economic activity would tend to spread evenly over space. Historically, all of these transport costs have been falling. For example, the steam engine, railways, the combustion engine and the use of containers for transportation have all worked to reduce the cost of shipping goods, while the automobile, commuter railways

Thanks go to Kwok Song Too for data, and to Filipe Lage de Sousa, Kay Shan and Stefanie Sieber for research assistance. Comments by Marcus Berliant, Antonio Ciccone, Anna Hardman, Vernon Henderson, Soks Kim, Jed Kolko, Diego Puga, Will Strange, Ping Wang and Zhu Wang, and other participants at presentations at Athens University of Economics and Business, University of Helsinki, University College London, Washington University in Saint Louis, Universitat Pompeu Fabra, Kansas City Federal Reserve Bank, and the International Workshop on “Agglomeration and Growth in Knowledge-based Societies” are gratefully acknowledged. Last but not least, Philippe Martin, the editor in charge and three referees made unusually generous and very perceptive comments on all aspects of the paper. We thank them as well as the Panel members and other participants at the Lisbon meeting.

and the airplane have performed a similar role for the cost of moving people. More recently, new information and communication technologies (ICT) have also significantly reduced the cost of transmitting and communicating information over both long and short distances. Such changes could lead us to predict the death of distance. That is, to suggest that location will no longer matter and that economic activity will, in the near future, be evenly distributed across space. This paper is concerned with one particular aspect of this prediction: the impact that less costly communication and transmission of information might have on cities, the urban structure and the spatial distribution of economic activity.

Two innovations in the twentieth century have changed dramatically the cost of communicating and transmitting information. The first is the widespread adoption of telephony (first fixed line, then mobile), which made possible oral communication over long distances. The second main innovation is the internet and E-mail which has played a similar role for written documents, voice and images. Both these technologies may require substantial upfront fixed investments, but once made they essentially eliminate the link between the cost of communication and the distance between locations.

What are the implications of these changes in ICT for urban structure and the distribution of economic activity in space? This paper provides a partial answer to this question. We begin with a brief description of the adoption path for a number of recent ICT innovations before turning to consider in more detail the ways in which ICT might affect urban structure. We next present the main theoretical argument and identification strategy. Our model suggests that improvements in ICT will increase the dispersion of economic activity across cities. That is, it will make city sizes more uniform. In the empirical section, we test this prediction using cross country data and find empirical support for this conclusion. A concluding section spells out a number of policy implications.

2. ICT and urban structure

Figure 1. Private Cars, Mobile Phones and Personal Computers

About here

Changes in ICT are very clear in the data, especially if we focus on technology adoption. Figure 1 presents the number of cars, phone lines, mobile phones, and personal computers during the last five decades, using Comin and Hobijn's "Historical Cross-Country Technological Adoption Dataset" [Comin and Hobijn (2004)]. The adoption of the telephone was well under way by the 1950's. By the end of the 1990's, the number of telephone lines exceeds 150 million in both Europe and the US. In contrast, changes in personal computers are all concentrated in the 1980's and 1990's. The US went from less than 5 million computers in the early 1980's to more than 140 million computers in the late 1990's. This is a remarkable change that is likely to have very important effects. The data show a similar pattern for the EU that went from less than 5 million personal computers to 100 million in the late 1990's. The EU and the US have also experienced similar dramatic changes in the use of cell phones, but with the growth occurring even

later than for personal computers. Mobile telephones technology was practically unused in 1985, but more than 150 million people owned a mobile phone in the 1990's in the EU, and more than 85 million in the US. Of course, these numbers alone do not reflect the costs associated with implementation of these technologies but they do show the dramatic growth in adoption.

Particularly significant is therefore to contrast the linear growth in number of cars or phones with the exponential growth in new forms of ICT. If part of the role of cities is to save on communication costs, and given that the dramatic growth in ICT adoption has clearly lowered these costs, it seems intuitive that these changes could have some significant implications for urban structure. It is to these implications that we now turn.

Economists use the phrase “agglomeration economies” to describe the advantages that occur when economic activity is densely concentrated. The first source of agglomeration economies is known as *human capital or knowledge spillovers*. In areas of dense economic activity, workers casually exchange knowledge about technology and production conditions at their places of employment. Such transfers happen fortuitously but also are sometimes sought out deliberately by firms, as anecdotal evidence about life in California’s Silicon Valley testifies. A second force for agglomeration is *labour market pooling*. At any point in time, firms are subject to idiosyncratic shocks (e.g. as a result of changing demand for their particular product) that makes them want to hire or fire workers. Because these shocks are idiosyncratic, when one firm is firing, another firm may well be hiring. If firms locate in close proximity to one another, it is easy for workers to move from firms experiencing bad times to those experiencing good times. As the saying in Silicon Valley goes, “people change jobs but not parking lots.” Agglomeration is thus attractive to workers because it helps insure them against idiosyncratic shocks. It is also attractive for firms because it weakens the impact of wages to their own idiosyncratic shock and thus mitigates in part the effect of uncertainty in the economic environment. This actually makes employment in bad times more costly than it would otherwise have been, but this is more than compensated for by the ease of expansion in good times. A third force for agglomeration comes from the greater variety of intermediate products and richer mix of labour skills and expertise that are available in larger urban areas. Greater variety of goods and services lowers prices and wages and also enhances firms’ options in choosing technologies for production and distribution of their products. The associated effects on firms are known as *pecuniary externalities* (as distinct from real externalities, the latter term being reserved for non-market interactions among economic agents’ output decisions). Finally, local amenities due to weather, physical attractiveness, culture or tradition are important factors in enhancing the appeal of particular urban agglomerations. These mechanisms, whose articulation essentially goes back to Alfred Marshall’s *Principles of Economics*, explain at least some of the spatial concentration that we observe throughout the world.

Of course, if these agglomeration economies were the only forces driving the location of economic activity then we would expect to observe extreme spatial concentration. In reality, we do not, because these agglomeration economies are offset by costs (dispersion or congestion forces) as activity becomes increasingly concentrated. These costs take

many forms but all arise from the fact that competition for local resources, broadly defined, increases with spatial concentration. For example, congestion occurs as a result of increased competition for space, firms pay higher rents and wages as result of increased competition for land and workers, while they receive lower prices for their output as a result of increased competition in goods markets. The balance between these agglomeration and dispersion forces determines the spatial structure of the economy.

The strength and importance of these agglomeration and dispersion forces depend on many things, including notably the cost of communicating information across space. Knowledge spillovers, for example, depend on the role that distance plays in inhibiting efficient communication of ideas. The importance of face to face communication shows just how dramatic these distance effects can be. But the telephone, the email and video conferencing, for example, are all reducing these costs of communicating ideas from a distance.

Changing communication costs may also affect the benefits from labour market pooling. Recall, these benefits require workers to move from firm to firm. ICT may increase the efficiency of this process as news about vacancies in one firm are more easily communicated to workers who may be looking for work. A similar story could be told about the benefits of pecuniary externalities. For example, falling information costs allow firms to more easily identify potential suppliers of intermediate goods or workers with particular skills. Turning to dispersions forces, ICT may facilitate e-working, allowing individuals to avoid the high costs of commuting in congested cities. Alternatively, it may increase the competition faced by firms as consumers find it easier to identify alternative sources of supply.

Thus, new ways to transport ideas and to communicate information are, in general, likely to affect all of the agglomeration and dispersion forces that urban and regional economists have identified as key determinants of the concentration of economic activity in space. Thus, independently of the sources of agglomeration forces, ICT will likely have an impact on spatial concentration. In what follows, we illustrate the potential impact of ICT by focusing on production externalities as the main source of agglomeration. This gives us a specific prediction about the impact of ICT that we then confirm using real world data. These effects may also be consistent with other models where ICT has a similar effect on different agglomeration economies. So, our exercise does not allow us to discriminate between different models which predict that ICT will disperse economic activity across cities. It does, however, suggest that models that predict changes to city structure in the opposite direction (i.e. increasing concentration) are not consistent with the data.

ICT can, in principle, have many distinct effects on the distribution of economic activity in space. On one hand, it can increase the spatial scope of knowledge spillovers — it is easier for any professional to acquire context that helps her assess information she casually receives from counterparts in other firms. Therefore, fewer person-to-person interactions may suffice to obtain a better understanding of what other firms are up to. To the extent that knowledge spillovers, whether deliberate (as among employees of the same firm) or fortuitous, are productive, we would expect that ITC would strengthen them

among individuals who are located further apart. Local increasing returns are thus less localized when a wider set of people across an entire country or across countries can interact with each other by using new technologies and while economizing on commuting costs. In this sense, improvements in ICT reduce the importance of the quantities of productive factors employed in a city on that city's productivity. This is the stand we take on the theory we present in the next section and in more detail in Appendix A.

This implies that local urban agglomeration effects become less important and lead to less concentration of people and jobs in a few successful (and larger cities) or urban agglomerations. Agents and firms obtain smaller benefits from locating close to each other and so they locate more evenly in space in order to economize on land rents (and other congestion costs). ICT, in particular, can help businesses create opportunities by improving their communications with other firms, suppliers and clients worldwide. For example, real estate, tourism and hotel operators may market their products directly, without relying on city-based intermediaries. This is important, as most of recent urban growth worldwide has been fuelled by growth in service sectors, while manufacturing has been relocating to smaller urban centres with good transportation links [Henderson (1997)] and often are outsourced to lower cost countries.

These arguments associate ICT with greater spatial dispersion in economic activity. This would, in turn, imply larger concentration of the city size distribution. That is, it would be associated with a reduction in the variance of city sizes. Arguably, this potential advantage may not be fully realized if the interurban transportation system does not develop sufficiently to serve a greater network of urban centres. However, at any given level of development, improvements in ICT increase the incentives for economic activity to relocate to smaller urban centres.

