

The systematic structure and predictability of urban business diversity

Hyejin Youn
Luís M. A. Bettencourt
José Lobo
Deborah Strumsky
Horacio Samaniego

SFI WORKING PAPER: 2014-09-033

SFI Working Papers contain accounts of scientific work of the author(s) and do not necessarily represent the views of the Santa Fe Institute. We accept papers intended for publication in peer-reviewed journals or proceedings volumes, but not papers that have already appeared in print. Except for papers by our external faculty, papers must be based on work done at SFI, inspired by an invited visit to or collaboration at SFI, or funded by an SFI grant.

©NOTICE: This working paper is included by permission of the contributing author(s) as a means to ensure timely distribution of the scholarly and technical work on a non-commercial basis. Copyright and all rights therein are maintained by the author(s). It is understood that all persons copying this information will adhere to the terms and constraints invoked by each author's copyright. These works may be reposted only with the explicit permission of the copyright holder.

www.santafe.edu



SANTA FE INSTITUTE

The systematic structure and predictability of urban business diversity

Hyejin Youn^{1,2,3*}, Luís M. A. Bettencourt³, José Lobo⁵,
Deborah Strumsky⁶, Horacio Samaniego⁷, Geoffrey B. West³

¹ Institute for New Economic Thinking at the Oxford Martin School, Oxford,
OX2 6ED, UK

² Mathematical Institute, University of Oxford, Oxford, OX2 6GG, UK

³ Santa Fe Institute, 1399 Hyde Park Road, Santa Fe NM 87501, USA

⁵ School of Sustainability, Arizona State University, Tempe, AZ, USA

⁶ Department of Geography & Earth Sciences, University of
North Carolina-Charlotte, Charlotte, NC, USA

⁷ Instituto de Conservacion, Biodiversidad & Territorio, Universidad Austral de Chile,
Valdivia, Chile

September 5, 2014

Abstract

Understanding cities is central to addressing major global challenges from climate and health to economic resilience. Although increasingly perceived as fundamental socio-economic units, the detailed fabric of urban economic activities is only now accessible to comprehensive analyses with the availability of large datasets. Here, we study abundances of business categories across U.S. metropolitan statistical areas to investigate how diversity of economic activities depends on city size. A universal structure common to all cities is revealed, manifesting self-similarity in internal economic structure as well as aggregated metrics (GDP, patents, crime). A derivation is presented that explains universality and the observed empirical distribution. The model incorporates a generalized preferential attachment process with ceaseless introduction of new business types. Combined with scaling analyses for individual categories, the theory quantitatively predicts how individual business types systematically change rank with city size, thereby providing a quantitative means for estimating their expected abundances as a function of city size. These

results shed light on processes of economic differentiation with scale, suggesting a general structure for the growth of national economies as integrated urban systems.

Subject terms: Physical sciences; Applied physics

Diversity is central to the resilience of complex adaptive systems whether ecosystems or economies (1–3). In particular, it has been argued that the success and resilience of cities, together with their role in innovation and wealth creation, are driven by their ever-expanding diversity (2–9). The presence and ever-changing admixture of individuals, ethnicities, cultural activities, businesses, services, and social interactions is a defining characteristic of urban life. Together with its counterpart, specialization, this is often cited as what makes a city unique and distinctive, and has consequently featured prominently in the study of cities across economics, geography and urban planning. Despite its acknowledged importance, however, there have been surprisingly few quantitative investigations into possible systematic regularities and underlying dynamics that govern the diversity of cities across an entire urban system. A recurrent goal in developing an overarching science of cities is to discover, and conceptually understand, general patterns for how people, infrastructure and economic activity are organized and inter-related (10–12). A compelling question therefore is how is diversity related to aggregate urban socio-economic and infrastructural metrics and how do these depend on city size (13–15, 17)?

