Evidence for Allometric Scaling of Government Services in American Cities

Tim Gulden[1,2], Matt Koehler[1], Steve Scott[1], Zoe Henscheid[1]

[1] MITRE Corporation, McLean, VA
{mkoehler, slscott, zhenscheid}@mitre.org
[2] now at RAND Corporation, tgulden@rand.org
Approved for Public Release; Distribution Unlimited. Case Number 15-3312

Abstract. This paper seeks to identify scaling patterns in data relating to government services. It provides a general overview of the kinds of scaling that are associated with various aspects of government service in cities. It is hoped that these patterns will help to identify areas where social simulation methods might be productively applied while also providing quantitative targets against which such models can be validated. Using data from a variety of sources, the paper looks for evidence of allometric scaling in poverty and means tested programs, income tax revenue, healthcare services, air travel, and legal services. While none of these surveys provides the depth required to develop computational models of demand in these sectors, they can be taken together to develop insight into the areas of government enterprise where different types of scaling occur and where different approaches to computational modeling might be appropriate.

Introduction

Cities are well conceived as complex systems of remarkable diversity. The human social networks, institutions, and culture of a city interact in almost unfathomably complex ways with its natural and built environment, infrastructure, and surroundings to produce a living entity that can exist at a tremendous range of sizes and take on many different characters from one part of the world to another. At the same time, as with other complex systems, cities display certain regularities that tie them together as a common phenomenon. A deeper understanding of these regularities is needed in order to support systematic investigation into the patterns of development and behavior that produce them. A better understanding of how cities actually work, in turn, offers the prospect of better planning and more effective urban policy with regard to everything from infrastructure to social welfare.

The study of allometric scaling has proven useful in understanding some of the fundamental principles of biological systems. Many properties of biological organisms have been shown to exhibit cross-species regularities wherein a property of the organism changes with the size of the organism in a way that is well described by a power-law relationship. Metabolic rates scale as a ¾ power of the mass of an organism over 27 orders of magnitude ranging from cellular levels up to the largest organisms (West and Brown 2005). Mammalian heart rates scale as the mass of the organism to approximately the -½ power – thereby relating the inner workings of a shrew and a blue whale. Over the past decade or so, work in complex systems theory has cast considerable light on the geometric and biological principles that underlie these regularities in Biology.

Analogies have long been made between urban and biological systems and there is almost certainly something apt about them. We often speak of the health of a city, and of its growth and development. However, when we look for the many allometric scaling principles from biology in urban systems, we do not find them. We do, however, find a different set of scaling principles. The analogy between biological and urban systems breaks down at some point (as analogies will do) but points to the possibility of uncovering a similar set of principles that underlie the functioning and development of cities.

Indeed, considerable progress has been made in this regard over the past few years (Bettencourt 2013). In particular, urban infrastructure has been shown to scale sub-linearly, with larger cities requiring less infrastructure per resident than smaller ones do. Various aspects of economic activity, in contrast, have been shown to scale super-linearly, reflecting higher output and income per capita in larger cities than in smaller ones. It has been noted that super-linear scaling is not something that is observed in the biological realm.

Whereas the sub-linear scaling of biological systems can be explained largely as a function of the geometry of organisms and their subsystems, and progress has also been made in explaining sub-linear infrastructure scaling in cities with similar approaches, the super-linear scaling of urban productivity would seem to be based fundamentally in the increased interactions of urban residents (Bettencourt 2013), something unique to social systems. This makes urban scaling an ideal target for the tools of computational social science – particularly the use of agent-based modeling to gain understanding into why these scaling phenomena are observed. The scaling phenomena, in turn, provide a novel validation target for agent-based models purporting to capture urban dynamics.

This paper seeks to identify scaling patterns in data relating to government services. It provides a general overview of the kinds of scaling that are associated with various aspects of government service in cities. It is hoped that these patterns will help to identify areas where social simulation methods might be productively applied while also providing quantitative targets against which such models can be validated.

Using data from a variety of sources, this paper looks for evidence of allometric scaling in poverty and means tested programs, income tax revenue, healthcare services, air travel, and legal services. While none of these surveys provides the depth required to develop computational models of demand in these sectors, they can be taken together to develop insight into the areas of government enterprise where different types of scaling occur and where different approaches to computational modeling might be appropriate.

Methods

This study takes the form of a data survey, pulling data from diverse sources and analyzing it within a common framework to identify scaling relationships. Areas covered are poverty and means-tested programs, income tax revenue, government subsidized healthcare, air travel, and legal services. These subjects are chosen to provide a cross-section of government services where appropriate data could be located without undue effort.

