Yannis Menelaos Ioannides and
Henry G. Overman

Spatial evolution of the US urban system

Originally published in Journal of economic geography, 4 (2). pp. 131-156 © 2004 Oxford University Press.

Spatial Evolution of the U.S. Urban System

Yannis M. Ioannides[*†]
Tufts University

Henry G. Overman[*‡]
London School of Economics

August 2003

ABSTRACT: We examine spatial features of the evolution of the US urban system using US Census data for 1900 – 1990 with non-parametric kernel estimation techniques that accommodate the complexity of the urban system. We consider spatial features of the location of cities and city outcomes in terms of population and wages. Our results suggest a number of interesting puzzles. In particular, we find that city location is essentially a random process and that interactions between cities do not help determine the size of a city. Both of these findings contradict our theoretical priors about the role of geography (physical and economic) in determining city outcomes. More detailed study suggests some solutions that allow us to restore a role for geography but a number of puzzles remain.

Key words: Urban growth, spatial evolution, economic geography, kernel estimation.
JEL classification: R00, C14

[*]Ioannides acknowledges generous research support from John D. and Catherine T. MacArthur Foundation and the National Science Foundation. We thank two anonymous referees and Richard Arnott, the editor for comments that have substantially improved the paper. We thank Danny T. Quah for giving us access to his tSrF package; Rob Axtell, Duncan Black, Gilles Duranton, Ed Glaeser, Anna Hardman, Michael Murray, Diego Puga, Atsushi Yoshida and other seminar participants at Brookings, Harvard, Kyoto University, the NBER Summer 2001 Institute, and the University of Toronto for their comments; Nick Gill and Tim Harris for research assistance and Duncan Black and Linda H. Dobkins for access to data.

[†]Department of Economics, Tufts University, Medford, Massachusetts 02155, U.S.A. e-mail: `yioannid@tufts.edu`

[‡]Department of Geography and Environment, London School of Economics, Houghton Street, London WC2A 2AE, United Kingdom (e-mail: `h.g.overman@lse.ac.uk`; website: `http://cep.lse.ac.uk/~overman`). Also affiliated with the Centre for Economic Policy Research, and the Centre for Economic Performance at the London School of Economics.

1. Introduction

This paper considers the spatial characteristics of the US urban system, along with several of its other attributes, as it evolved over the twentieth century. As any urban system develops some cities prosper while others decline. Further, as the system expands, we may see new metropolitan areas form while other declining cities die. That is, both the size and number of cities can change as the urban system evolves. At the same time we may also observe changes in the *geography* of the urban system. For example, changes in the location of different activities, in the location of larger and smaller cities or in the nature of the spatial relationships between cities.

Both the size distribution and the geography of the urban system have been the subject of considerable empirical investigations. Work on the size distribution has tended to centre on the question of whether Zipf's law or its deterministic equivalent, the rank size rule, holds for cities. Carroll (1982) and Cheshire (1999) provide an overview of the earlier literature while Gabaix and Ioannides (2003) consider more recent work. As the latter makes clear, work by Gabaix (1999) on the relationship between Gibrat's law and Zipf's law has provided an alternative basis for structuring empirical work. Recent work on the city size distribution does not just update the older literature but also provides new insights and methods for studying the issues. See, for example Eaton and Eckstein (1997), Dobkins and Ioannides (2000; 2001), Black and Henderson (2002), Overman and Ioannides (2001) and Ioannides and Overman (2003).

In contrast, our understanding of geographical features of the urban system is limited to either intra-metropolitan spatial structure or very specific features of inter-metropolitan geography, as emphasised in particular by central place theory a la Christaller and Losch. Thus, while recent advances in empirical methods have increased our understanding of the evolution of the city size distribution, spatial features of the urban system remain largely unexplored.[1] This paper seeks to redress this balance.