On the other hand, ICT may also make certain local public goods more important as a share of consumption or as a share of inputs. Also, changes in the industrial composition of cities, which have been favouring services, may on balance foster concentration of certain services due to increasing returns at the plant level. London, New York or Paris are attractive in part because there are certain products and services that can only be found there. Similarly, urban living affords better consumption prospects. As individuals spend more on amenities, such as theatres and other artistic activities, certain large cities would become relatively more attractive and therefore likely to grow relative to smaller or medium size cities. On top of increasing the share of some of these goods and services in consumption, better ICT may make these goods and services more readily available and cheaper to consume. Clearly, to the extent that public goods (and other forces of urban concentration) become effective and far-reaching with ICT, we should observe a more dispersed size distribution of cities and a more concentrated spatial distribution of economic activity.

From this verbal discussion -- and other ones in the literature, like Glaeser (1998) which is more extensive but still qualitative -- it should be clear that there are two key paths through which ICT may affect the urban structure. But it is still not clear which effect is likely to dominate. Thus our next step is to develop a theoretical model which will make all these connections clear. In particular, it will connect changes in ICT with changes in

the size distribution of cities. As we will show, the effect on urban structure generated by the model will depend on the particular assumption made on how ICT affects agglomeration forces. This relationship is monotonic and so it helps us design an empirical exercise that is informative on which of these different effects dominates in reality.

3. A Model of ICT and Urban Evolution

As discussed above, any reasonable model of urban structure or of the role of space in economic activity, more generally, would predict that improvements in ICT should have effects on the distribution of economic activity in space. However, no model of urban systems seems to have explicitly incorporated the effects of ICT. We use the theory in Rossi-Hansberg and Wright (2007), from now on RHW, to illustrate how ICT may lead to more urban dispersion.

To reiterate, the trade-off between agglomeration effects —the benefits that firms and individuals obtain from being close to each other— and congestion costs determines the size of cities. Most of these agglomeration effects are related to interactions of different types among individuals. These interactions will be affected by the communication and information technology used by these agents. But how will ICT affect agglomeration forces? And how will these changes in agglomeration forces change the distribution of economic activity in space? The paper addresses this question by considering the impact of these changes on the steady state distribution of city sizes. This requires in turn the specification of a fully dynamic model that also accommodates a stochastic structure.

The potential effect of ICT on the dispersion of the size distribution of cities is not obvious. Essentially, we need to understand whether the consequences of productivity shocks, or of other shocks that cities may experience, will be more or less persistent, and will have larger or smaller effects, the larger agglomeration forces. The model in RHW views the connection between agglomeration effects and productivity as mediated by industry-specific physical and human capital. Agglomeration effects are the result of an externality generated by the amount of human capital and labour employed in the city.

To illustrate this mechanism, suppose that an industry receives a positive and persistent productivity shock. Naturally, firms in the cities that produce in those industries will want to produce more. This implies that they want to use more of the industry specific factors. But the total availability of those factors is given in the current period. So the price of the industry specific physical and human capital increases. The positive productivity shock also implies that cities specialized in that industry will grow as firms employ more workers. The higher price of the industry specific human and physical capital will create incentives to accumulate more of these factors. So next period the industry will have more industry specific factors. Because of the agglomeration effects (and this is the key) having more of these factors will imply more workers being hired and higher productivity, which in turn will elicit further accumulation of factors and induce larger cities, even if next period's productivity shock is lower. That is, the effect of the original productivity shock on city size will be persistent through its effect on the accumulation of industry specific

factors. The stocks of these factors are determined by the accumulated history of the industry's productivity shocks and, therefore, the size distribution of cities is determined by the history of these shocks. It is the long run invariant distribution of these factors across industries which then determine the long run invariant size distribution of cities.

The mechanism described above relies crucially on the impact that the level of human and physical capital has on the level of productivity in a city — that is, on the strength of the agglomeration effects. The stronger these effects, the larger the impact of past productivity and the larger the reaction of city size and industry specific factors to an idiosyncratic productivity shock. If agglomeration forces are very small and the productivity of an industry producing in a given city is essentially independent of the level of human capital and employment in the city (and therefore the level of physical capital), today's productivity shock will have only a temporary effect on city size and no effect on the long term stock of these factors. Hence, cities will not grow and may even decline substantially depending on the history of shocks to an industry. This implies that all cities will have similar sizes and so the distribution of city sizes will be extremely concentrated. If all cities are of similar sizes, the distribution of economic activity in space will exhibit a lot of dispersion. Note that the more concentrated the size distribution of cities the more dispersed the distribution of economic activity in space.

In contrast, if agglomeration effects depend heavily on the amount of factors employed in a city, the effect of past shocks on the stock of industry specific factors will be very important. Cities specialized in industries that received a history of good shocks will be very large, and cities that received a history of bad shocks will be small. Hence, the size distribution of cities will be very dispersed and the distribution of economic activity in space will be very concentrated.

In Appendix A we present the details of the model in RHW that yields the result discussed above in an economy. The model allows for accumulation of physical and human capital and city creation and yields a realistic size distribution of cities. However, the main economic mechanism can be illustrated in a much simplified model, albeit a static and partial equilibrium one, which we sketch here. Consider a city with an aggregate production function that implies output per person in the form

$$Y = A(N)T(N),$$

where $A(N) = \tilde{A}N^\varepsilon$, N is total employment in the city, \tilde{A} is a productivity shock and $T(N) = TN^{-\tau}$ the time agents have left for work after commuting. The latter specification is a short hand in order to express the combined effect of declining marginal productivity of labour and congestion. For an exact derivation within a dynamic general equilibrium model, please see Appendix A. An alternative derivation that is also helpful in understanding interurban spillovers is given in the comment by one of our discussants, Omer Moav, following the paper. We assume that the respective elasticities satisfy $\tau > \varepsilon$ so that in effect city production is subject to diminishing returns.

The parameter ε plays a key role in our analysis. It captures the extent to which the size of the city affects its productivity through knowledge spillovers. More specifically, ε denotes the elasticity of productivity with respect to city size. We model this elasticity as

a result of knowledge spillovers, but in fact it can be the result of any other agglomeration mechanism. Note that the level of ε will also affect the average size of spillovers. So an increase in ε will tend to induce concentration in large cities.

With free labour mobility across cities and identical productivity shock process in all cities it has to be the case that income per person is the same across cities. Call this equilibrium income level w . Then by using the above specifications we have:

$$\tilde{A}TN^{\varepsilon-\tau} = w,$$

which, given w , determines the equilibrium value of population in a particular city N as

$$N = \left(\frac{\tilde{A}T}{w} \right)^{\frac{1}{\tau-\varepsilon}}. \quad (1)$$

So the elasticity of N with respect to \tilde{A} is equal to $1/(\tau-\varepsilon)$. Now let the quality of ICT determine the value of ε . If ε decreases with better ICT, an improvement in ICT leads to a decrease in the elasticity of city sizes with respect to idiosyncratic shocks and, in the long run (under some technical conditions), to a decrease in the variance of the size distribution of cities. If, in contrast, ICT increases ε then it increases the elasticity of city sizes with respect to idiosyncratic shocks and, in the long run, it increases the variance of the size distribution of cities. So the actual effect of ICT on ε is, ultimately, an empirical question that we try to settle in this paper. A similar analysis may be cast in terms of a possible role of telecommuting in affecting the elasticity of time left for work with respect to total city employment, but will not be discussed here for reasons of brevity. Again, the model presented in Appendix A allows for much greater generality and detail.

Of course, changes in ICT will in general change the level of technology too. Local interactions may become less important with ICT, but global ones will then become more important. The latter change would be reflected in an increase in the mean of \tilde{A} .

Ideally one would model explicitly the decision of agents in a city to adopt ICT. This has not been done in the literature in a way that may be readily adopted for our purposes. Here we only study the effect on urban structure given exogenous technology adoption decisions. We refrain from developing such a theory because it is unlikely that in view of our data we would be able to distinguish between alternative theories of the precise role of ICT in affecting the elasticities of the spillover effects.

As it is easy to note, the role played in the dynamic version by the persistence of shocks through their effect on the accumulation of industry specific factors is substituted in this simple model by the distribution of productivity shocks itself. Therefore this simple model cannot, by design, generate a realistic size distribution of cities or a realistic pattern of city dynamics. However, it is important to make sure that the logic of the example survives in a model that can generate basic features of the urban hierarchy. This is not immediate since changes in ε (through ICT) will in general change the accumulation of factors and therefore the distribution of productivities in our simple model above. The dynamic general equilibrium model in Appendix A confirms that the basic insight of the simple model in the previous paragraphs, and the logic outlined at the beginning of this section, goes through in a more complete economic environment.

4. Empirical Methodology

Our theoretical model, as developed in Appendix A, predicts that ICT should make the distribution of city sizes in the long run more concentrated if it weakens agglomeration effects. We study this prediction empirically by looking at the effect of ICT on the distribution of city sizes across different countries.

Box 1

About here

Unfortunately, as will be clear when we discuss our data in Section 5 below, the available data only tend to cover the larger cities in each country. This is a problem for our empirical implementation because such truncated data (i.e. data that do not cover the smaller cities) do not allow us to calculate the mean and variance of the entire city size distribution directly. To get around this problem, we proceed as follows. First, we assume that the city size distribution is Pareto (alternatively referred to as following a power law). Given this assumption we can express the log of the proportion of cities that are larger than S , that is the log of the counter-cumulative of the size distribution of cities, as a linear function of log city size:

$$\ln P(s > S) = -\zeta \ln S_0 + \zeta \ln S, \quad (2)$$

where S_0 denotes the minimum city size, which defines the lower bound of the city size distribution, and ζ the elasticity of the proportion of cities larger than S with respect to S . The latter is a negative number that is commonly referred to as the Zipf coefficient. See Box 1 for details. Given a set of cities and their sizes, an estimate of the Zipf coefficient is provided by running a regression of log rank on log city size. When the distribution is Pareto the Zipf coefficient can be consistently estimated by running the regression only on the sample of the larger cities, that is the upper tail of the distribution.