Emphasis has been placed on measuring the size of the problem or service that the government is trying to address, rather than on specific government spending in a given place. The indicators chosen, therefore, are not always direct measurements of government expenditure. We focus, instead, on reliable indicators that are closely correlated with issues that are associated with government services of various sorts.

We examine the linearity of relationships by plotting the various indicators against urban population size in log-log coordinates. While there are many ways to determine the goodness of fit of a power law distribution to discrete data (see for example, Clauset, Shalizi, and Newman 2009), a good linear fit in log-log coordinates may indicate a power-law relationship with an exponent equal to the fitted slope and is the traditional method used in allometric scaling work (Bettencourt et al. 2007). Given the exploratory nature of this work, it was this, "least squares," method that was chosen for use. A slope of unity indicates a linear relationship, which can also be thought of as a power-law with an exponent of 1. A slope less than one indicates a power-law exponent less than one and a sub-linear relationship. The indicator in question increases with city size at a rate that is less than proportional to population. A slope greater than one indicates a super-linear relationship, where the indicator increases with population at a rate that is more than directly proportional to population.

Except where indicated, we use the Core Based Statistical Area (CBSA) definitions developed by the US Office of Management and Budget based on the 2010 US Census (US Census Bureau 2010a). These boundaries supplant the older Metropolitan Statistical Area definitions. They are commonly referred to as "Metropolitan Areas" (50,000 people or more) and "Micropolitan Areas" (10,000 to 49,999). These areas are defined largely in terms of commuting patterns and are designed to capture an urban area as a coherent economic entity. CBSAs are collections of whole counties; thus making possible precise comparisons to data that are released at the county level.

Results & Analysis

Poverty and Means Tested Programs

Many government services are aimed at assisting people with low income. When such programs require beneficiaries to demonstrate financial need, they are collectively labeled "means-tested" programs. This section of the study examines a variety of means-tested programs to develop a sense of how poverty and poverty assistance scale with city size.

We begin by plotting the number of people below the poverty line against total population for each CBSA (Figure 1). We take these data from two sources: both the Small Area Income and Poverty Estimates (SAIPE)(US Census Bureau 2012c), and the American Community Survey 5-year estimates for 2008-2012 (US Department of Housing and Urban Development 2013).

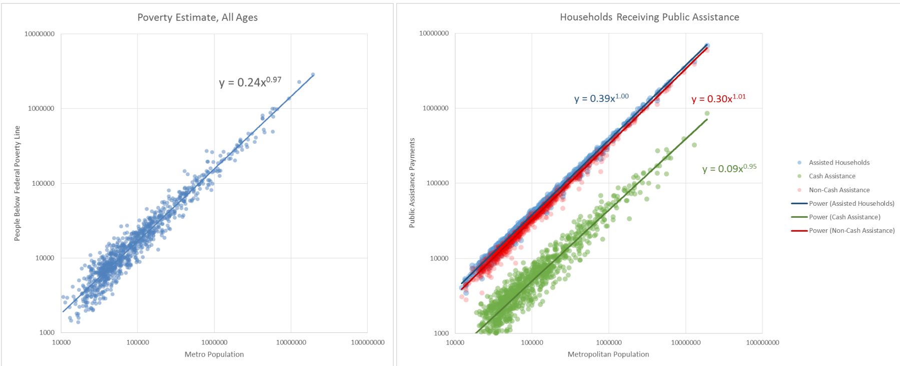

Fig. 1. People with income below the federal poverty level vs. city size (left), households receiving public assistance (right).

Across the full size range, we observe that about 15% of US city dwellers live below the federal poverty level. We find that the number of people living below the poverty line is remarkably linear with respect to city size, displaying a power-law exponent very near unity. While there is variation from city to city, this variation is not significantly correlated with city size.

The amenity draws of urban living, as well as the larger labor market would appear to precisely offset these higher costs to produce a rather stable percentage of people living on poverty incomes. It should be noted that the federal poverty level is weighted for family size, but not for geographic location. The same income that qualifies as poor in New York City or San Fransisco also qualifies as poor in a rural area with very low housing costs. It is thus remarkable that considerable housing and other expenses that increase the cost of living in large cities are not reflected in higher or lower poverty rates. The social and economic dynamics that undergird this stability are likely to be a productive target for social and economic modeling efforts.