The systematic quantitative understanding of diversity requires two issues be addressed. Measuring diversity typically involves identifying different (business) types and counting their frequency for a given unit of analysis, such as a city or a nation (1). It should be immediately clear, then, that such a task can be problematic because any systematic classification scheme is subject to an arbitrary recognition of specific categories, since any business type can be subdivided further as long as a defining distinction is made. Restaurants, for example, can be decomposed into fine dining, fast food, etc, and into their cuisine, price, quality etc. In general, therefore, urban diversity is scale-dependent and we should seek a resolution-independent characterization. Secondly, we need to deconvolute the intricate relations between scale, diversity and economic productivity, and between diversification and specialization. New business types must involve a larger number of people, both as workers and clients, and, to be sustained, should lead to greater economic productivity, by permitting, for example, greater specialization and interdependence. Thus we may expect that the larger scale of bigger cities should afford greater economic diversity (at least in absolute terms) but such an expectation is at odds with the idea that specialization drives increase in efficiency.

In this paper we present a novel approach to measuring and characterizing economic diversity in order to clarify its underlying role in urban economic development. Our analysis reveals a surprising systematic behaviour common to all cities. We show how this can be derived theoretically and present a simple model for understanding its structure based on a variant of preferential attachment for introducing new business types. The

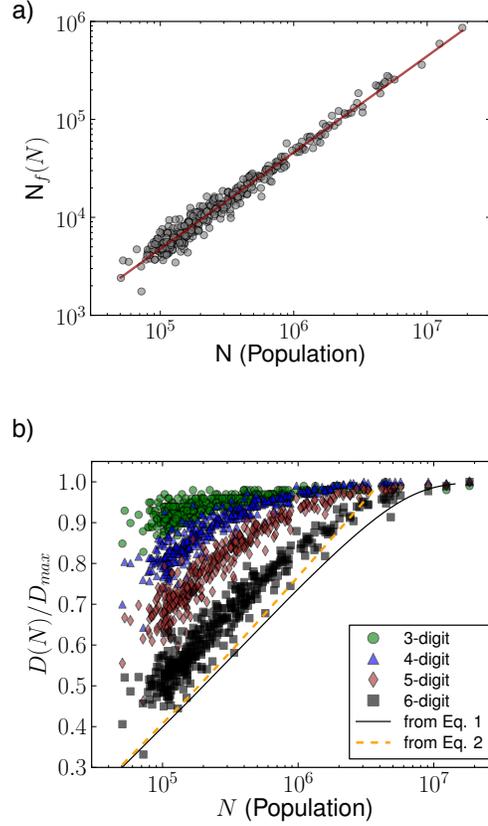


Figure 1: **The total number of establishments and their unique types as a function of city size.** (A) The total number of firms N_f scales linearly with city size: $N_f \sim N^\alpha$ where $\alpha = 0.98 \pm 0.02$ with $R^2 = 0.97$. (B) The number of distinct business types D normalized by its maximum value D_{max} at various levels of classification, r , based on the NAICS scheme, from the lowest resolution (3-digit) to the highest (6-digit) denoted by green circles, blue triangles, red diamonds and black squares, respectively (corresponding values of D_{max} are 317, 722 and 1160). All values are scaled by the corresponding size of the classification scheme at that resolution, D_{max} , such that all values fall in between 0 and 1. Note that $D(N)$ behaves very similarly to data of Japan (16). The black solid line and orange dashes are the predictions from Eq. 1. with and without ϕ .

model quantitatively predicts how individual business types systematically change rank with city size, shedding light on processes of economic differentiation with scale.

We focus on the frequency distribution of business types (the number of “species”) and first ask how this varies across cities (the “ecosystem”). We identify our unit of analysis as the *establishment*, which is defined as a single physical location where business is conducted, so that, for example, individual stores of a national chain would be counted separately. Establishments are nowadays seen as fundamental units of economic analysis because innovation, wealth generation, entrepreneurship and job creation all manifest themselves through the formation and growth of workplaces (18). We explore a unique dataset, the National Establishment Time-Series, a proprietary longitudinal database built by Walls & Associates, to capture economic life at an extraordinarily fine-grained level (19,20). This dataset includes records of nearly the entire set of establishments (work places) in US urban areas (over 32 million) each of which is classified according to the North American Industry Classification System (NAICS). We aggregate such information into the standard definition of functional cities: the 366 Metropolitan Statistical Areas (MSAs), which are defined by the census bureau as unified labor markets centered on a single large city wielding substantial influence over its surrounding region (21). These MSAs account for over 90% of the economic output of the US and almost 85% of its population (the lower limit on MSA size is $\sim 50,000$).