We have underscored the model's prediction that improvements in ICT will decrease the variance of the cross-sectional distribution of cities. In Appendix A we show that this leads to a higher absolute value of the Zipf coefficient, $|\zeta(S)|$. Hence, the absolute value of the Zipf coefficient increases with improvements in ICT, at least when attention is restricted to the upper tail of city sizes. In other words, improvements in the quality of ICT decrease the variance of the size distribution while they increase the absolute value of the Zipf coefficient. Since the Zipf coefficient is negative, they decrease its algebraic value. This result holds independently of whether or not the city size distribution is Pareto. Of course, if it is not Pareto, the Zipf coefficient will not be a constant, but the model predicts that its value will change in the same direction for all, large enough, city sizes. It is important to note that even though our basic theory implies a Pareto distribution of city sizes only for particular cases (see RHW for details), we approach the data using Zipf coefficients that are specified as independent of city size. That is, we assume that the size distribution is always Pareto.

So our assumption of a particular distribution for city sizes gives us a specification that can be estimated given the data that we have at our disposal and that allows us to directly test our theory on the impact of ICT on the city size distribution. The crucial question is then, of course, whether this is an appropriate assumption. It turns out that empirical evidence gives us good reason to think that the Pareto distribution is a reasonable fit for real world city size distributions across a large numbers of countries and at many different points in time. Even if the distribution is not Pareto, it is likely that reductions in variance will be associated with change for the Zipf coefficient (as discussed above) in the same direction. In the discussion of our empirical results, section 6 below, we also address this concern by presenting robustness checks using other measures of dispersion in the upper tail.

Once we have the estimated Zipf coefficient, we may use it as a summary of the city size distribution for different countries and examine how variations in the city size distribution may be attributed to the observed changes in ICT. Clearly a large number of other factors will determine the city size distribution and we will need to control for these if we want to isolate the effect of ICT. We discuss and motivate these additional controls in the results section below.

To capture the effects of ICT we use data on the number of telephone lines and on access to the internet. Our main focus is on the number of telephone lines for several reasons. First, because we have more data. Second, because the impact on the urban structure will take time and, as Figure 1 shows, the rapid adoption of internet technology has only occurred relatively recently. Third, because for telephone lines we have a way to deal with the endogeneity problem for ICT. That is, we have a way to control for the fact that the number of telephone lines may be driven by the urban structure, rather than the other way round. As no such instrumental variable is available for the internet, our results when we use access to the internet as an additional explanatory variable are necessarily more speculative (although both sets of results point in the same direction). We discuss this issue further below after we provide details of our empirical strategy.

The Zipf coefficient for a country c in year t is obtained from the following regression:

$$R_{ict} = \lambda_{ct} + \zeta_{ct} P_{ict} + e_{ict}, \quad (3)$$

where $R_{ict} = \ln(\text{rank}_{ict})$ is the log of the rank of city i in country c in year t , P_{ict} is the log population of that city, λ_{ct} is a country-year specific intercept, and ζ_{ct} is a country-year specific Zipf coefficient.

We follow Rosen and Resnick (1980) and subsequently Soo (2005) in seeking to understand the determinants of the Zipf coefficient ζ_{ct} in (3) in terms of country characteristics:

$$\zeta_{ct} = \theta_c + \delta t + X_{ct} \eta + \varepsilon_{ct}, \quad (4)$$

where θ_c is a constant which may be country specific, δ is a linear time trend and η is a vector of unknown coefficients. The parameter θ_c , sometimes known as a country fixed effect, controls for the fact that the Zipf coefficient may be higher or lower for some particular country in all years for reasons that are unobservable to us as researchers. The

linear time trend controls for the fact that the Zipf coefficient may be increasing or decreasing over time for all countries for reasons that, again, are unobservable to us as researchers. Finally, the coefficients in η capture the effect of changes in explanatory variables (including ICT) on the Zipf coefficient. The coefficients in equation (4) can be estimated by means of a regression of Zipf coefficients for each country and year, $\hat{\zeta}_{ct}$, that are estimated according to (3), on a collection of explanatory variables X_{ct} that are thought to determine the city size distribution.

We make three modifications to this standard approach. The first deals with a potential bias in the estimation of the Zipf coefficient. In a recent paper, Gabaix and Ibragimov (2006) return to a known bias of the estimate of the Zipf coefficient from Equation (3). This bias arises from the fact that ranks and sizes are obviously correlated. The bias is strong in small samples and their proposed correction¹ is to use $R_{ict}^* = \ln(\text{rank}_{ict} - 0.5)$ in place of the log rank in the left hand side of Equation (3).

The second modification exploits the fact that we have panel data (i.e. data on cities and country characteristics for several years) to control for unobserved country specific determinants of differences in the city size distribution. This is captured by the country specific intercept, or fixed effect, θ_c in equation (4). Of course, as usual, if X_{ct} contains time invariant observed characteristics then their coefficients cannot be separately identified via econometric procedures because the country specific intercept captures *all* time invariant differences across countries whatever their source.

The third modification deals with the likely endogeneity of telephone lines with respect to the urban structure. We want to be able to interpret the coefficient on telephone lines from equation (4) as capturing the impact on the Zipf coefficient of changing the number of telephone lines per capita. That is, we want to be able to talk about the *causal effect* of ICT on urban structure. But what if changes to the urban structure come from some other source (say the increasing use of the automobile, which incidentally we do control for, in part via the road density variable) and countries respond to this by changing the number of telephone lines per capita? Then, the number of telephone lines per capita would be driven by a different underlying factor and thus a function of urban structure, and thus be endogenous. In such a case, our regression will capture an association between ICT and urban structure which includes both the direct causal effect (from ICT to urban structure) and any indirect feedback effects (from urban structure to ICT). In the worst case scenario, there may actually be no causal effect of ICT on urban structure, but we may reach the erroneous conclusion that there is an effect, because our regressions pick up the reverse feedback effect from urban structure to ICT. To control for this we adopt an instrumental variables approach. That is, we look for some variables (i.e. country characteristics) that (1) do not change if the urban structure changes for some exogenous reason (e.g. the increase in the availability of the automobile) and that (2) do not themselves independently affect the urban structure but that are (3) correlated with the

¹ According to Gabaix and Ibragimov (2006), Theorem 1, the correction of the bias in the OLS estimate is optimal in the sense that the proposed transformation of the dependent variable reduces the bias to leading order only.

number of telephone lines. We can then use these instrumental variables to capture the causal effect of the number of telephone lines on urban structure. To do this, we first predict the number of telephone lines in a given country and year on the basis of the values of the instrumental variables. We then include these predicted values in place of the actual values of telephone lines in equation 4. If we find that these predicted values of telephone lines still have an affect on urban structure, then we can reason as follows. We know that the predicted values are not capturing any direct effect of the instrumental variables on urban structure (because we assumed – condition (2) – that these variables do not independently affect the urban structure). We also know that the predicted values cannot be capturing any feedback effect from urban structure to ICT because we are predicting phone lines per capita as a function of the instrumental variables (and we assumed – condition (1) – that these variables were independent of exogenous changes to urban structure). Thus, the predicted values must be capturing the causal impact of ICT on urban structure. Of course, for this strategy to work our instrumental variables must be correlated with ICT – condition (3) – so that the predicted values bear some relation to the actual values.

Our idea is to construct instrumental variables based on the market structure in the telecommunications sector. Clearly market structure in the telecommunications sector should affect the number of phone lines thus satisfying the third condition for a valid instrument. Results that we report below show that this is indeed the case. Telecommunications market structure is also unlikely to have a direct effect on the urban structure (independent of its affect on ICT), thus satisfying the second condition for a valid instrument. What about the first condition that telecommunication market structure is not affected if urban structure changes for some exogenous reason? We find it hard to come up with a convincing story where market structure is directly affected by changes to the urban structure. But it is possible to come up with stories where both market structure and urban structure are being driven by some common factor that we have omitted from our estimation. For example the general trend towards liberalisation of the economy during the 1980's affected many sectors in addition to the telecommunications sector. Perhaps it was the reform of one of these sectors (e.g. transport) that changed urban structure. In that case, our instrument will fail to satisfy the first condition because it is actually changes in attitudes to liberalisation (exogenous to our model) that are changing both urban structure and market structure. By including other measures of liberalisation, we are able to provide some indirect evidence that this is not the case and that our instruments are likely to be valid. But, formally, conditions (1) and (2) are “maintained assumptions” that cannot be easily tested and the instrumental approach relies on their holding. We find these assumptions reasonable, but some readers may not do, in which case our assertion of a causal effect of ICT on urban structure will need to be interpreted with caution.

We can identify three broad market structures for the countries in our sample during the years that we study: competitive, public or private monopoly. Two of the instruments that we use are dummies for whether the country has a public monopoly or a private monopoly with a competitive structure as the excluded category. The other two

instruments measure the time that has passed since the end of the private and public monopolies for countries that have liberalised their telecommunications sector. As no such instrumental variables are available for the internet, our results when we use access to the internet as an additional explanatory variable are necessarily more speculative (although both sets of results do point in the same direction).

Finally, in addition, we report in Appendix B a fourth modification which increases the efficiency of the estimators by implementing a one-step procedure that simultaneously estimates equations (3) and (4).

Before turning to the implementation of our approach, we re-iterate that, although we rely on the Pareto law to motivate our econometric approach, our estimations should capture more generally the impact upon the entire distribution of city sizes that emanate from changes in underlying determinants of interest.