When we examine the number of federally subsidized housing units using the HUD Picture of Subsidized Households 2013 dataset (US Department of Housing and Urban Development 2013), we find evidence for the same strongly linear scaling in housing subsidies that we found for poverty statistics. While at first the relationship appears marginally superlinear, we observe a collection of outliers in a handful of Metro- and Micropolitan areas with populations below 100,000. These cities report fewer than 100 federally subsidized units. It is not clear whether this is a problem with reporting, or whether there is some “critical mass” below which some cities decide to largely forgo federal housing assistance. If these cities with fewer than 100 units are excluded from the dataset, then the power-law exponent drops from

1.03 to 0.99, which is indistinguishable from unity. In either case, the difference from an exponent of one is extremely small and the relationship between city size and its number of federally subsidized housing units scales in a linear fashion.

While we find significant variation between cities with respect to both poverty and housing subsidy, we find that larger cities are no more or less poor than smaller cities and, in a related finding, that larger cities receive no more or less housing assistance than smaller ones.

Overall, it appears that there are neither synergies nor efficiencies with regard to the lower income parts of the economy in cities. While cities are often conceived as huge collections of poor people, we find little support for that framing: urban residents are no more likely to be poor in a large city than they are in a small city. This may be because subsistence of lower income workers is roughly proportional to the population: a given number of people need only so many restaurants, have only so many houses to clean, yards to maintain, etc. People migrate in and out of the city to occupy these economic niches and it appears that the urban economy creates these niches in a way that is directly proportional to population.

Income Tax Revenue

We now shift from poverty-related statistics to an examination of the full spectrum of income levels by examining data from the Internal Revenue Service to determine how tax revenue scales with city size. Our base file for this analysis is the IRS Statistics of Income (SOI) file for 2011 (Internal Revenue Service 2011). Because these data are complied by zip code some reallocation was required to allocate them to CBSA city boundaries that follow county boundaries. We do this by performing a GIS overlay between CBSA boundaries and zip code (ZCTA) boundaries (US Census Bureau 2010b). Zip codes having half or more of their land area within the CBSA are counted as part of the CBSA, while zip codes having less than half of their area in the CBSA are excluded from it. While this approach is not as precise as a population-weighted allocation, it is simple, fast, and reasonably unbiased. We begin by plotting aggregate income and aggregate income tax paid against city population (Figure 2).

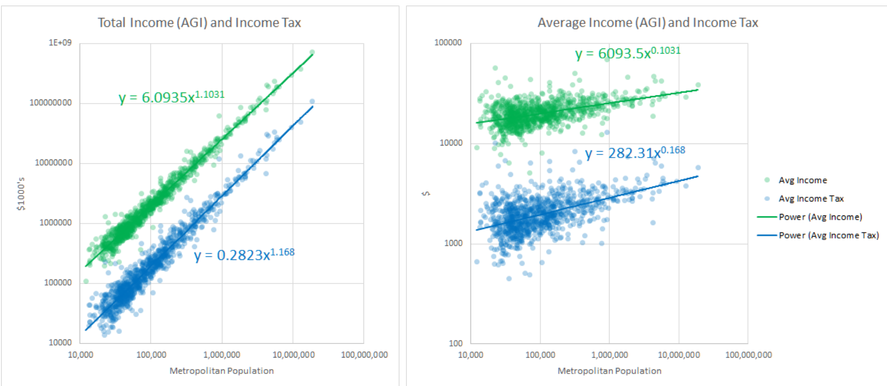

Fig. 2. Total adjusted gross income (AGI), and Average AGI vs. city size.

Both income, and income tax increase more quickly than city size, demonstrating allometric scaling exponents of 1.10 and 1.17 respectively (left figure). This plot makes more intuitive sense if we plot average income against city size (right figure). This has the effect of subtracting one from the slope. Average income in a city of 10,000 is just under \$43,000/y, while average income in a city of 10,000,000 is just over \$70,000/y.

It has been established that GDP scales with an exponent of about 1.15 [Bettencourt or West] – a bit faster than we observe here. This may be because this analysis includes only personal income, excluding corporate income and other forms of value added.

Note that while per capita income is higher in larger cities, per capita tax is higher still. This is presumably due to the progressive nature of the US income tax system. Further examination of the data indicates that the average tax rate in a city of 10,000 is about 9%, while the average tax rate in a city of 10,000,000 is 14%.

We next break these incomes out by Adjusted Gross Income (AGI) Class. The data include six categories: tax returns reporting income under \$25K, \$25K to \$50K, \$50K to \$75K, \$75K to \$100K, \$100K to \$200K, and over \$200K.