Such an undertaking is timely given that recent theoretical advances have highlighted the importance of spatial dimensions in understanding the evolution of urban systems. The theorists who have developed the New Economic Geography, including Masahisa Fujita, Paul Krugman and Anthony Venables [Fujita, Krugman and Venables (1999)], have added important new spatial

[1]Two papers do consider the issue. Dobkins and Ioannides (2001) examine the basic dynamics of spatial interactions among US cities and its impact on their populations focusing particularly on the entry of new cities. Black and Henderson (2002) consider the importance of both first and second nature geography in explaining the growth rates of cities. They find that both factors are important in explaining city growth.

insights to the established literature on systems of cities, represented most notably by the work of J. Vernon Henderson [Henderson (1974; 1988)]. The system of cities approach, as the latter has come to be commonly known, featured powerful models of intrametropolitan spatial structure, but neglected intermetropolitan spatial structure. Subsequently, intermetropolitan spatial structure played a key role in the new economic geography literature, starting with the work of Paul Krugman [Krugman (1991)].[2] Further, as shown by Fujita and Thisse (2002), the importance of spatial dimensions is not just restricted to the New Economic Geography. Rather, it is a general feature of recent theoretical advances in our understanding of the economics of agglomeration.

This recent theorising has formalised thinking about two fundamental features of any given location – the *first* and *second natures* – that determine the extent of development at that location [Krugman (1993)]. First nature features are those that are intrinsic to the physical site itself, independent of any development that may previously have occurred there. For example, locations on navigable rivers, with favourable climates have first nature features that might encourage development. The second nature features of a location are those that are dependent on the spatial interactions between economic agents. These second nature features might depend on previous development at the location (e.g. the availability of specialist suppliers) or on the spatial structure of the economic system more generally (e.g. the benefits of good access to a large market).

We had to resolve two key questions in undertaking our analysis of the spatial evolution of the US urban system. The first is how this theoretical work should help structure our analysis. Our conclusion was that, at the current stage of development, the precise implications of these models for growth in a system of cities are pretty fuzzy. Real life geography, the tendency for all cities to grow, the gradual convergence to some kind of equilibrium in the westward expansion of the country, the movement of population towards the sunbelt, and changes in the US urban system induced by a shift, over the period of study, in industrial structure away from manufacturing and towards services are all important features in the spatial evolution of the US urban system that have not yet been elaborated in the formal theory. Thus, in what follows, we seek to increase our understanding of first and second nature features of the US urban system without restricting our analysis to specific functional forms. Towards the end of the paper, we consider what this sort of analysis can tell us about recent theoretical work.

[2] Tabuchi (1998) attempts to synthesise the system of cities literature with the economic geography-based theories by incorporating intrametropolitan commuting costs as well as intermetropolitan transport costs.

A number of other authors have taken the opposite approach and attempted to estimate structural parameters of Krugman type models. Thomas (1996), Hanson (2000) and Combes and Lafourcade (2001) all undertake this type of analysis. While interesting, the extreme complexity of these analyses make it difficult to ascertain what we actually learn about the spatial features of the economies that these authors consider.

The second key question that we have to resolve related to the appropriate empirical tools to use. We have chosen to focus predominantly on non-parametric methods proposed by Quah (1993; 1996a,b; 1997; 1999). These methods have been widely applied in the growth literature for studying the evolution of income disparities across countries or regions. As we will see below they are ideally suited to capturing and simplifying the spatial features of the evolution of quite complex urban systems. In particular, they do not impose the sort of structure that parametric analyses rest on, namely fitting a representative city into the data.

The evidence that we consider falls in to two broad categories. The first relates to the location of cities. The second to the size and growth of cities. The paper is structured as follows. In section 2 we describe the data that we use. In section 3 we consider spatial features of the location of cities while in section 4 we turn from location to spatial elements of city outcomes in terms of population and wages. Finally, section 5 relates our findings to theory and concludes.