5. Data

We use the same city data as that used in Soo (2005), which were taken from Thomas Brinkhoff's City Population project (<http://www.citypopulation.de>). Soo's paper provides a fairly extensive discussion of the nature of the data, particularly with regard to the issue of the definition of cities.² Data on population, GDP per capita in 2000 Purchasing Power Parity (PPP), trade and government expenditure as a percentage of GDP, non-agricultural economic activity and land area come from the World Bank World Development Indicators (online). Data on kilometres of roads come from the International Road Federation World Road Statistics.³ GDP growth is calculated as log difference of GDP in 2000 PPP; its volatility is measured as the empirical standard deviation over the observed time period.

We use two different measures to evaluate the role of ICT in determining the city size distribution: telephones and the internet. As already discussed, our main focus is on telephones, because we have data for a longer time period and can construct a set of potential instruments. Data on the number of telephone land lines and internet users per 1000 also come from the World Bank World Development Indicators. We multiply these numbers by 1000 and use per capita measure of phone lines and internet users in the estimations.

The information on the market structure of the telecommunications sector (private or public monopoly versus competitive provision) that we use to instrument the number of

² We use data on cities as opposed to urban agglomerations because they are more consistently available internationally.

³ Further details on all these variables are provided in Soo (2005). A large number of our explanatory variables were also kindly provided by Kwok Tong Soo and supplemented by the authors.

telephone lines comes from the OECD International Regulation Database.⁴ This same source provides the data that we use to capture the general degree of competition in the economy when considering the validity of our instruments.

We start with city-level data for 73 countries covering 7,530 different cities recorded at various time periods between 1972 and 2001. There are 197 country-year pairs meaning that, on average, we observe each country 2.7 times. Detailed inspection of the data reveals that the relevant explanatory variables are missing for many countries. Fortunately, the variables are available for three blocks of countries -- North America, Europe and some of the former countries of the Soviet Union.⁵ Deleting countries with missing data leaves us with 24 countries covering 2,955 cities recorded at various time periods between 1980 and 2000. There are now 63 country-year pairs meaning that, on average, we observe each country 2.6 times. For the internet regressions we have to drop Mexico and restrict the time period to 1990 to 2000. This gives us 23 countries and 41-country year pairs covering 2,792 cities. **Error! Reference source not found.** presents the descriptive statistics of the variables used in our empirical analysis below. Additional details on our data variable definitions may be found in the online Data Appendix.

Table 1: Descriptive statistics

About here

6. Results

We start by estimating Equation (3) for each country-year pair to get some idea of the distribution of Zipf coefficients for the countries in question. The mean Zipf coefficient across our 63 country-year pairs is -1.370. The maximum value is -0.928 for Belarus in 1998, the minimum is -1.714 for Belgium in 2000. Interestingly, for all country-year pairs we strongly reject the null hypothesis that the Zipf coefficient is equal to one.⁶ That is, we strongly reject Zipf's law, strictly construed as an estimated Zipf coefficient (ζ - the

⁴ See the Indicators of Product Market Regulation Homepage at: <http://www.oecd.org/eco/pmr> and described in Conway (2006). Missing data points for Eastern European countries were filled by the authors based on media coverage.

⁵ Specifically, we use data on Austria, Belarus, Belgium, Bulgaria, Canada, Denmark, Finland, France, Greece, Hungary, Italy, Mexico, the Netherlands, Norway, Poland, Portugal, Romania, the Russian Federation, the Slovak Republic, Spain, Sweden, Switzerland, the United Kingdom and the United States. We have to drop Germany because of the reunification of 1990 and the ensuing adoption of Federal Republic of Germany institutions in former German Democratic Republic territory causes definitional problems.

⁶ Our result are broadly in line with those reported in Soo's 2005 table 1 which are for the last year for each country-pair and which do not implement the Gabaix-Ibragimov correction.

elasticity of the proportion of cities larger than S with respect to S) being equal to 1. See Box 1, for more details.

We now turn to consider the effect that ICT has on the city size distribution as summarized by the Zipf coefficient. As discussed in the data section, our main focus is on the impact of phone lines. Figure 2 plots our estimates of Zipf coefficients for different countries at different times against the number of telephone lines. Two things stand out from the plot. The first is that, overall, the relationship between Zipf coefficients and the number of telephone lines is as predicted. Second, looking at individual countries, we can see the same effect replicated within the country in terms of changes over time. This is, perhaps, easiest to see for Belgium (BEL) and Austria (AUT), but careful inspection will convince the reader that a similar pattern is observed for many countries. The rest of this section formalises the findings that emerge from this scatter plot, showing that they are robust to the introduction of additional explanatory variables and to controlling for the endogeneity of telephone lines.

**Figure 2. Estimated Zipf Coefficients against Telephone
Mainlines, different countries and years.**

About here

We begin by estimating the simplest possible specification where we treat phone lines and a time trend as the only variable that explains differences in the Zipf coefficient across countries and time. That is, we estimate Equation (4) with phone lines and time as the only explanatory variable (X_{ct}). The results are shown in column [1] of **Error! Reference source not found.** The positive time trend picks up the fact that, on average, across all countries, variance in the city size is increasing over time. Our focus, however is on phone lines and we find that these have a significant negative impact on the Zipf coefficient. That is, the more phone lines a country has, the more concentrated is its city size distribution. This is consistent with our theoretical result that improvements in ICT (or, in the data, an increase in the number of phone lines) lead to smaller local external effects and therefore a more concentrated city size distribution.

Table 2: Phone lines and the city size distribution

About here

Clearly, there are many omitted characteristics of countries that could be correlated with both phone lines and the degree to which population is spread out across the city size distribution. Column [2] begins to address this problem by including several additional explanatory variables. Before turning to discussing the empirical results, we briefly motivate each of the additional control variables.

We include the inverse of road density as a proxy for transport costs within the country. Countries with a low road density are likely to have high transport costs encouraging population to concentrate in just a few cities. Thus, we expect the coefficient on inverse

road density to be positive (fewer roads imply higher inverse road density, higher transport costs and urban population that is more concentrated in fewer cities).

We include three variables to control for the economic and geographic size of the country: population, income and land area. More densely populated countries are likely to have more equal city size distribution, so we expect the coefficient on population to be negative while that on land area should be positive. Although we do not constrain the coefficients to be equal, these two variables could pick up other effects, thus introducing some ambiguity about the expected signs on their coefficients. Given our focus on more developed countries we have no strong prior on the sign of the coefficient on GDP. Models in the New Economic Geography tradition predict that our measure of trade openness (trade as a percentage of GDP) should have a negative effect on spatial concentration and hence on the Zipf coefficient because international trade weakens agglomeration forces (See Chapter 18 of Fujita et al., 1999). Similarly, we would expect higher agricultural production to lead to less concentration and a flatter city size distribution. That is, we expect the coefficient on non-agricultural sectors as a percentage of GDP to be positive. A measure of the size of government (government expenditure as a share of GDP) allows for the possibility that larger governments may imply higher population concentration. That would indeed be the case if (as Ades and Glaeser (1995) emphasize) rent seeking behaviour encourages citizens to locate close to policy makers in the capital city. Conversely, large governments have more means to work against agglomeration forces and support peripheral regions through regional policies. Thus we have no strong priors on the sign of the coefficient on government expenditure.

Finally, we include the standard deviation of the rate of growth of real GDP since the theory underlying our approach (discussed in Appendix A) indicates that a higher volatility of total factor productivity shocks should lead to a larger variance of the size distribution and therefore larger Zipf coefficients. We also include the number of cities as a convenient way of allowing for non-linearities in the Zipf regression.⁷

Results reported in column [2] are in line with our expectations for all variables except, trade as a percentage of GDP (which has the wrong sign), and the inverse of road density and the volatility of GDP (which are insignificant). Introducing country fixed effects and instrumenting for phone lines per capita alters these findings so we consider the issue no further for now. Instead, we draw attention to the fact that introducing all of these controls does not change our conclusion on the role of phone lines. The coefficient is still negative and significant, albeit slightly smaller in absolute value. Thus, introducing a large number of additional controls does not change our conclusion that telephone lines are associated with more concentrated city size distributions.

⁷ Several studies, including notably Black and Henderson (2003), have emphasised that the log rank - log population relationship is roughly concave. The relationship therefore exhibits a steep (negative) slope for the highest ranked cities and a flatter (still negative) slope for lower rank cities. Increasing the number of cities may therefore have a positive impact on the coefficient.

Columns [3] and [4] of **Error! Reference source not found.** report results after introducing a country-specific fixed effect. Column [3] reports results when phone lines and a time trend are the only explanatory variable in X_{ct} . Column [4] reports results when we include the additional control variables. Note that the coefficients for four time invariant variables cannot be identified in the fixed effect specification. These variables are time invariant either because of data availability (government expenditure as a percentage of GDP, standard deviation of GDP growth) or because they show very little time series variation (land area and number of cities). Moving from column [2] to column [3] we see that introducing unobservable country-specific effects decreases the absolute value of the coefficient on phone lines, although the coefficient remains negative and significant. Introducing additional controls in column [4] now leaves the effect of telephone lines per capita essentially unchanged. Note that, as alluded to above, the introduction of fixed effects now gives us signs on inverse road density and on trade as a percentage of GDP that are consistent with our theory (although only trade is significant) Thus, our finding that telephone lines are associated with less dispersed city size distributions is robust to controlling for other country characteristics both observed and unobserved.

Of course, one may still worry that the relationship is being driven by time varying unobserved characteristics of countries and that it is changes in these unobserved characteristics that drive changes in urban structure which, in turn drive changes in our explanatory variables. For example, increasing car ownership may lead to the dispersion of population and telephone lines then respond to that dispersion (rather than vice versa). To control for this, we adopt the standard solution of looking for instrumental variables as discussed, in depth, at the end of section 4.