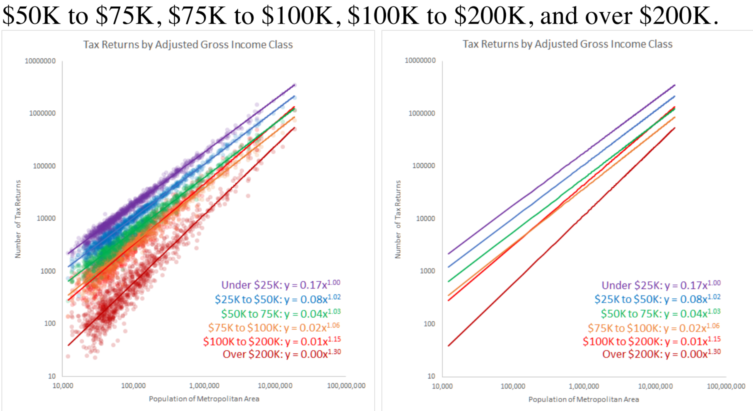

Fig. 3. Returns with income falling into six different AGI classes. Data points and trend lines (left). Trend lines only (right).

We find overall that the scaling exponent rises as city size increases; however, there is a major discontinuity beginning at \$100K/y. The four classes below that level have similar exponents, ranging from 1.00 to 1.06 – all quite close to unity. The two higher income classes, however, show exponents of 1.15 and 1.30 respectively.

This finding is probably related to an earlier observation that income in the US (and many other countries) displays a composite distribution, with the vast majority of incomes falling into an exponential distribution, but the top 1% to 3% of incomes follow a Pareto power-law distribution (Drăgulescu and Yakovenko 2001). The cutoff between these two distributions is thought to be around \$125K. The current result reflects not only this discontinuity in the income distribution, but also considerable clustering of higher incomes in larger cities.

Whereas the percentage of lower and middle incomes remains nearly constant as we vary city size, the (small) percentage of upper income households increases dramatically. For example, the percentage of tax returns showing more than \$200K in the average US city of 10,000 is less than one half of one percent. In a city of 10,000,000 (New York), that percentage grows to nearly six percent.

Further exploration of the dataset indicates that the difference in scaling exponents of incomes is not due to a difference in the scaling of labor income vs. wealth-based income (interest, dividends, capital gains, rents, retirements withdrawals, etc.), with both types of income displaying similar scaling properties in all income classes. Seen in this light, large cities can be seen not so much as concentrations of poor people as they are concentrations of wealthy people.

Healthcare: Medicare and Medicaid

We next look at how government health programs scale with city size. We will compare data from two sources, the CMS Medicare Geographic Variation Public Use File for 2013 (Centers for Medicare and Medicaid Services 2013) and the US Census American Community Survey (US Census Bureau 2012a), both reaggregated to the CBSA level.

First we plot both overall Medicare beneficiaries and Medicare beneficiaries with diabetes (Figure 4).

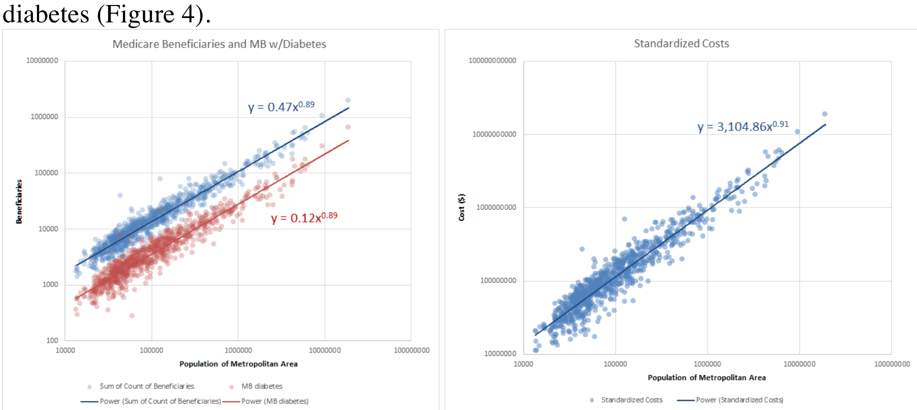

Fig. 4. Medicare beneficiaries, beneficiaries with diabetes, and standardized Medicare costs vs. city size.

These data show the number of Medicare beneficiaries to demonstrate strongly sublinear scaling. The same holds for standardized Medicare costs. This would seem to indicate that the costs incurred by Medicare beneficiaries are similar in cities of different sizes, but that Medicare is used less in larger cities. The data show no major scaling effect for Medicare eligibility – with about the same percentage of the population eligible in cities across the range of sizes.