2. Data

There are a variety of ways to define cities empirically. In this paper we use contemporaneous Census Bureau definitions of metropolitan areas, whenever possible. From 1900 to 1950, we use metropolitan areas as they were defined by the 1950 census. For years before 1950, we use Bogue's reconstructions of what each metro area's population would have been with the metropolitan areas defined as they were in 1950 [Bogue (1953)]. From 1950 to 1980, we use the metropolitan area definitions that were in effect for those years. However, between 1980 and 1990, the Census Bureau redefined metropolitan areas. The effect of the redefinitions were that the largest U.S. cities took a huge jump in size, and several major cities were split into separate metro areas. While this might be appropriate for some uses of the data, it would introduce "artificial" differences in growth patterns for the 1980–1990 period. Therefore, we reconstructed the metro areas for 1990, based on the 1980 definitions, much as Bogue did earlier. We believe that this gives us the most consistent definitions of US cities (metropolitan areas) that we are likely to find.

The method raises a question as to which cities, as defined or reconstructed, should be included. In the years from 1950 to 1980, we use the Census Bureau's listing of metropolitan areas. Although the wording of the definitions of metropolitan areas has changed slightly over the years, the number 50,000 is minimum requirement for a core area within the metropolitan area. Therefore, we used 50,000 as the cutoff for including metropolitan areas as defined by Bogue prior to 1950. Adopting the same cutoff facilitates comparability with a fair amount of previous research.

Table 1 presents summary statistics for the data in the ten different time periods. The second column reports the US population while the third reports the urban population. The trend towards increasing urbanisation is clear, with the percentage of population classified as urban rising from 38% in 1900 to 77% in 1990. The nearly seven fold increase in urban population is not mirrored by comparable increases in city sizes. As we see from column four mean city size has little more than doubled over the period. The average city in 1990 was 2.2 times the average city size in 1900. From column five we see that median city size has grown slightly slower doubling in size between 1900 and 1990. Column six shows how we can reconcile the huge growth in urban population with the relatively small growth in average city size - the number of cities has almost tripled over the period from 112 in 1900 to 334 in 1990. While it is often difficult to deal econometrically with an increasing number of cities, it is clear that the entry of new cities is a key aspect of the evolution of the US urban system.

In addition to population we also have data on earnings in all cities in the sample for all years drawn from Census reports. Data on earnings are not available for non-urban areas, but population data are available for all counties (urban and non-urban) in all periods.[3] We use these county based data to consider some second nature features of the evolution of the urban system.

As we are interested in spatial features, we need to be able to specify the location of cities and counties. We take the centre of the city as the latitude and longitude given in the *1999 Times World Atlas*. For counties we take the latitudes and longitudes of the largest human settlement. We use this information when calculating distances between cities. The distance between cities is calculated as the geodesic distance between them on the basis of great circle distances.[4] We feel that this is the most appropriate distance metric for studying the spacing of cities as the urban system

[3]We are grateful to Duncan Black for sharing with us his county-based data. More details on the county data are available in Black and Henderson (2002).

[4]Great circle distances are calculated as follows. For any two locations A and B, we can calculate the angle formed by a ray joining the two points A and B and a ray joining A to the centre of the earth as follows:

$$\text{angle} = (\sin(\text{latA}) \times \sin(\text{latB})) + (\cos(\text{latA}) \times \cos(\text{latB}) \times \cos(\text{longA} - \text{longB}))\,,$$

evolves. Clearly some topographical features may make geodesic distances a poor approximation to economic distance for a few cities. However, the only feasible alternative measures using transport networks are very problematic when studying the evolution of the urban system over such a long time period. Clearly the location of the highway network is highly endogenous with respect to city location as construction begins in the middle of the sample period. The rail network circa 1900 might be useful as an exogenous measure of distances between the 112 cities extant at that time, but would provide no information for new cities, that entered from 1900 to 1990 and are nearly twice as many as those extant in 1900.