One might have similar concerns about inverse road density as a proxy for transport costs. That is, more dispersed population leads to more roads and lower transport costs rather than transport costs driving population dispersion. We have experimented with lagged road density as an instrument but this resulted in considerable reductions in sample size and little change in the coefficient on road density. As our main interest is in the ICT variables, which we are able to instrument, we do not worry about this further other than to note that the coefficients on inverse road density should be interpreted with caution.⁸ We assume that all other right hand side variables are exogenous.

Table 3: First stage regressions for Table 2

About here

⁸ In an earlier version of this paper, we also tried to instrument for trade as a percentage of GDP using two dummies indicating when a country joined EU or NAFTA. These two variables turned out to be very weak instruments for trade, but reasonable additional instruments for ICT. For this reason, we no longer instrument trade, but continue to use these two additional instruments. We note that instrumenting trade using these regional trade agreement dummies does not change our main finding on ICT. Using only the market structure instruments for ICT also does not change our results.

First stage regression results from the regression of phone lines per capita on the exogenous and instrumental variables are reported in Table 3. In the cross section (column [5]), both public and private sector monopolies significantly decrease the number of telephone lines. Adding additional control variables (column [6]) substantially weakens the effect of the instruments. However, once we have included a fixed effect (columns [7] and [8]) we see that public monopolies are significantly positively associated with the number of phone lines. The F-test on the joint significance of the instruments for our preferred specification (column [8]) show that the model is fairly well instrumented (i.e. our instruments are fairly highly partially correlated with ICT).

Some of the time series variation in these variables comes from liberalization that moved countries from private monopolies to competition. Most of the variation, however, comes from privatization coupled with liberalization which moved countries from public monopolies to competition.⁹ Our results in the most relevant specification in column [8] suggest that, at least in terms of the *number* of phone lines, the efficiency effects of liberalization were outweighed by changes to public service agreements and the tendency for newly privatized firms to reduce the cross-subsidisation of residential lines by business users. Note that in all cases the predicted phone lines from the first stage regressions are increasing over time. The general increase of telephone mainlines is explained by the time trends as well as increases in population, road density, and GDP.

Columns [5]-[8] in Table 2 show what happens when we use these variables to instrument for the number of phone lines. Column [5] ignores country fixed effects and includes instrumented phone lines as the only explanatory variable. Comparing to column [1] we see that our results change very little. The effect of phone lines turns insignificant as we introduce more explanatory variables (compare column [6] to column [2]) but it is again significant if we introduce fixed effects with phone lines on their own (column [7] versus column [3]) and if we introduce fixed effects and time-varying explanatory variables (column [8] versus column [4]).

In sum, we find a robust negative significant effect of the number of phone lines per capita on the Zipf coefficient. Over our study period, increasing phone lines per capita have tended to cause the dispersion of population across the urban structure resulting in a more concentrated city size distribution.

Columns [1] and [2] of Table 4 show that we reach a similar conclusion for the impact of the internet on the city size distribution. Column [1] presents results from a regression of Zipf coefficients on the number of internet users per capita and a time trend. That is, from estimating equation (4) with internet users per capita and a time trend as the only explanatory variable (X_{ct}). We see a negative significant effect on the Zipf coefficient, although the effect is smaller than that of phone lines. Column [2] shows what happens when we introduce the same additional controls as we did for phone lines. Introducing additional controls more than halves the absolute value of the coefficient on internet users

⁹ During the time period we consider there were no privatisations that replace a public monopoly with a private monopoly, although this had certainly happened in earlier periods (e.g. in the United Kingdom).

per capita and turns it insignificant. Columns [3] and [4] show that the negative effect of the internet vanishes once we also consider telephone mainlines. Results, not reported here, show that instrumenting by means of the same set of instrumental variables gives very similar coefficients, but at a slightly lower level of significance (5% instead of 1% in column [1]). Only 16 countries have more than one year of data for internet usage so, not surprisingly, implementing the fixed effects specification gives insignificant results. Finally, introducing fixed effects and instrumenting leads to coefficients that are essentially zero.¹⁰ This is hardly surprising given the limited number of observations and the fact that market structure in the telecommunications sector does not provide good instruments for the number of internet connections.

Our results about the effect of internet connections on urban structure are encouraging for another reason, too. As Figure 1 documents, the incidence of mobile telephony adoption is similar to that of the internet. While direct measures of mobile telephony are available for some countries, the pattern of availability as of now would reduce substantially sample size. We think it is difficult to find good instruments, especially since the development of mobile telephony markets have been influenced by the different standards adopted by providers. Therefore, we take the results of the internet connections, where in fact telephone use is also controlled for, as suggestive of the impact of mobile telephony, a technology that is conceptually akin to a combination of telephony and the internet.¹¹

6.1. Additional robustness checks

A question concerning the nature of our instrumental variables strategy arises from considering the literature on the political economy of urbanization. This literature suggests that urban concentration may be driven by the general degree of competition in the economic and political spheres. This raises the possibility that some measure of the degree of competition in the economy should be included directly in the regression, violating the second requirement for the validity of an instrument (that it should have no direct effect on urban structure). Using this line of reasoning, one could argue, for example, that the industrial organization of the telecommunications sector is actually just proxying for an overly centralized public sector which favours larger cities. To summarise, if industrial organization of the telecommunications sector is capturing other factors that have a direct effect on urban concentration, then it is inappropriate as an instrument. We explore this possibility by controlling for the general degree of competition in the economy using the OECD “RegRef” indicators of regulatory

¹⁰ These results are available on request and documented in the Data Appendix.

¹¹ Modern mobile, that is cellular telephony was not invented in the same countries where it was first commercialized, such as Japan, but did indeed spread fast in sparsely populated countries, such as the Nordic European countries, with whom is indeed nowadays closely associated. Still, adoption per capita is greatest in densely populated but small countries, like Hong Kong and Luxemburg.

conditions in the airlines, telecommunications, electricity, gas, post, rail, and road freight industries of member countries (Conway and Nicoletti (2006)). The OECD claims that the RegRef indicators are a good proxy for the overall degree of competition in the economy. The RegRef indicators are, unfortunately, only available for a subset of the countries in our data set. Specifically, data are not available for Belarus, Bulgaria, Hungary, Mexico, Poland, Romania, Russian Federation, and Slovak Republic leaving 45 country-year pairs instead of the 63 pairs originally. Re-estimation of the original model (column [8] in Table 2) with the reduced set of countries produces a stronger effect for phone lines, -0.308, which is significant at 1%. Including the average of all RegRef indicators as an additional explanatory variable only slightly reduces the effect of phone lines to -0.280 and its significance to 2%. The coefficient on this new explanatory variable is virtually zero and highly insignificant.¹² Taken together these results suggest that concerns about the validity of our instrument as a result of political economy stories are theoretically interesting, but empirically invalid.

Another question concerns what happens with other measures of urban concentration. We have experimented with the Gini index, the normalized Herfindahl concentration index and the coefficient of variation as dependent variables. These measures reflect different aspects of dispersion and are defined at the country level. The coefficient of variation, a standard measure of dispersion, averages the squared deviations from the mean and then divides by the mean. The normalized Herfindahl concentration index reflects squared normalized city sizes. The Gini coefficient is the mean absolute deviation among all pairs of cities, relative to the mean city size. Table 5 reports results for the coefficient on telephone lines for the same eight specifications that appear in Table 2.

Remarkably for such different measures of dispersion, the results for the coefficients of telephones line per capita with the Gini and Herfindahl indices as dependent variables are – up to a scale factor – very similar to the two-step using Zipf (although just statistically insignificant in column [8]). The coefficient estimates with the coefficient of variation as a dependent variable are also negative and thus consistent with our findings using the Zipf coefficient. It is worth recalling, however, that our choice of Zipf coefficient as a measure for urban concentration is not arbitrary, being driven instead by a desire to link our empirical results firmly with our theoretical model of urban structure. In contrast, while all three of these alternative measures of urban concentration appear intuitive, when estimated on the *truncated* sample of larger cities they are not consistent estimators of the population variables and cannot be linked to our theoretical prediction of decreasing variance for all cities. For this reason, the Zipf coefficient results represent our preferred specification and we take these robustness checks as broadly consistent with our overall findings.

¹² These results are available on request and documented in the Data Appendix.

7. Policy Discussion and Conclusion

We find robust evidence that increases in the number of telephone lines per capita lead to a more concentrated distribution of city sizes and so correspondingly to more dispersion in the distribution of economic activity in space. The basic model that underlies our approach rationalizes this empirical result. As access to telephones improves, the ensuing changes in city size distributions imply that local production externalities decrease. That results in an urban structure that is less dependent on past shocks and hence a size distribution of cities with smaller variance. This smaller variance is reflected in Zipf coefficients that are larger in absolute value.

Figure 3 illustrates the magnitude of our empirical results (using column [8] in Table 2). It assumes that the size distribution of city sizes is Pareto, that the smallest city has 100000 inhabitants and plots the distribution for the average Zipf coefficient in our data (labelled 'actual' in the figure) and the distribution we would expect if the log of phone lines per capita increases by one standard deviation (all other effects are set to zero as this is just for illustration purposes). The figure also compares these two distributions with the one associated with Zipf's law, that is, a Pareto distribution with coefficient minus one, according to Equation (2). The increase in phone lines per capita concentrates the distribution, by making the Zipf relationship steeper. If ICT improves, cities are not as large. For example, the share of cities with more than a million inhabitants is reduced by 0.6 percentage points. Since the share of cities with populations larger than a million is about 4.3%, this implies about a 14% decrease in the number of these large cities. This is a significant change in urban structure!

Figure 3. Effect of Phone Lines per Capita on the Size Distribution of Cities.

About here

We argue that the internet is likely to have similar, or even larger, effects on urban structures once its use has spread more thoroughly through the different economies. So far the evidence on internet usage is more speculative, although it goes in the same direction. The data suggest that as the number of internet users increases we should see effects that are about one tenth of the size of the ones we observe for phone lines. As we argue in the introduction, massive internet adoption is a fairly recent phenomenon, and at least in so far as our measurement of the extent of its adoption is concerned, urban structure may take some time to adjust. This may explain the small coefficients we find in the data.