Results

The data reveals that the total number of establishments, N_f , in each MSA is linearly proportional to its population size, N : that is,

$$N_f \approx \eta N \tag{1}$$

with the proportionality constant $\eta \simeq 21.6^{-1}$ (Fig 1A). Thus, on average, there are about 22 people per establishment in a city, regardless of its size. Combined with the fact that the number of employees, N_e , also scales linearly with N , the average size of establishments is also independent of population size with $N_e/N_f \simeq 11.9$. This remarkable constancy of the average employment rate and number of establishments across cities is puzzling when viewed in light of agglomeration effects, and *per capita* increases in productivity, wages, GDP, or patent production, with population size (14,15). Clearly, then, the increasing returns to scale characterizing the benefits of urbanization and increasing city size are not simply due to bigger cities having more establishments. This suggests that investigating the diverse composition of different economic types is an important ingredient for understanding how larger cities realize gains in their aggregate economies.

The simplest measure of the level of economic diversity counts distinct establishment types, $D(N)$, within a metropolitan area of population N (1). Figure 1B shows that $D(N)$ increases logarithmically with N , but eventually levels off. This saturation, which

is also observed in Japanese metropolitan areas (16), may be attributed to the failure of the finite NAICS classification scheme which for the largest cities cannot fully capture the true extent of economic diversity: When means of differentiation among work places is insufficient, saturation is inevitable. For example, the category *restaurants* in NAICS could be further deconstructed into different cuisines: Thai, Indonesia, Korean, Japanese, Chinese, etc. This effect can be traced out by utilizing the hierarchical structure of NAICS; see the Supporting Information (SI) text for the detailed methods. Fig. 1B (colored circles) shows D for different resolutions, r , normalized by maximum numbers $D_{\max}(r)$ for each level. The logarithmic increase remains intact while the saturation appears at larger and larger city sizes as the resolution becomes finer. This suggests that saturation is, in fact, an artifact of insufficient detail in the description of business types. We return to this point below.

A more insightful way of assessing economic diversity is to examine the constituent types of $D(N)$ for individual cities. The abundance of the hundred leading business types for a selection of cities is shown in Fig. 2A; see SI for more examples and further rank abundance figures. In New York, the most abundant business type is *offices of physicians*, followed by *offices of lawyers* and *restaurants*; Phoenix ranks *restaurants* first and *real estate* second (perhaps not surprising in a rapidly growing city); San Jose, which includes Silicon Valley, predictably ranks *computer programming* second only to *restaurants*. Indeed, the composition of economic activities in cities has its own distinctive characteristics reflecting the individuality of each city. *It is therefore all the more remarkable that, despite the unique admixture of business types for cities, the shape of these distributions is universal; so much so that, with a simple scale transformation, their rank abundances collapse to a single unique curve common to all cities* (Fig. 2B). Note that the curve is robust to changes in levels of wealth, density, and population size, which vary widely across the U.S; see, SI.

This universality can be derived from a sum rule for the total number of establishments as follows. Let $F_i(N)$ be the number of the i^{th} most abundant business type in a city of size N , as shown in Fig. 2A. When summed over all ranks, this must add up to the total number of establishments, $N_f(N)$: thus, $\sum_{i=1}^{D_{\max}} F_i(N) = N_f(N) = \eta N$. But, from Eq. (1) $N_f(N) = \eta N$, so introducing the per capita distribution, $f_i(N) \equiv F_i(N)/N$, this can be re-expressed as

$$\sum_{i=1}^{D_{\max}} f_i(N) = \eta \quad (2)$$

As N increases, any growing or diverging dependences of the $f_i(N)$ on N in Eq. (2) cannot be cancelled against each other because of their positive definiteness, whilst any decreasing dependence vanishes. Since η is a constant this implies that each f_i must itself become independent of N for sufficiently large N , predicting that the *per capita* frequency of business types must be the same for all cities. Note that the derivation is independent of the underlying dynamics. Fig. 2B verifies the predicted invariance of f_i across all cities.