We also assign cities to one of nine regions according to the Census Bureau division of the country. Kim (1997) argues that the census regions are likely to serve well as economic regions, at least over the first half of the century.[5]

Finally, we use the date of settlement for each city, as obtained by Dobkins and Ioannides (2001). At first glance, one would suppose that the east to west settlement of the country would determine settlement dates, but we find early settlement dates in the west and late ones along the east coast. Settlement here refers to historical references to settlement in a location, and our variable is compiled by sifting through historical records. In a number of cases, the dates are references to military forts. We use those dates because often the site of the fort determined the site of the city that grew up nearby. The earliest date is that of Jacksonville, Florida, in 1564, and the latest is Richland, Washington, originally the site of a nuclear facility settled in 1944.

3. The Location of Cities

This section deals with issues pertaining to spatial features of the location of cities. We begin with some basic facts about the spacing of cities as the US urban system has evolved. To do this, we examine the evolution of average bilateral distances between cities relative to the average distance of a city from its nearest neighbor. Referring to Table 2, we see that the average bilateral distance

---

where latA and longA are the latitude and longitude of location A measured in radians. Similarly for latB and longB. The distance is then

$$\text{distance} = 3954 \times \text{acos}(\text{angle}).$$

acos(angle) gives us the approximate distance if the two points were located on a circle of radius one. We then need to multiply by the radius of (a circular) earth (3954 miles) to get an estimate of the distance. The assumption of a spherical earth leads to an error of approx 0.2% on an area the size of the US.

[5] Kim (1997), p. 7–9, discusses the original intention of the definition of U.S. regions as delineating areas of homogeneous topography, climate, rainfall and soil, but subject to requirement that they not break up states. By design, the definitions were particularly suitable for agriculture and resource-based economies. The role of those industries as inputs to manufacturing would make them likely to serve well as economic regions.

between cities rises by nearly 25%, from 802.5 miles in 1900 to 1005 miles in 1990, as the US urban system expanded over the North American land mass. In contrast, the average distance of a city from its nearest neighbor falls by 35%.[6] Put together, these numbers clearly show that US urban system both expanded and thickened during the twentieth century.

But what drives the location of cities that determines this spacing of the urban system? Theoretical reasoning points to first or second nature features of potential locations as the key determinant of whether or not a city is located there. Let us assume, for the moment, that the first and second nature features that are most favourable for cities are *not* evenly distributed across the US. This means, in turn, that cities themselves will not be evenly distributed across the US. In empirical terms, we can think of an uneven distribution as a departure from randomness, where randomness would imply that cities are equally likely to locate at all locations. This discussion suggests that a first step should be to test whether the urban system has evolved so as the location of cities is non-random.

To do this, we first need a precise definition of randomness in a spatial context. We will assume that under the null hypothesis of randomness, city location is randomly distributed according to a spatial Poisson process where the probability of a city locating in any given area is proportional to that area.[7] Clearly such a definition ignores the large topographical variation that we see across the US, but this variation is one of the key determinants of first nature differences across locations and thus should not be taken in to account in our definition of randomness.

The test that we use is a very simple test for non-randomness first proposed by Clark and Evans (1954). The basic idea is to assume that some underlying spatial probability process, in our case spatial Poisson, determines the distribution of cities and then to compare the actual distance between cities to the distance that we would expect if cities were located randomly according to this distribution. Although a full matrix of intercity distances is available, Clark and Evans (1954) show that a "sufficient" test of non-randomness can be based on the distance to nearest neighbor city only.[8]. However, even if location of cities is non-random, we may fail to reject the null-hypothesis of randomness. Non-randomness might be manifested in other dimensions than

[6]The dispersion of the former declines while that of the latter slightly grows, as evidenced by the coefficient of variation and nonparametric densities that we have estimated but do not report here. Both those distributions become more symmetric, as evidenced by the ratio of medians to means and the nonparametric densities.

[7]That is, we treat cities as points and ignore their own areas. For details, see Cliff and Ord (1975) and Ripley (1979).