We find that public and private monopolies on average increase the number of phone lines per capita. An average country with a public monopoly will therefore have a more concentrated size distribution of cities. These effects are also economically substantial. If the average country with a public telephone monopoly transitions to a competitive telecommunications sector, our results indicate that the change in the distribution of city sizes should be about a tenth of the change in Figure 3, but in the opposite direction. That is, the Zipf curve will become flatter as the number of larger cities increases. According to our calculations above, this change would lead to an increase of around 9% in the number of cities with more than a million residents.

Even though the analysis in this paper allows us to derive conclusions about the effect of ICT on the urban structure, as it stands it is not designed to derive implications on welfare. So far we can conclude that according to the theory above, ICT causes a decrease in the strength of intra-urban spillovers. That is, ICT reduces the importance of local factors of production on the city's productivity and has led to the evolution of a more uniform distribution of cities. So far, it looks like ICT trades off spatial concentration of economic activity for total factor productivity, and thus bears a conceptual similarity to infrastructure and other regional policies, that are discussed by Martin (1999) and Baldwin *et al.* (2003), Ch. 17, and typically generate tradeoffs. However, increases in the scope (from urban to possibly regional, national or international) of spillovers or factor complementarities are likely to be associated with ICT as well. As the scope of externalities increases, we should also expect increases in the growth rate of total factor productivity everywhere. ICT implies not only smaller local spillovers but also larger national or international spillovers. In this sense, ICT would behave more like the “win-win” policies discussed in *ibid.*, p. 444. In fact, if only the scope but not the average size of externalities changes, by construction ICT will lead to a welfare gain.

This paper studies the first effect of ICT that leads to changes in the size distribution of cities. This effect by itself is likely to have a negative impact on welfare as it reduces the local externality. However ICT will also have the second effect, on spillovers at a larger spatial scope, which will not affect the distribution of economic activity in space, but is likely to have important implications for the growth of aggregate total factor productivity over time. This effect will have positive effects on welfare. In order to understand the ultimate welfare effects of ICT, one needs to account for both the local and national or international effects. As a first step, we have studied the local implications of ICT only.

The theory presented above implies that we can expect to see agents reallocating across cities as a result of improvements in ICT. Moreover, our empirical results imply both that the reallocation that we have observed, and that will likely observe in the future, are substantial. Moving costs are also important and we have not commented on them so far in this paper. Some of these costs are due to regulation and lack of flexibility in labour markets. Others reflect frictions in the adjustment of urban public infrastructure. Yet other costs are due to the cost of selling and buying homes, and the actual transport and organizational costs involved in moving across cities. Our results highlight the importance of government policy in reducing this type of costs. If moving costs are artificially high, because of government regulation and various types of interference, economies will not be able to take advantage of the gains associated with the ICT improvements, since agents will not respond by redistributing themselves accordingly in space.

It is imperative to allow for the natural reallocation that will result from further improvements in ICT. After all, this reallocation across space is the adjustment by the individuals and by firms to an economic environment where physical location is becoming less important.

REFERENCES

- Ades, Alberto F., and Edward L. Glaeser (1995), "Trade and Circuses: Explaining Urban Giants", *Quarterly Journal of Economics*, 110, February, 195-227.
- Baldwin, Richard, Rikard Forslid, Philippe Martin, Gianmarco Ottaviano, and Frederic Robert-Nicoud (2003), *Economic Geography and Public Policy*, Princeton, NJ: Princeton University Press.
- Black, Duncan and Vernon Henderson (2003), "Urban Evolution in the USA", *Journal of Economic Geography* 3(4), 343-372.
- Comin, Diego A., and Bart Hobijn (2004), "Cross-Country Technological Adoption: Making the Theories Face the Facts," *Journal of Monetary Economics*, 39—83. Data accessible at: <http://www.nber.org/hccta/>.
- Conway Paul and Giuseppe Nicoletti (2006), "Product Market Regulation in Non-manufacturing Sectors in OECD Countries: Measurement and Highlights", OECD Economics Department working paper 58; data accessible at: <http://www.oecd.org/eco/pmr/>.
- Fujita, Masahisa, Paul R. Krugman and Anthony J. Venables (1999), *The Spatial Economy: Cities, Regions, and International Trade*, Cambridge, MA: MIT Press.
- Gabaix, Xavier (1999), "Zipf's Law for Cities: An Explanation," *Quarterly Journal Economics*, 114 (3), 739-767.
- Gabaix, Xavier and Rustam Ibragimov (2006), "Log(Rank-1/2): A Simple Way to Improve the OLS Estimation of Tail Exponents", Harvard Institute of Economic Research Working Papers No. 2106.
- Gabaix, Xavier and Yannis M. Ioannides (2004), 'The Evolution of City Size Distributions,' pp. 2341-2378, in Henderson, J. V., and J.-F. Thisse, eds., *Handbook of Regional and Urban Economics*, Vol. IV, *Cities and Geography*, Amsterdam: North-Holland.
- Glaeser, Edward L. (1998), "Are Cities Dying?" *Journal of Economic Perspectives*, 12 (2), 139-160.
- Henderson, Vernon (1974), 'The Sizes and Types of Cities', *American Economic Review*, 64, March, 640-656.
- Henderson, Vernon (1997), 'Medium Size Cities', *Regional Science and Urban Economics*, 27, 449-470.
- Martin, Philippe (1999), 'Public Policies, Regional Inequalities, and Growth', *Journal of Public Economics*, 73, 85-105.
- Rosen, K. T. and M. Resnick (1980), "The Size Distribution of Cities: an Examination of the Pareto Law and Primacy", *Journal of Urban Economics* 8(2), 156-186.
- Rossi-Hansberg, Esteban and Mark Wright (2007), "Urban Structure and Growth", *Review of Economic Studies* 74, April, 597-624.
- Soo, Kwok Tong (2005), "Zipf's Law for Cities: a Cross-country Investigation", *Regional Science and Urban Economics* 35(3), 239-263.
- Wooldridge, Jeffrey (2007), *Introductory Econometrics: A Modern Approach*, 3rd edition, south-Western.
- Zipf, George K. (1949). *Human Behavior and the Principle of Least Effort*. Addison-Wesley, Cambridge, MA.

BOX 1: Zipf's law and the Pareto distribution

Zipf's law for cities [Zipf (1949)] is an empirical regularity that has attracted considerable interest by researchers. In its strict version, which is also known as the rank-size rule, the law is a deterministic rule that states that the second largest city is half the size of the largest, the third largest city is a third of the size of the largest city, etc. To illustrate, let us take a country (for instance the US), and order its cities by population: New York as the largest has rank 1, Los Angeles as the second largest has rank 2, etc. We then draw a graph, known as Zipf's plot (see Fig. B.1): on the y -axis, we place the log of the rank (New York has log rank $\ln 1$, Los Angeles log rank $\ln 2$); on the x -axis, the log of the population of the corresponding city (which will be called the size of the city). If the rank-size rule holds, this produces a downwards sloping line with slope equal to -1.

Generally, and to a remarkable extent, statistical analyses for many different countries, as Gabaix (1999) emphasizes and Gabaix and Ioannides (2004) discuss in detail, obtain estimated coefficients that are concentrated often around one. This indicates that the size distribution of cities is well approximated by Zipf's law with coefficient one. Nevertheless, there is substantial variation in Zipf's coefficients across time and across countries, a fact that ought to cause some doubts as to full validity of the law.

Consider the three Zipf plots on Fig. B.1. They look quite similar to one another, yet the slopes of ordinary least squares lines fitted to them are not equal to -1. Note that the plot for France is steeper than that for UK which in turn is steeper than that of the US; the respective estimates are -1.55, -1.46, and -1.37, are all estimated with very high precision and using 96, 232 and 552 observations, respectively. Note also that the plot for the US is furthest to the right because its cities are larger than those of the UK with the same rank, whose plot in turn is further out than that of France, for the same reason. The techniques employed in the main part of the paper are aimed at backing out from such differences the effect of ICT across countries and over time.

Figure B.1. Zipf's plots for three countries.

About here.

Can we obtain Zipf's law by means of theoretical arguments? The simplest direct theoretical argument one could make would be by invoking Gibrat's law. If different cities grow randomly with the same expected growth rate and the same variance (Gibrat's Law for means and variances of growth rate), then the limit distribution of city sizes converges to Zipf's law. See Gabaix and Ioannides (2004) for an extensive discussion of this issue. Empirically, on the other hand, Zipf's law for cities is an instance of a power law (see further below for details). Power laws are attractive in various sciences, especially in physics, because they are "scale free", in that they do not depend on the definitions of units of measurement. Naturally, this is an important concern in physics. Rossi-Hansberg and Wright (2007) provide a rigorous justification for a power law that is directly rooted in economic theory. It follows as a special case of the model outlined in Appendix A.

A power law of cities states that the proportion of cities that are greater than a particular city of size S , the counter-cumulative of the size distribution of cities, is of the form:

$$P(s > S) = \left(\frac{S}{S_o} \right)^{\zeta}, \quad (\text{B.1})$$

where ζ denotes a negative parameter, and S_o the lower boundary of the distribution, also a parameter and itself a function of the various determinants of city sizes as discussed in the main body of the paper and Appendix A. The mean city size and the variance associated with the law given by Equation (B.1), which is also known as a Pareto probability distribution, are given by:

$$\bar{S} = E\{S\} = \frac{\zeta}{\zeta+1} S_o; \quad V_0\{S\} = \frac{\zeta}{(\zeta+1)^2(\zeta+2)} S_o^2. \quad (\text{B.2})$$

The mean is finite, if $|\zeta| > 1$; the variance is finite, if $|\zeta| > 2$. Zipf's law is the case of a Pareto law with $\zeta = -1$. These properties also help underscore that the rank-size rule cannot correspond to a reasonable probability distribution, strictly speaking, as such a distribution would have neither a finite mean nor a finite variance.