The invariance remains valid when we treat the discrete rank, i , as a continuous variable, x , and correspondingly f_i as a continuous function, $f(x)$; see SI. The surprise in the data is that this predicted collapse to a single curve extends all the way down to relatively small cities (that is, up to relatively high ranks) mirroring a similar precocious scaling observed in urban metrics.

The universal form of this scaled rank-size distribution, $f(x)$, has three distinct regimes: for small $x (< x_0$, say), it is well described by a Zipfian power law with exponent γ , as shown in the inset of Fig. 2B; for larger $x (> x_0)$, it is approximately exponential; and finally, as x approaches the maximally allowed value for the total number of categories, D_{\max} , $f(x)$ drops off suddenly. To a very good approximation, these can be combined into a single analytic form:

$$f(x) = Ax^{-\gamma}e^{-x/x_0}\phi(x, D_{\max}). \quad (3)$$

The parameters are $\gamma \approx 0.49$ and $x_0 \approx 211$. The overall normalization, A , is not an independent parameter and is determined from the sum rule for $f(x)$, Eq. (2), which gives $A \approx f(1) \approx 0.0019$ in excellent agreement with the empirical data.

The function, $\phi(x, D_{\max})$, parametrizes the cut-off that is enforced by the finite resolution, D_{\max} . It satisfies two conditions: (i) it is only important when x approaches D_{\max} , and completely dominates when $x = D_{\max}$; (ii) its effect vanishes if the resolution is sufficiently fine. A simple phenomenological function that satisfies all of these conditions is $\phi(x, D_{\max}) = e^{[1-(1-x/D_{\max})^{-1}]}$; see, SI for generalizations. For comparison, we show fits to the data both with and without ϕ in Fig 1B and Fig 2B to illustrate the effect of ϕ .

Recall that $D_{\max}(r)$ is the maximum possible number of business categories that can appear in a city using a given classification scheme with resolution r , and therefore increases with increasing resolution; see SI. In the limit of the finest grained resolution, $r \rightarrow \infty$, we expect $D_{\max} \rightarrow \infty$, in which case the term ϕ becomes constant, turning off the saturation phenomenon. This effect is already visible in Fig. 1B, where saturation sets in at smaller city sizes for coarser resolution. This leads to the conclusion that the highest present level of resolution of the NAICS classification scheme may be insufficient to capture the actual business type diversity of the largest cities.

The cut-off, ϕ , is, therefore, attributed to the saturation observed in Fig 1B. This can be shown by deriving the diversity function $D(N)$ from the rank-size distribution $f(x)$ (including the cut-off ϕ). This construction is possible because $D(N)$ corresponds to the lowest possible rank, x , containing only a single establishment. In the continuum limit this gives $F[D(N)] = 1$, or $f[D(N)] = 1/N$, leading to

$$D(N) = f^{-1}(1/N) \quad (4)$$

Solving this numerically using Eq. (3) gives excellent agreement with data in Fig. 1B (the black solid line). When the classification resolution r is sufficiently fine, D_{\max} becomes large and ϕ becomes 1. In such a limit, the numerical solution gives no saturation in $D(N)$