[8]Sufficient in the sense that the statistical test is asymptotically valid for a large number of underlying spatial probability processes obeying a number of standard assumptions. See also Ripley (1979).

the distance to nearest neighbor. We return to this possibility below.[9]

Define: d_i as the distance to city i's nearest neighbor; d_A as the mean nearest neighbor distance; d_E as the expected mean nearest neighbor distance under the assumption that locations follow a spatial independent Poisson distribution; σ_E^2 as the expected variance of mean nearest neighbor distance under the assumption that locations follow a spatial independent Poisson distribution; ρ as the spatial density of cities; and, I as the number of cities. Then the Clark-Evans test for non-randomness is based on the simple test statistic $CE = (d_A - d_E)/\sigma_E$ which is distributed asymptotically $N(0,1)$. To calculate the statistic, we need to use our specific assumption on the spatial process that governs the random location of cities to allow us to calculate the expected mean and variance. Under our assumption that cities are located according to a spatial Poisson distribution, these take a simple form: the expected mean nearest neighbor distance is $d_E = 1/(2\sqrt{\rho})$, and the standard deviation is $\sigma(d_E) = 0.26136/(\sqrt{I\rho})$.

Table 3 shows the results for the US as a whole for each of the census years. The final column reports $R = d_A/d_E$, the ratio of actual to expected distance. A number less than one indicates that cities are closer together than would be expected if they were randomly located. Conversely, a number greater than one indicates that cities are further apart than would be expected if they were randomly located. The CE column reports the Clark-Evans test statistic that tells us whether this departure from randomness is significant. From the table, we see that US cities are spaced closer than we would expect if they were randomly located, but that this non-randomness is only significant at the beginning of the century. We find this result surprising for two reasons. First, we had strong theoretical priors that first and second nature features would matter for city location. Second, casual observation suggests that cities are very clustered in certain parts of the country.

For those who find the theoretical reasoning convincing, there are a number of possible reactions to this finding that, for the US urban system as a whole, city locations are essentially random. The first is to question our underlying assumption on the uneven distribution of first and second nature features. The second is to argue that non-randomness might show up in more detailed (higher order) features of the spacing of cities. The third possibility is that the error occurs because we consider the US urban system as a whole when we know there are strong regional variations particularly with respect to the distribution of first nature features. As we see next, all of these

[9]Indeed, a large number of additional test statistics, including extensions to k-nearest neighbor methods, have been developed since the original Clark and Evans test used here. See, for example, Diggle (1983) for a description of these methods.

arguments have some bite.

Let us start with the possibility that first and second nature features are evenly, rather than unevenly distributed across the US. The evidence we present in section 4 shows that second nature features are clearly not evenly distributed across space.[10] What about first nature features? The problem is that these first nature effects are largely unobserved and may change over time thus making it impossible to directly assess whether these features are evenly or unevenly distributed.

We can make some progress however, by considering indirect evidence on the role of first nature. A theoretical model that explains city location on the basis of first nature should see good locations settled first and the largest cities developed at these good locations. Given our discussion above, we cannot assess whether good locations are evenly distributed, but we can consider whether good sites are settled first. If better first-nature sites are settled earlier, on average – arguably, a rather simplistic view of history – then early settlement would confer a permanent advantage in terms of city size. To test this relationship we would like to consider the relationship between relative city size and the date at which a city was settled. We do this by constructing the distribution of city sizes relative to the US average (US relative) and the distribution of city sizes relative to cities settled at roughly the same date (Same date relative) and then studying how the two relate to one another.

A general way to look at these kinds of relationships between two distributions of interest, has been proposed by Danny Quah in a series of papers [Quah *op. cit.*]. Quah proposes estimating stochastic kernels that give the distribution of one of the variables (Same date relative) conditional on the distribution of another variable (US relative). These tools have been widely used to study income inequalities across countries and regions but have not been widely used to study urban systems.[11] There are several advantages in using this approach to study the spatial evolution of the urban system. Most importantly, it imposes no structure on the underlying relationships. The estimated relationships can be non-linear and are allowed to change over time. A further attractive property is that, in the evolution of the urban system, no city is truly representative of the entire distribution of cities. Standard parametric tools rely on the assumption that there is some average (representative) unit whose outcomes can be modelled in a concise way as the function of a limited

[10]For example, given the distribution of economic activity across the US, some cities clearly have better access to large markets.