With these caveats, it is still interesting to note that with our data, reasonable good statistical fits are obtained when we regress the log rank against log city size and a constant, which is the so-called Zipf regression. Interestingly, for all country-year pairs we strongly reject the null hypothesis that the Zipf coefficient is equal to minus one. In other words, we strongly reject Zipf's law, strictly construed.

"Every cloud has a silver lining," however. Starting from an empirical law, like Zipf's law, one may motivate a more general and far reaching inquiry into urban structure and growth and therefore on the determinants of city size distributions. This is what we have sought to do in this paper!

Our use of the Zipf coefficient as a measure of dispersion in this paper is both original and may be easily defended on the basis of the properties of a power law for cities (B.2). Specifically, the coefficient of variation, defined as standard deviation divided by the mean, is given by $(\zeta / (1 + \zeta))^2$, and is therefore a monotone increasing function of the Zipf coefficient only. Similarly, the Gini coefficient is given by $-.5\zeta + 1$.

We think that rigorous research along the lines of our paper helps caution economists, sociologists, urban scientists and econophysicists against undue predictions. E.g., as groups of countries integrate, like the EU, economic forces are unleashed which reshape their urban systems. What is likely to happen to the sizes of their larger cities and their ranks? Zipf's law offers a straightforward prediction. But is it reliable? We think not, and have instead proposed a way to make predictions that rely on underlying determinants of city sizes in a dynamic world.

APPENDIX A: A Model of ICT, Urban Evolution and City Size Distributions

We illustrate how to study the role of ICT on the urban hierarchy by considering the basic theoretical model in RHW. Total factor productivity, that is the level of technology, in industry j at time t is given by

$$\tilde{A}_j = A_j \tilde{H}_j^{\gamma_j} \tilde{N}_j^{\varepsilon_j},$$

where \tilde{H}_j and \tilde{N}_j denote the city's total employment and industry j specific human capital, and $\ln A_j$ is an independent and identically distributed (i.i.d.) productivity shock with mean zero and variance ν across all industries j and time periods t . Thus parameters γ_j and ε_j determine the importance of knowledge spillovers from total employment in industry j and industry j -specific human capital in the economy, which are external to individual firms in the industry but internal to the urban economy due to the presence of city developers. If both parameters γ_j and ε_j are equal to zero there are no external effects and economic activity has no incentive to agglomerate in cities. The larger both of these parameters the more important are a city's total human capital and employment in determining \tilde{A}_j city-specific total factor productivity industry j .

A very simple way to introduce the effect of ICT is therefore to let these two parameters vary with the quality of information technology, t . Namely, let $\gamma_j(t)$ where $\partial\gamma_j(t)/\partial t < 0$ and, similarly, let $\varepsilon_j(t)$ be such that $\partial\varepsilon_j(t)/\partial t < 0$. Essentially, this assumption amounts to ICT's increasing the importance of agglomeration effects since people located far away can now interact at a smaller cost and so people living in the city are less important in determining the city's productivity level. Conversely, we could assume that both $\gamma_j(t)$ and $\varepsilon_j(t)$ depend positively on the quality of ICT, which would be consistent with arguments that emphasize the greater importance of public goods as a result of changes in ICT. Which effect dominates is, ultimately, an empirical question that we try to settle in this paper.

In order for city sizes to be well defined, it will be clarified shortly below that we need to guarantee that the knowledge spillover parameters γ_j and ε_j satisfy $\gamma_j(t) + \varepsilon_j(t) < 1/2$ for all t . Otherwise, cities would, in a sense, be too productive and therefore would grow unboundedly since agglomeration effects would dominate congestion costs at all population levels. As long as this condition is satisfied, as a city grows eventually congestion costs become more important than agglomeration costs and so city sizes are finite.

Cities consist of a central business centre, where all agents work and all production is located, and residential areas surrounding it. Each agent consumes the services of one unit of land per period. For spatial equilibrium within each city agents should be indifferent about where to live in the city. Therefore, equilibrium rents at a distance z from the centre should obey $R(z) = \tau(\bar{z} - z)$, where \bar{z} denotes a city's radius, where rent is equal to 0. Hence, total rents in a city of radius \bar{z} are given by

$$TR = \int_0^{\bar{z}} 2\pi z R(z) dz = \frac{\pi\tau}{3} \bar{z}^3 = \frac{b}{2} \tilde{N}^{\frac{3}{2}}.$$

since everyone in the city lives in one unit of land, a city of population \tilde{N} and $b \equiv 2\pi^{-\frac{1}{2}}\tau/3$. Total commuting costs are given by

$$TCC = \int_0^{\tilde{z}} 2\pi z \tau z dz = b\tilde{N}^{\frac{3}{2}}.$$

Assuming the presence of city developers or governments that internalize city-wide externalities, RHW show that in this framework the unique equilibrium allocation may be obtained as a solution to the following planning problem¹³: Choose state contingent sequences $\{C_{ij}, X_{ij}, N_{ij}, \mu_{ij}, u_{ij}, K_{ij}, H_{ij}\}_{t=0, j=1}^{\infty, J}$ to maximize

$$(1-\delta)E_0 \left[\sum_{t=0}^{\infty} \delta^t N_t \left(\sum_{i=1}^J \ln C_{it} / N_t \right) \right] \quad (\text{A.1})$$

subject to, for all t and j ,

$$C_{ij} + X_{ij} + b \left(\frac{N_{ij}}{\mu_{ij}} \right)^{\frac{3}{2}} \mu_{ij} \leq A_{ij} \left(\frac{K_{ij}}{\mu_{ij}} \right)^{\beta_j} \left(\frac{H_{ij}}{\mu_{ij}} \right)^{\alpha_j + \gamma_j} \left(\frac{N_{ij}}{\mu_{ij}} \right)^{1-\alpha_j-\beta_j+\epsilon_j} u_{ij}^{1-\alpha_j-\beta_j} \mu_{ij}, \quad (\text{A.2})$$

$$N_t = \sum_{j=1}^J N_{tj}, \quad (\text{A.3})$$

$$K_{t+1j} = K_{tj}^{\omega_j} X_{tj}^{1-\omega_j}, \quad (\text{A.4})$$

$$H_{t+1j} = H_{tj} \left[B_j^0 + (1-u_{tj})B_j^1 \right], \quad (\text{A.5})$$

where: N_{ij} , K_{ij} , X_{ij} and H_{ij} denote total employment, total physical capital, physical capital investment, and total human capital in industry j in the economy and α_j and β_j are positive parameters satisfying $0 < \alpha_j + \beta_j < 1$; C_{ij} denotes total consumption of representative household i , and μ_{ij} the number of cities producing goods in industry j ; u_{ij} denotes the fraction of time agent i devotes to work. Thus, the maximization problem above amounts to maximizing the sum total of households' lifetime utilities, (A.1), subject to: a resource constraint, Equation (A.2), which expresses that the use of resources for consumption, investment, and commuting costs may not exceed current output (and presupposes that rents are redistributed back to the city residents); a labour market equilibrium condition, Equation (A.3), according to which the total labour force is allocated to all industries and all cities (free labour mobility); and the two factor accumulation equations, for physical and human capital. That is, respectively, Equation (A.4), according to which current investment and the existing capital stock produce capital stock in the next period (with ω_j being a parameter satisfying $0 < \omega_j < 1$), and

¹³ RHW describe how to solve for an equilibrium in which some of the externalities are not internalized by city planners or developers and, therefore, cities are inefficient. Independently of the ability to internalize all potential urban externalities or not, this problem can be solved using a pseudo-planner problem and in all cases yields results we underscore in this paper. The main role that city developers or city governments play is to coordinate agents to work in the equilibrium number of cities. Land developers and city governments traditionally have played this coordination role by offering guarantees and incentives to locate in a particular area.

Equation (A.5), where human capital is augmented at a rate that depends on the portion of each individual's endowment of leisure not allocated for work.

The problem of maximizing with respect to the number of cities μ_{ij} is a static problem with first order condition

$$s_{ij} = \frac{N_{ij}}{\mu_{ij}} = \left[\frac{2(\gamma_j + \varepsilon_j)}{b} A_j K_{ij}^{\beta_j} H_{ij}^{\alpha_j + \gamma_j} N_{ij}^{-\alpha_j - \beta_j + \varepsilon_j} u_{ij}^{1 - \alpha_j - \beta_j} \mu_{ij}^{-\gamma_j - \varepsilon_j} \right]^2.$$

So the size of a city, s_{ij} , with core industry j is then given by

$$s_{ij} = \left[\left[\frac{2(\gamma_j + \varepsilon_j)}{b} \right]^{\frac{1}{1 - 2(\gamma_j + \varepsilon_j)}} F_j \hat{A}_j H_{ij}^{\hat{\alpha}_j} K_{ij}^{\hat{\beta}_j} N_{ij}^{-\hat{\alpha}_j - \hat{\beta}_j} u_{ij}^{\hat{\phi}_j} \right]^2, \quad (\text{A.6})$$

where the auxiliary variables \hat{A}_j , $\hat{\alpha}_j$, $\hat{\beta}_j$, and $\hat{\phi}_j$ are defined as:

$$\hat{A}_j = A_j^{\frac{1}{1 - 2(\gamma_j + \varepsilon_j)}}, \quad \hat{\alpha}_j = \frac{\alpha_j + \gamma_j}{1 - 2(\gamma_j + \varepsilon_j)},$$

$$\hat{\beta}_j = \frac{\beta_j}{1 - 2(\gamma_j + \varepsilon_j)}, \quad \text{and} \quad \hat{\phi}_j = \frac{1 - \alpha_j - \beta_j}{1 - 2(\gamma_j + \varepsilon_j)}.$$

Given this log-linear specification, RHW show that capital investments and consumption in each industry are constant fractions of output net of commuting costs, and the fraction of time devoted to work u_{ij} is constant across time. Taking natural logarithms of both sides of Equation (A.6) that defines city size, we get

$$\ln s_{ij} = 2 \left(\psi_{ij} + \frac{1}{1 - 2(\gamma_j + \varepsilon_j)} \ln \frac{2(\gamma_j + \varepsilon_j)}{b} + \frac{1}{1 - 2(\gamma_j + \varepsilon_j)} \ln A_j + \hat{\beta}_j \ln K_{ij} \right) \quad (\text{A.7})$$

where the auxiliary variable ψ_{ij} includes all non-stochastic variables that enter the planning problem, including N_{ij} and H_{ij} .