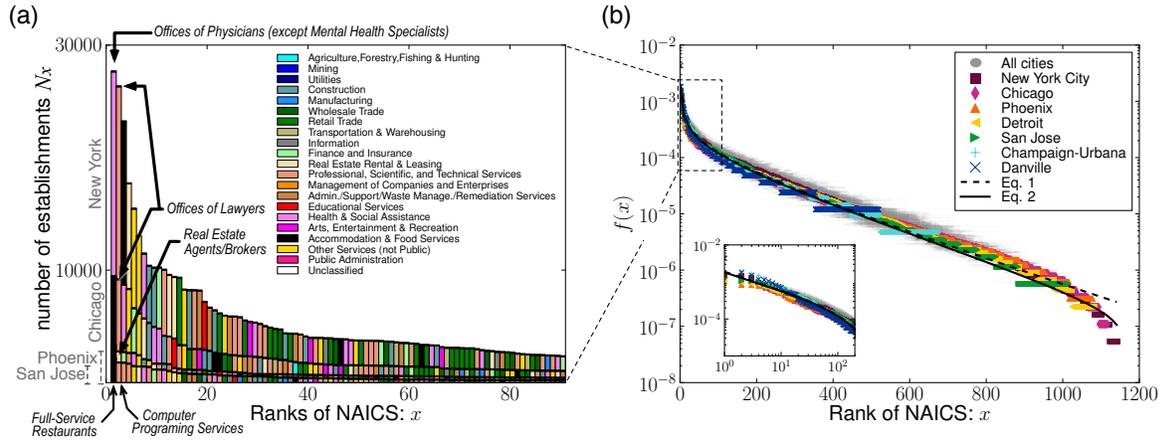


Figure 2: **Rank-abundance of establishment types** (A) The number of establishments at rank x ranging from 1 to 90 in descending order of their frequencies (from common to rare) for New York city, Chicago, Phoenix and San Jose. Establishment types are color coded by their classification at the 2-digit level. (B) Universal rank-abundance shape of the establishment type by dividing N_x by the population size of city in semi-log for all ranges. All metropolitan statistical areas are denoted by gray circles. Seven selected cities are denoted by various colors and shapes; New York city, Chicago, Phoenix, Detroit, San Jose, Champaign-Urbana, and Danville are, respectively, marked by red squares, pink diamonds, orange triangles, yellow left triangles, green right triangles, sky blue pluses, and blue crosses. The black dash line and the black solid line are fits predicted from Eq. 1 without and with ϕ respectively. The inset shows the first 200 types on a log-log plot showing an approximate Zipf-like power law behavior.

as the orange dashed curve in Fig 1B shows. The analytic solution for infinite resolution reduces to

$$D(N) \approx x_0 \ln N/\mathcal{N} \quad (5)$$

where \mathcal{N} depends weakly on N . This represents an open-ended ever-expanding diversity with population growth and confirms that the cut-off, ϕ , is associated with the saturation observed in $D(N)$.

We can understand the structure of $f(x)$ and $D(N)$ in the context of generalized preferential attachment or growing models (22, 23). The growing model is a widely accepted mechanism for generating rank-size distributions, whether for words, genes or cities. It is based on a stochastic growth process in which new elements of the system (business types in this case) are attributed a probability, α , of adding a new type, or adding to an existing type (24, 25). In the classic Simon-Yule model, for example, the attachment probability, α , of being an existing type is proportional to the existing frequency. As a result, the model exhibits a feed-back mechanism in which more-frequent types acquire new elements with higher probability than less-frequent types. Such a model provides a plausible mechanism for the observed Zipfian behavior in the high frequency regime of $f(x)$: namely, a power-law distribution $x^{-\gamma}$ with $\gamma = (1 - \alpha) < 1$ when $x < x_0$ as in Fig 2B. Since γ is estimated as 0.49 in the Eq. (3) from our empirical data, α is 0.51. This number suggests quite a rapid expansion so that small cities go through a stage of significant increase in business diversity as they grow: One in two new establishments entering the system, on average, spawns a new type. This fast pace of diversification contrasts to the pace in the exponential regime of $f(x)$ where x is larger than x_0 . The Yule-Simon model can be still used to explain this exponential regime by relaxing the condition for α : The probability is no longer constant but a function of system size N . In our case, the probability of each business being a new type decreases with city size. In fact, this slow diversification is observed in our Fig 1B where the number of business types $D(N)$ logarithmically increases, that is, slower than the linear. Because $\alpha(N) = dD(N)/dN$, the logarithmic function $D(N)$ in the Eq. (5) gives the $\alpha(N) \approx x_0/N$. This variation in $\alpha(N)$ can be used in the model of aggregate growth to predict the exponent $\gamma = \frac{\alpha(N)N}{D(N)} \frac{1}{1-\alpha(N)} = [(1 - x_0/\mathcal{N}) \ln(N)]^{-1}$, which vanishes for large N , and results in the observed exponential behavior in $f(x)$; see SI for details.