[11]In a companion piece, Overman and Ioannides (2001), we use these tools to study non-spatial features of the evolution of the US urban system.

number of variables and unknown parameters. The stochastic kernel method does not rely on this representative agent assumption. This flexibility makes the method ideal for studying highly non-linear evolving systems. For readers unfamiliar with the methodology, Appendix A provides additional technical information on estimation.

Figure 1 shows the stochastic kernel mapping from the distribution of US relative city sizes to the distribution of same date relative city sizes [*c.f.,* Quah (1999)]. The US relative city sizes are constructed by taking the (log of the) ratio of city size to the US average city size. The same date relative city sizes are defined as the (log of the) ratio of city size to the mean city size for cities that were settled at a similar period. Settlement dates are constructed as outlined in section 2 and grouped in to similar settlement dates using twenty year bands. Both distributions are normalized by subtracting their mean and dividing by their standard deviation, so that each univariate distribution has a variance of 1 and a mean of 0.

The way to interpret this stochastic kernel is as follows. Take a point on the US relative axis, say 1.0, which corresponds to a city with (log) city size that is one standard deviation above the US (log) mean. Now imagine taking a cross-section across the stochastic kernel orthogonal to the US relative axis and parallel to the same date relative axis. This cross-section traces out a conditional distribution giving the same date relative city sizes for all the cities whose population is one standard deviation above the US mean. The stochastic kernel plots these conditional distributions for all values of US relative city size.[12]

Consider the end of the century first. To continue with our illustrative example, take the point 1.0 on the US relative axis and read across the stochastic kernel parallel to the same date relative axis to get the conditional distribution. The mass of this conditional distribution is tightly centred around 1.0 on the same date relative axis. That is, cities one standard deviation above the US mean also tended to be one standard deviation bigger than cities settled at the same date. Looking at the kernel we see that the same applies for cities one standard deviation below the mean and for US relative city sizes more generally.

From the stochastic kernel for 1990 it is clear that initial benefit conferred no advantage at the end of the century. To understand why we reach this conclusion it is useful to ask what would we have expected to see if early settlement *had* conferred a permanent advantage? If this were the case, cities that were large relative to the US average, would be better first-nature sites, settled

[12]Actually, for technical reasons the kernel is not plotted for the very largest cities. See Appendix A for more details.

earlier. Thus, although they are large relative to the US, we would expect them to be a similar size to sites that were settled at the same time (on similarly good sites). Likewise, smaller cities would be located on poorer sites settled later. However, although they are small relative to the US, we would expect them to be a similar size to cities on similarly poor sites that were settled at similar late dates. That is, if first nature characteristics matter most, then the stochastic kernel should map cities to approximately zero in the same-date relative distribution. Cities settled at similar dates should be of similar sizes.

Therefore, from the stochastic kernel for 1990 one can only conclude that early settlement in good first nature locations has conferred no permanent advantage. One could object to this interpretation by arguing that changes in transportation and infrastructure have rendered inland and temperate climate locations much better first nature sites than they were 100 hundred years ago. This is clearly true, but what stands out is that this relationship also held at the beginning of the century, as the figure for 1910 shows. If good sites were settled first, the advantages of these sites had already largely diminished by the beginning of the twentieth century.

These findings are actually consistent with the idea that economic geography might play some role in determining the location of cities, despite our finding that location was essentially random. If first nature effects do not play a particularly large role in determining city outcomes, then it is possible that there are many reasonably good sites fairly evenly spread across the US. If good sites are spread out, then location may appear to be essentially random. Viewed from a dynamic perspective, it may be that what makes a good location is changing over time and with so few cities entering there are plenty of good first nature sites in terms of these new factors that make a particular location desirable. Both of these stories reintroduce first nature geography as a determinant of city location, but only at the expense of downplaying first nature importance for city outcomes in the twentieth century!