This sublinear effect seems to be driven by declining participation as city size increases. When the number of beneficiaries is expressed as a percentage of the eligible population (Figure 5, left), we see a statistically weak, but significant fall off from 77% of eligible people participating in the average city of 10,000, to 42% of eligible people participating in a city of 10,000,000. Note also, that the data for some smaller cities show significantly more people participating than are shown to be eligible. When the same data are expressed as a percentage of total population, we see a decline from 18% in the average city of 10,000 to about 8% of the population in a city of 10,000,000.

We can check these somewhat puzzling findings by comparing to self reported Medicare status from the US Census American Community Survey (ACS), using the 2012 3-year estimates. The two sources are compared in Figure 5 (right).

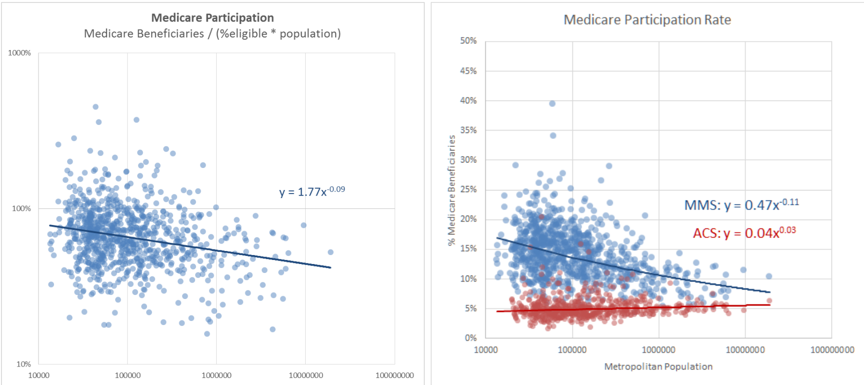

Fig. 5. Right: Medicare participation as a fraction of those eligible according to Center for Medicare and Medicaid Services. Left: Medicare participation as a percentage of total population reported by the Center for Medicare and Medicaid Services (blue), and American Community Survey (red).

As reported by the ACS, Medicare use actually increases slightly with increased city size, rising from about 4% of the population in a city of 10,000 to about 6% of the population in a city of 10,000,000. Medicaid, in keeping with the earlier findings on poverty, is almost precisely flat. These numbers would appear to be reporting on the same indicator – Medicare beneficiaries, but produce strikingly different results. Resolving this discrepancy is beyond the scope of this study.

Travel Demand

Government plays many important roles in transportation, from road building to port security and air traffic control. We examine one aspect of this by looking at overall air traffic originating in cities of different sizes. We analyzed the Airline Origin and Destination Survey (DB1B) from the Bureau of Transportation Statistics (Bureau of Transportation Statistics 2014) to identify the airport market origins of a 10% sample of itineraries on commercial carriers in the first quarter of 2014. This involved approximately six million records of airline tickets. The nature of airport markets makes comparison to metropolitan populations difficult, as people living in smaller cities often drive considerable distances to use airports in larger cities. Comparing the air travel data to metropolitan populations required a certain amount of hand matching and subjective judgment. This matching can likely be improved with additional work; however, the results stemming from the initial effort are suggestive. We find that the relationship between city size and air travel demand is complicated, displaying different scaling behavior in different classes of cities.

Examining the full set of metropolitan areas for which airport markets can be established, we see what appears at first to be evidence of superlinear scaling, displaying a scaling exponent of 1.12 (Figure 6).

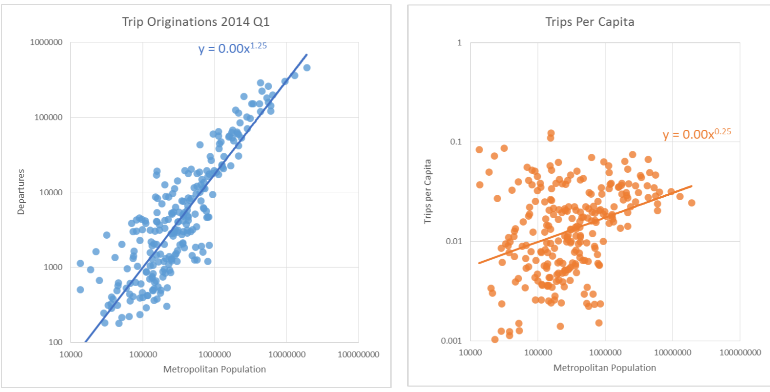

Fig. 6. Commercial airline passenger trip departures vs. city size for all US airline markets.