The empirical findings (Figs. 1-2) coupled with the predictions of the model described above suggest that all cities, as they grow, exhibit the same underlying dynamics in the development of their business ecology. Initially, small cities, with a limited portfolio of economic activities, need to create new functionalities at a fast pace. These basic activities constitute the economic core of every city, big and small. Later, as cities grow, the pace at which new functionalities are introduced slows down dramatically, but never completely ceases. Large cities, then, presumably rely primarily on combinatorial processes for developing new relationships among their many existing functionalities, which in turn

is the source of observed increases in economic productivity. This is a general feature of combinatorial growth process: Once the set of individual building blocks is large enough, their combination is sufficient to generate novelty even when the set itself expands slowly or not at all.

The universal distribution of frequencies does not, however, account for the entire developmental process of economic functionalities in cities. The stochastic Simon-Yule model, for example, does not predict what business compositions sit in what ranks. If, during growth, the introduction and success of each establishment were independent of business type (but dependent on frequencies), there would be no structure in how ranks are occupied. This is in clear disagreement with the following observation as well as with the pattern that “creative” activities and innovation concentrates disproportionately in large cities (5, 14).

The process by which specific business types assume different ranks in different cities may be particular to the ecology of specific places; or it may also be a property of scale. To distinguish between these two cases, we perform a multi-dimensional allometric scaling analysis of the number of specific establishments in each type. The super (or sub)-linearity of specific business types represents a systematic *per capita* increase (or decrease) of their abundances with city size. Fig. 3A shows an example: the number of *lawyers’ offices*, N_{lo} , scales as $N_{lo} \sim N^\beta$, with $\beta \approx 1.17$. That the exponent β is greater than one (super-linear) means that larger cities systematically have more lawyers *per capita*. Because lawyers’ offices typically appear at high frequencies ($x < x_0$) we can approximate $f(x)$ in Eq. (3) by its power law behavior and write, $N_{lo}/N = f(x_{lo}) \sim x_{lo}^{-\gamma}$ and thereby derive how the rank of lawyers changes with city size: $x_{lo} \sim N^{(1-\beta)/\gamma}$. This predicts $x_{lo} \propto N^{-0.4}$, which is in good agreement with the actual scaling shown in Fig. 3B. We can similarly predict how the ranks of low abundance business types scale. The rank shift can be expressed as:

$$x \sim \begin{cases} N^{\frac{1-\beta}{\gamma}} & \text{for small } x < x_0 \\ x_0(1 - \beta) \ln N & \text{for large } x > x_0 \end{cases} \quad (6)$$

Thus, business types whose abundances scale super-linearly with population size systematically increase their rankings, whereas those that are sub-linear systematically decrease, as expected.

Fig. 3C summarizes the values of scaling exponents for business sectors at the 2-digit-level (see *SI* for the histogram of exponents at the 6-digit level). Most primary sectors such as, *agriculture*, *mining*, and *utility* scale sub-linearly, predicting their systematic suppression, in relative terms, as cities get larger. On the other hand, informational and service businesses, such as *professional*, *scientific*, and *technical services* and *management of companies and enterprises* scale super-linearly, and are therefore predicted to increase disproportionately with city size, as observed. There are also sectors such as *restaurants*, for example, that do not change ranks. Note that sectors that deviate from linearity tend to be tradable industries that may be exchanged across cities (26). Because markets for these industries are not restricted to their immediate spatial location, comparative

advantages may generate agglomeration effects resulting from city size and/or to specific places.