What about second nature geography? As we mentioned above the evidence we present in section 4 shows that second nature features are clearly not evenly distributed across space. This observation, coupled with the finding of random location of cities, suggests that second nature might not play a particularly important role in determining the location of cities. We can, however, identify some role for second nature if we consider the location of different sizes of cities.

Figure 2 shows a stochastic kernel mapping the distribution of population (US relative city size) to the distribution of distance to nearest neighbors, $\hat{f}(d_i|P_i)$. Both variables are normalised exactly as before. Starting this time with the beginning of the century, the figure for 1910 shows that smaller cities tended to locate far away from their neighbors. A city one standard deviation below mean population size tended to be one standard deviation above the mean in terms of distance to their nearest neighbor. The reverse holds for larger cities. That is, big cities are only ever found close to other cities, while small cities may be close to other cities but are much more likely to be by themselves. By 1990 the relationship has changed. Smaller cities still tend to be further from their nearest neighbor, but the relationship is not as stark as in 1910. While it is possible to contrive a first nature story that might lead to this particular pattern, it seems much more likely that second nature interactions between cities explain the different relative location patterns of large and small cities. Again, this story reintroduces a role for geography in determining the spacing of different types of cities rather than the spacing of cities per se.

In our discussion of the randomness result we suggested three issues that we might like to consider further. We have provided some evidence for the possibility that first nature and second nature may still play some role in determining the location patterns of cities despite the fact that, for the US urban system as a whole, city locations are essentially random. We now consider a final possibility, that the system as a whole looks random because of offsetting factors when we pool different regional systems that are non-random. To give a concrete example, cities in one region may be too close together because in that region closely spaced locations on the seacoast make good sites for cities, while cities in another region may be too spread out because in that region mountains and deserts force cities apart. When we pool these two regions it is possible that the overall pattern may look random.

Table 4 shows that this may be part, but not all of the story. The table shows the same nearest neighbor statistic, reported in table 3, but calculated for census regions rather than the entire US. (Figure 3 clarifies the designation of US census regions). The table shows that randomness can be rejected for two out of the four regions for which we report results. In particular, the South and West regions show strong evidence of non-randomness. Cities in the South are too far apart. An explanation could be that although mountain ranges there are distinct but not particularly massive, the South's urban development might have been influenced by its plantation economy past.[13] Cities in the West are too close together, in part because deserts and the sea coast restrict the area of urban development. We still cannot reject randomness for the Mid West and North East

[13] We owe this suggestion to a referee.

however. In these two regions cities are essentially randomly located suggesting that we still find a puzzle about the role of first and second nature at other spatial scales.

To conclude, most theories of the location of cities allow some role for first or second nature in driving location. In light of this theoretical reasoning, our results on the randomness of city locations present somewhat of a puzzle. We can rescue a role for first nature by considering smaller spatial scales and allowing for the fact that first nature might not actually be that important during the twentieth century. Likewise, we find some role for second nature geography in explaining the location patterns of different size cities. Clearly, more work remains to be done. However, for now, we turn from the issue of city location to the issue of city sizes and wages to assess the role that second nature might play in determining city outcomes.

4. Spatial Features of The City Size and Wage Distributions

We have already said as much as we are going to say on the role of first nature and the size of cities. To reiterate, if early settlement of good sites conferred an advantage, then this advantage had already faded by the beginning of the twentieth century. In this section, our focus is on the role that second nature features may play in determining the distribution of city sizes. We again use the same tools, developed by Danny Quah [Quah *op. cit.*] and employed above, to characterise some key spatial aspects of the evolution of the US urban system. We start with a key question and consider whether second nature features help determine the distribution of city sizes and wages. In particular, we will consider whether spatial interactions between locations (both urban and non-urban) help determine city outcomes. This focus is driven both by our interest in spatial features per se and by the type of data that we have available.