On closer examination of individual cities, however, it appears that this is an artifact of spillover between smaller cities and larger ones. While a small city may have a minor airport, it will tend to attract few flyers not so much because its citizens fly less, but rather because commercial air flights are much more numerous and less expensive when they originate from larger cities that are within a few hours drive. This has the effect of substantially understating travel demand from small cities and somewhat overstating travel demand from larger cities.

When we restrict the analysis to metropolitan areas with 100,000 people or more (top 75 markets), we see a slightly sub-linear trend – though the small sample makes the trend less certain, and it is hard to rule out that this is a linear trend (Figure 7).

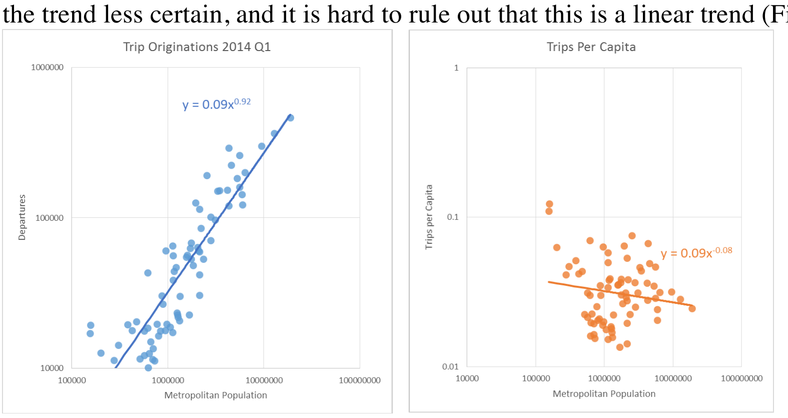

Fig. 7. Commercial airline passenger trip departures vs. city size for US airline markets in metropolitan areas with over 100,000 residents.

The deviation from linear in this mid-sized sample is produced largely by the top few cities, which display a lower slope. If we isolate that group by pulling out the top 20 metropolitan areas in the US, we see a substantially different trend, with a scaling exponent of only 0.56 (Figure 8). This indicates a substantial drop in trips per capita as city size increases among the top 20 cities in the US. It should be noted here that this type of analysis is better suited to larger sample sizes and these findings should probably be taken more as fodder for hypothesis formation than as established empirical findings.

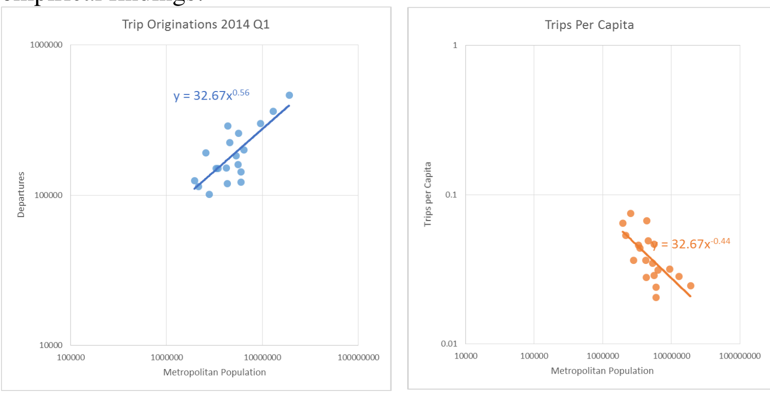

Fig. 8. Commercial airline passenger trip departures vs. city size for 20 largest US airline markets.

These findings suggest that the higher incomes and greater business activity in larger cities do not translate to more air travel – particularly for very large cities. It is possible that the tremendous diversity of the economies of such cities makes air travel to other cities relatively less important as they contain within themselves a greater fraction of the goods and services that their citizens require in order to conduct business. A more through analysis of this issue would benefit from an effort to separate business and leisure travel. The greater wealth of larger cities might be expected to produce a greater demand for leisure travel – if that is the case, the implied reduction of demand for business travel might be even greater than these preliminary statistics indicate.

Legal Services

Legal services touch on a host of government programs, ranging from the provision of courts, to control of crime, and the regulation of commerce. While considerable work is possible in unpacking the different parts of legal services, this study will examine the industry as a whole as a natural starting point.

We analyze the US Census 2012 Business Patterns dataset (US Census Bureau 2012b), taking employment and payroll in the legal services industry as proxies for the size of the legal services industry overall. We further assume, at this stage in the analysis, that the level of government involvement is a relatively constant proportion of the industry. This merits further investigation at a later date.