Discussion

We have shown that the distribution of business in U.S. cities is characterized by a universal rank-size curve in which specific types predictably increase or decrease their relative rankings and frequencies as a function of city size. The results constitute a first general picture of the properties of the economic diversity of US cities measured in terms of business types. We show how statistics of business composition may be a result of a general mechanism of business creation based on existing frequencies and how the rate of introduction of new types slows down with city size. Our results, together with the theoretical framework presented, provides an important contribution to developing a theory of cities and urbanization that can encompass their dynamics, organization, and economic diversity. It provides quantitative support for the hypothesis advanced by classic central place and locational theories of cities (27, 28) that a general hierarchy of economic activities exists. These classic observations need to be interpreted not only in terms of the appearance of specific new sectors with city size, but also with the *disproportionate* growth of *certain* types versus others. Our work also gives quantitative support to business life-cycle theories (29, 30), where some types of business may be more prevalent in larger cities, but over time tend to move down the urban hierarchy as they mature and internalize more of their business model. We believe that the present results, together with further analyses of revenue, employment, temporal patterns, provide the foundation for a mechanistic understanding of how large cities realize greater economic productivity and how urbanization tends to promote nationwide economic growth (31).

Materials and Methods

The National Establishment Time-Series is a proprietary longitudinal database built by Walls & Associates. This dataset includes records of nearly the entire set of *establishments* in US urban areas (over 32 million) each of which is classified according to the North American Industry Classification System (NAICS). The establishment is a term to define a single physical location where the business is conducted (work place). Therefore one city can have multiple number of Starbucks establishments. We aggregate these establishments of various NAICS categories within a city. We use Metropolitan Statistical Area (MSA) as a functional city for our unit of analysis.

acknowledgments

We thank Ricardo Hausmann, Coco Krumme and Marcus Hamilton for discussions. This research is partially supported by the Rockefeller Foundation, the James S. McDonnell Foundation (grant no. 220020195), the National Science Foundation (grant no. 1312294), the John Templeton Foundation (grant no. 15705), the U.S. Department of Energy through the LANL/LDRD Program (contract no. DE-AC52-06NA25396), and by gifts from the Bryan J. and June B. Zwan Foundation and the Eugene and Clare Thaw Charitable Trust.

References

1. Whittaker RH (1972) Evolution and measurement of species diversity. *Taxon* 21:213-251.
2. Quigley JM (1998) Urban diversity and economic growth. *Journal of Economic Perspectives* 12:127-138.
3. Puga D, Duanton G (2000) Diversity and specialisation in cities: why, where and when does it matter? *Urban Studies* 37:533-555.
4. Glaeser E, Kallal HD, Scheinkman JA, Shleifer A (1992) Growth in cities. *J Polit Econ* 100:1126-52.
5. Henderson JV (1988) *Urban Development: Theory, Fact, and Illusion* (Oxford Univ. Press).
6. Henderson JV, Kuncoro A, Turner M (1995) Industrial development in cities. *Journal of Political Economy* 103:1067-1690.
7. Henderson JV (1974) The sizes and types of cities. *Am Econ Rev* 64:640-56.
8. Glaeser E (2011) *Triumph of the City: How Our Greatest Invention Makes Us Richer, Smarter, Greener, Healthier, and Happier* (Penguin Press).
9. Jacobs J (1984) *Cities and the Wealth of Nations: Principles of Economic Life* (Random House, New York).
10. Batty M (2008) The size, scale, and shape of cities. *Science* 319:769-771.
11. Pan, W. *et al.* Urban characteristics attributable to density-driven tie formation. *Nat. Commun.* 4:1961 doi: 10.1038/ncomms2961 (2013)
12. Bettencourt LMA, West GB, (2010) A unified theory of urban living. *Nature* 467:912-913.