If $\ln A_j$ and K_{ij} are the sole stochastic variables in Equation (A.7), then the mean and variance of city sizes are easily obtained and given, respectively, by

$$\ln s_{ij} = 2 \left(\psi_{ij} + \frac{1}{1 - 2(\gamma_j + \varepsilon_j)} \ln \frac{2(\gamma_j + \varepsilon_j)}{b} + \hat{\beta}_j E(\ln K_{ij}) \right), \quad (\text{A.8})$$

$$V_0(\ln s_{ij}) = 4 \left(\frac{1}{1 - 2(\gamma_j + \varepsilon_j)} \right)^2 v + 4(\hat{\beta}_j)^2 V(\ln K_{ij}).$$

It is now clear why condition $\gamma_j(t) + \varepsilon_j(t) < 1/2$ must hold. They are to ensure that the mean and variance of the city size distribution are mathematically well defined. RHW show that as $t \rightarrow \infty$ the variance of the log of physical capital in industry j is:

$$V_0[\ln K_{ij}] = \frac{v}{(1 + \hat{\beta}_j)^2},$$

so that the variance of the long run log-city size distribution may be obtained from (A.8) and given by

$$V_0 [\ln s_{ij}] = 4v \left(\left(\frac{1}{1-2(\gamma_j + \varepsilon_j)} \right)^2 + \left(\frac{\beta_j}{1-2(\gamma_j + \varepsilon_j) + \beta_j} \right)^2 \right). \quad (\text{A.9})$$

Note that the variance of the city size distribution is then increasing in $\gamma_j(t_i) + \varepsilon_j(t_i)$. Therefore, any assumption that we make about the dependence of these elasticities on ICT is reflected on changes on the invariant distribution of city sizes.

In the example of section 3 in the main text, roughly speaking, the parameter $\gamma_j(t_i) + \varepsilon_j(t_i)$ corresponds to ε here. The parameter τ corresponds to b here.

We would also like to connect the variance of the distribution of city sizes to the Zipf's coefficient, in order to be able to connect our theoretical results with the data through this coefficient. The local Zipf coefficient is given by the elasticity of the counter-cumulative of the city size distribution, $P(s > S)$, with respect to city size,

$$\zeta(S) = \frac{S}{P(s > S)} \frac{\partial P(s > S)}{\partial S} < 0.$$

Given the mean of the distribution of city sizes, as we increase the variance we are shifting mass to the tails of the distribution. This implies that for S high enough (large enough city sizes) the term $|\zeta(S)|$ will be smaller the larger the variance. As the variance goes to infinity, $\zeta(S) > -2$, $\lim_{s \rightarrow \infty} \zeta(S)$ converges to the Pareto coefficient.

APPENDIX B: One-step estimation

Earlier versions of this paper reported one-step estimation results for the relationship between telephone mainlines and the Zipf coefficient. This appendix reports these earlier results for comparison. While the one-step procedure is more efficient, it is less intuitive and requires slightly stronger assumptions on the error term.

In the one-step procedure we substitute $\hat{\zeta}_{ct}$ from Equation (4) into equation (3) and estimate directly:

$$R_{ict}^* = \lambda_{ct} + \theta_c P_{ict} + \delta t P_{ict} + [X_{ct} P_{ict}] \eta + e_{ict}^* , \quad (5)$$

where $e_{ict}^* = e_{ict} + \varepsilon_{ct} P_{ict}$ is now composed of an i.i.d. error plus the interaction between rank and error from the first step, while all other coefficients and variables are as defined before. The interacted variable $X_{ct} P_{ict}$ is instrumented by using the predicted values from a regression of $X_{ct} P_{ict}$ on $Z_{ct} P_{ict}$ and the other exogenous explanatory variables (see Wooldridge (2007, pp. 235) for why this is the correct approach).

Results for telephones and internet usage per capita are reported in Tables 6 and 7 respectively. Focusing only on our key variables of interest, we see that this one-step procedure gives the same substantive effect of telephones and internet usage on urban structure although the more efficient estimation procedure increases the significance of the coefficients. As the one-step estimation procedure requires stronger assumptions on the error term, and as these particular empirical results do not correct for the resulting complicated error structure in (5), we view these results as placing an upper bound on the significance of the two core explanatory variables of the urban structure, telephone lines and internet users per capita.




Figure 1. Private Cars, Mobile Phones and Personal Computers.
Source: Comin and Hobijin (2004).

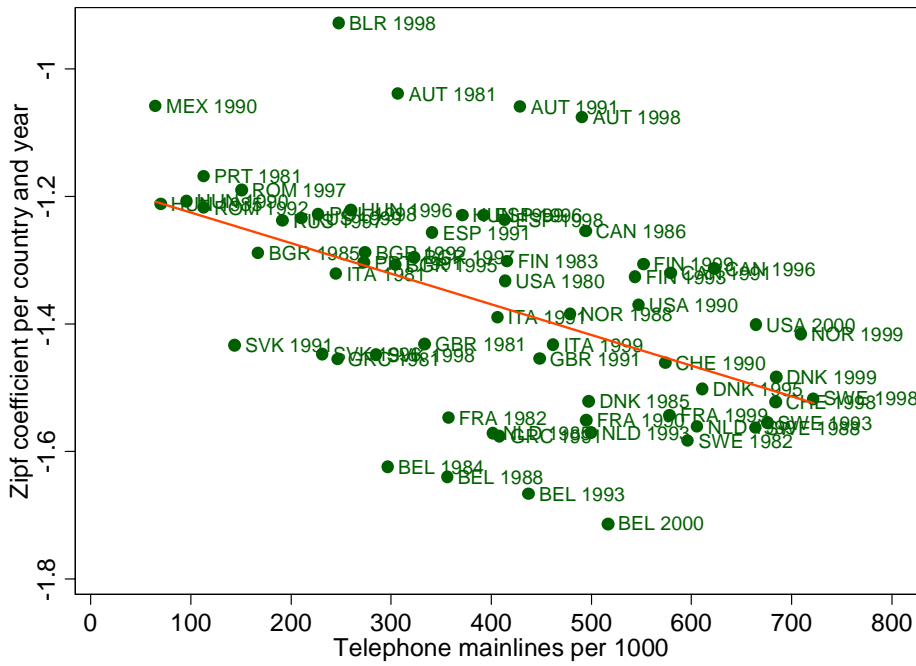


Figure 2. Estimated Zipf Coefficients against Telephone Mainlines per 1000 Inhabitants, different countries and years. Source: Authors' own calculations.

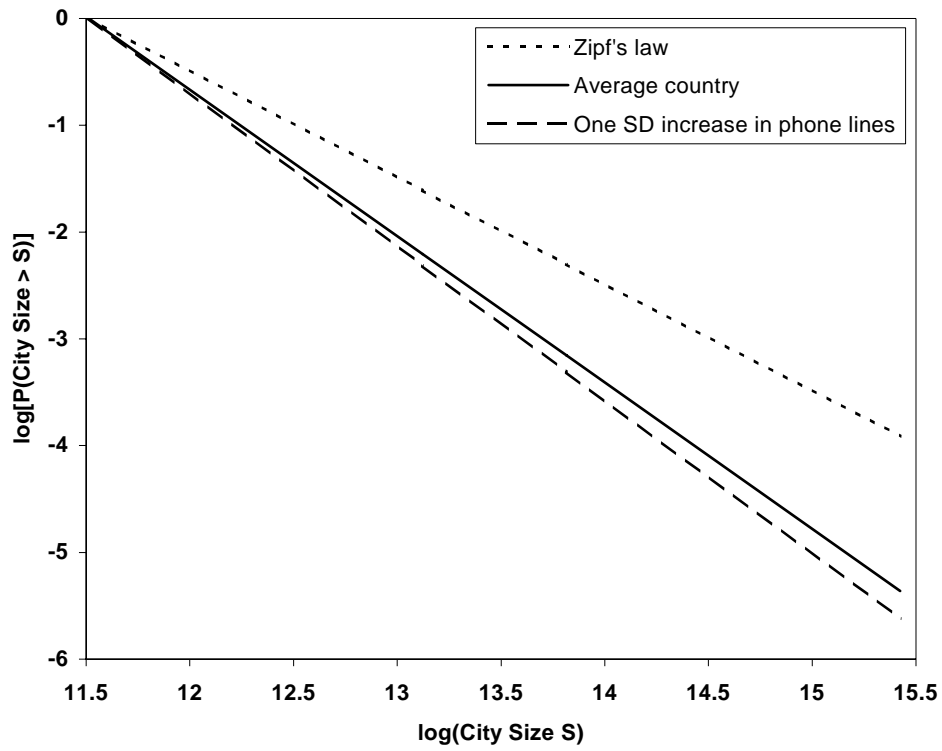


Figure 3. Effect of Phone Lines per Capita on the Size Distribution of Cities.

Source: Authors' own calculations.




Figure B.1. Zipf's plots for three countries.

Source: Authors' own calculations.