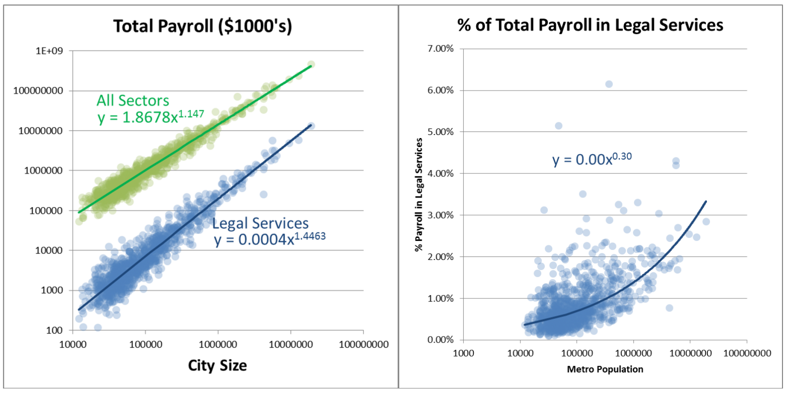

Fig. 9. Total payroll for all sectors and for legal services (left). Legal services payroll as a percentage of all sector payroll (right).

We first note that total payroll – aggregated across all sectors – grows faster than city size, demonstrating a scaling exponent of 1.15. This finding is precisely in line with earlier findings by Bettencourt, et al. Legal services payroll, however, can be seen to scale much more quickly, with an exponent of 1.45 (Figure 9). We can gain some intuition into what this scaling factor indicates by separating the number of workers in the industry from the average salary in the industry. The extremely rapid scaling of legal services payroll is driven by superlinear scaling in both the number of workers in the industry, and superlinear scaling in the average pay per worker.

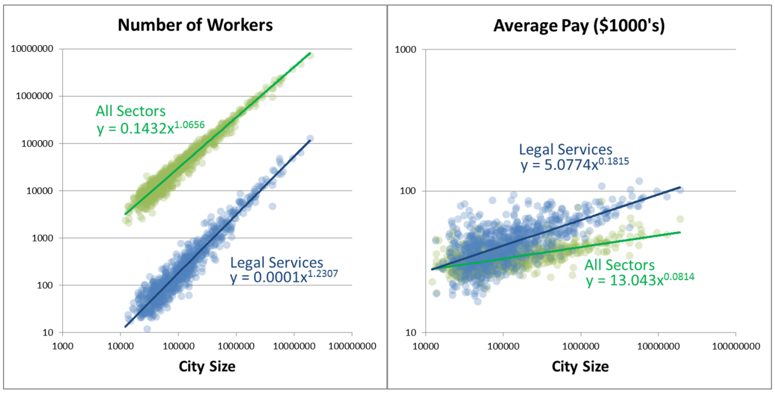

Fig. 10. Number of workers in all sectors and in legal services sector (left). Average pay in all sectors and legal services sector (right).

While labor force participation is somewhat higher in larger cities, demonstrating a scaling exponent of 1.07, we find that participation in the legal services industry scales much faster, with an exponent of 1.23 (Figure 10, left).

A further effect is seen when examining average pay. In all sectors taken together, we see average pay increasing with an exponent of 0.08 (roughly consistent with the

IRS data discussed earlier). Note that this is equivalent to total payroll increasing with an exponent of 1.08. Average pay in the legal services industry, scales with an exponent of 0.18. In more intuitive terms, this means that the average legal services worker in a city of 10,000 can expect about the same pay as the average worker – about $29,000 per year. In the largest US cities, however, a major gap has opened; with the average worker commanding $50,000 and the average worker in legal services making something closer to $100,000.

These data indicate that legal services play a more valuable role in the economies of larger cities. As we compare cities of various sizes, legal services command a larger percentage of the workforce, and an even larger percentage of the payroll.

Looking at legal services in the context of government services, it should be noted that this is an extremely diverse industry, and the government plays very different roles in its subsectors. For example, legal services related to business and real estate are largely private sector activities, whereas criminal justice involves substantial government involvement. While there has been some work looking at scaling of criminal activity (Alves et al. 2013) much work remains to be done to analyze easily available data in the service of understanding how these various sectors break down. Additional exploration of US Census Bureau County Business Patterns data, as well as FBI Uniform Crime Reports are likely to be quite productive in understanding the various parts of the legal system as they relate both to cities of different sizes, and to government services.