13. Bettencourt LMA, Lobo J, Strumsky D, West GB (2010) Urban scaling and its deviations: revealing the structure of wealth, innovation and crime across cities *PLoS ONE* 5:e13541.
14. Bettencourt LMA, Lobo J, Helbing D, Kühnert C, West GB (2007) Growth, innovation, scaling, and the pace of life in cities. *Proc Natl Acad Sci USA* 104:7301-7306.
15. Bettencourt LMA (2013) The origins of scaling in cities. *Science* 340:1438-1441.
16. Mori T, Nishikimi K, Smith TE (2008) The number-average size rule: a new empirical relationship between industrial location and city size. *J Rg Science* 48:165-211.
17. Gomez-Lievano A, Youn H, Bettencourt LMA (2012) The Statistics of Urban Scaling and Their Connection to Zipf's Law. *PLoS ONE* 7(7): e40393. doi:10.1371/journal.pone.0040393
18. Davis SJ, Haltiwanger JC, Schuh S (1998) *Job Creation and Destruction* (The MIT Press)
19. Walls, Donald W., National Establishment Time-Series Database: Data Overview (November 2, 2007). 2007 Kauffman Symposium on Entrepreneurship and Innovation Data. Available at SSRN: <http://ssrn.com/abstract=1022962> or <http://dx.doi.org/10.2139/ssrn.1022962>
20. Bettencourt LMA, Lobo J, West GB, Youn H (2013) The hypothesis of urban scaling: formalization, implications and challenges, arXiv:1301.5919 [physics.soc-ph]
21. U.S. census bureau webpage, available at <http://www.census.gov/population/metro/>
22. Barabási A-L, Albert R (1999) Emergence of scaling in random networks. *Science* 286:509-512.
23. Newman MEJ (2005) Power laws, Pareto distributions and Zipf's law. *Contemporary Physics* 46:323-351.
24. Yule GU (1925) A mathematical theory of evolution, based on the conclusions of Dr. J. C. Willis, F.R.S. *Philos Trans R Soc London Ser B* 213:21-87.
25. Simon HA (1955) On a class of skew distribution functions. *Biometrika* 42:425-440.
26. Sachs DJ, Larrain FB (1994) *Macroeconomics in the Global Economy* (Prentice Hall Business Pub.).
27. Christaller W (1966) *Central Places in Southern Germany* (Prentice-Hall, Englewood Cliffs, NJ).

28. Lösch A (1954) *The Economics of Location* (Yale Univ. Press, New haven).
29. Klepper S (1997) Industry Life Cycles, *Industrial and Corporate Change* 6(1):145-182.
30. Duranton G, Puga D (2001) Nursery cities: urban diversity, process innovation, and the life cycle of products. *Ame Econ Rev* 91:1454-1477.
31. Bettencourt LMA, Samaniego H, Youn H (2012) Professional diversity and the productivity of cities. arXiv:1210.7335 [physics.soc-ph].

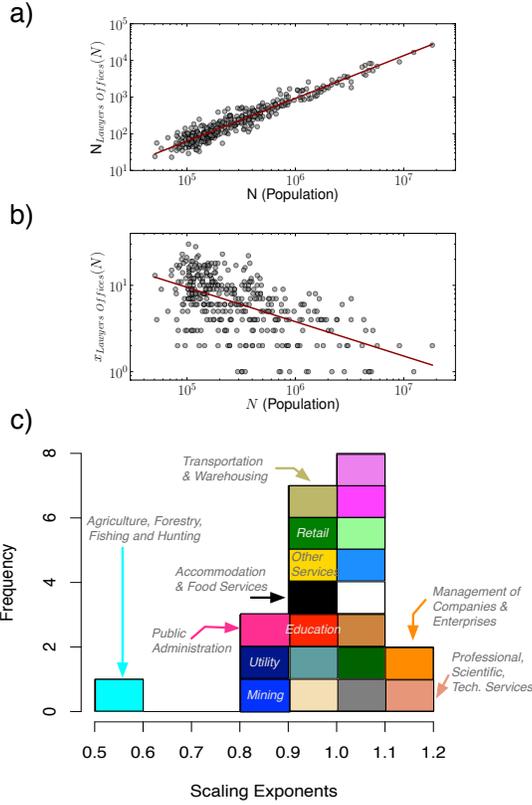


Figure 3: **Multi-dimensional allometric scaling of industry types.** (A) The number of lawyers' offices scales superlinearly with population size, $N_{lo} \sim N^{1.17 \pm 0.04}$ with $R^2 = 0.92$. (B) The rank of lawyers' offices increases with population size, expressing an increase in their relative abundance: $x_{lo} \sim N^{-0.4 \pm 0.06}$ with $R^2 = 0.32$. (C) Histogram of scaling exponents γ for all establishment types at the 2-digit level. While primary sectors disappear, managerial, professional, technical and scientific firms increase in relative abundance, helping to explain the increased productivity of larger cities despite the slow addition of new business types. Each scaling analysis is shown in SI.