Discussion

We have found that the scaling of indicators relevant to government services is consistent with previous findings on cities, productivity and infrastructure. While cities are often conceived as great concentrations of poor people, we did not find evidence of this. Rather, cities might be thought of as great concentrations of wealthy people. While the percentage of people using means tested-programs is uncorrelated with city size, we find that tax revenues and workers in high value-added services increase markedly with city size. All else being equal, we would expect this increase in wealth to create a greater demand for air travel, but we do not find this to be the case. Rather we find some evidence that the largest cities travel less on a per capita basis. This is likely due to the fact that business people in large cities have less need to travel. The ultra-diverse economies of the largest cities contain many things that a businessperson might travel for: client firm headquarters, international patent attorneys, management consultants, etc. Air travel therefore works much like other kinds of infrastructure, with higher density producing greater efficiency that shows up both as higher output and less need for infrastructure.

This study has laid the foundation required to ask a number of additional questions. Is the super-linearity in higher incomes due to greater inequality in larger cities, or simply due to an upward shift in the median income? Would a more multi-dimensional examination of travel patterns support the hypothesis that people in very large cities make less use of air travel? Similarly, would we see a wealth effect producing more leisure travel, but less business travel originating from large metropolitan regions? Why do we such a large discrepancy between administrative records-based and survey-based counts of Medicare beneficiaries? Is this simply a matter of differing definitions, or is one of the sources consulted (CMS and Census) producing numbers that are systematically biased? These and other questions have been raised and could be pursued using straightforward techniques based on the findings of this study to date.

References

Alves, Luiz G. A., Haroldo V. Ribeiro, Ervin K. Lenzi, and Renio S. Mendes. 2013. "Distance to the Scaling Law: A Useful Approach for Unveiling Relationships between Crime and Urban Metrics." *PLoS ONE* 8 (8): e69580. doi:10.1371/journal.pone.0069580.

Bettencourt, Luís M. A. 2013. "The Origins of Scaling in Cities." *Science* 340 (6139): 1438–41. doi:10.1126/science.1235823.

Bettencourt, Luís M. A., José Lobo, Dirk Helbing, Christian Kühnert, and Geoffrey B. West. 2007. "Growth, Innovation, Scaling, and the Pace of Life in Cities." *Proceedings of the National Academy of Sciences* 104 (17): 7301–6. doi:10.1073/pnas.0610172104.

Bureau of Transportation Statistics. 2014. "Origin and Destination Survey DB1B." http://www.transtats.bts.gov/Tables.asp?DB_Short_Name=Origin%20and%20Destination%20Survey.

Centers for Medicare and Medicaid Services. 2013. "Medicare and Medicaid Public Use File." https://www.cms.gov/Research-Statistics-Data-and-Systems/Statistics-Trends-and-Reports/Medicare-Geographic-Variation/GV_PUF.html.

Clauset, A., C. Shalizi, and M. Newman. 2009. "Power-Law Distributions in Empirical Data." *SIAM Review* 51 (4): 661–703. doi:10.1137/070710111.

Drăgulescu, Adrian, and Victor M. Yakovenko. 2001. "Exponential and Power-Law Probability Distributions of Wealth and Income in the United Kingdom and the United States." *Physica A: Statistical Mechanics and Its Applications*, Application of Physics in Economic Modelling, 299 (1–2): 213–21. doi:10.1016/S0378-4371(01)00298-9.

Internal Revenue Service. 2011. "Statistics of Income." http://www.irs.gov/uac/SOI-Tax-Stats-Statistics-of-Income.

US Census Bureau. 2010a. "American FactFinder." http://factfinder.census.gov/faces/nav/jsf/pages/index.xhtml.

———. 2010b. "Zip Code Tabulation Areas." https://www.census.gov/geo/reference/zctas.html.

———. 2012a. "American Community Survey." https://www.census.gov/programs-surveys/acs/.

———. 2012b. "County Business Patterns." http://www.census.gov/econ/cbp/.

———. 2012c. "Small Area Income and Poverty Estimates." http://www.census.gov/did/www/saipe/publications/index.html.

US Department of Housing and Urban Development. 2013. “Picture of Subsidized Households.” http://www.huduser.org/portal/datasets/picture/yearlydata.html.
West, Geoffrey B., and James H. Brown. 2005. “The Origin of Allometric Scaling Laws in Biology from Genomes to Ecosystems: Towards a Quantitative Unifying Theory of Biological Structure and Organization.” *Journal of Experimental Biology* 208 (9): 1575–92. doi:10.1242/jeb.01589.