Traditionally, models of the urban system have investigated the spatial interaction between locations using the concept of market potential. As we see in its definition below, this concept measures whether a location has good access to markets. Cities should be large and pay high wages if their location has high market potential [See for example Harris (1954)]. New economic geography models [Krugman (1992); Fujita *et al.* (1999)] have formalised this reasoning and shown that market potential should be a function of city incomes, distances between cities and the city price indices for manufactured goods. These models suggest that the effect of high market potential at a location might not be unambiguously positive. Given those theoretical foundations, we begin our analysis of spatial interactions between cities by using the concept of market potential.

In what follows, we examine the relationship between city sizes or wages and market potential using a series of stochastic kernels. However, before turning to details on the construction of the market potential, we briefly consider the advantages of using our non-parametric approach to study these spatial interactions. All of these stem from the fact that we do not need to impose any restrictions on the mapping from market potential to city sizes or wages.[14] In particular, we do not have to impose any form of linearity. Nor do we have to impose monotonicity so, for example, we can allow the distribution of city size or wages conditional on market potential to be twin peaked. This sort of flexibility is important given that theoretical reasoning suggests that competition effects may sometimes dominate demand effects implying that high market potential may be associated with both small and large cities. Finally, we do not have to restrict the mapping to be stationary over time. As we show below, this flexibility allows us to identify features of the spatial evolution of the US urban system that would be completely obscured were we to adopt a more standard parametric approach.

A. *City sizes conditional on market potential*

Given the available data we can construct three different definitions of market potential for city i at time t based on the following formula:

$$mp_{it} = \sum_{j \neq i} \frac{P_{jt}}{D_{ij}}. \tag{1}$$

The first two measures differ depending on whether the summation is across all cities or all counties in the US. In words, city i's market potential is the sum over all other cities (counties) j of population in city (county) j, P_{jt}, weighted by the inverse of distance between i and j, D_{ij}.When the summation is across all cities, we will refer to this as market potential (cities), and when it is across counties as market potential (counties)[15]. Taking different definitions is interesting because it allows us to see whether spatial interactions between cities differs from general spatial interactions between cities and other (non–city) locations in the US.

In addition to population data, we also have wage data for cities, but not counties, back until 1900. These wage data allow us to construct a third measure of market potential, where cities are weighted by average wages as well as distance: $mp^W_{it} = \sum_{j \neq i} \frac{W_{jt} P_{jt}}{D_{ij}}$. We will refer to this as market

[14]That is, over and above some regularity conditions. See Appendix A.

[15]For the county based market potential measure, note that the sum is over all counties that are not part of that metropolitan area in 1990.

potential (wages). This measure may better capture the importance of demand from other cities and regions than the measures that only consider population, and is thus closer to the Krugman version of the market potential model [Krugman (1992)].

The definition of market potential involves a somewhat arbitrary choice on the importance of distance. Results do not, however, appear to be too sensitive to the assumptions on how distance enters. In addition, in common with many authors, we are assuming that transport costs are directly related to the distance between cities without any consideration of actual transport networks and costs. Finally, given the lack of data on actual transport costs and changes in sectoral composition of output, we have chosen to take the "neutral" viewpoint that general transport costs are unchanged over the sample period. Again, without any further information on transport costs over the period, it is unclear what alternative assumption would be better.[16]

As before, when calculating the stochastic kernels both population and market potential are taken as (the log of) ratios to the contemporaneous mean, and the distributions are normalized by subtracting the mean and dividing by the standard deviation. The stochastic kernels are read exactly as before. For example, take a point on the market potential axis, say 1.0, which corresponds to a city with log market potential that is one standard deviation above the log mean. Cutting across the stochastic kernel parallel to the city size axis gives the conditional distribution of relative city sizes for cities with market potential one standard deviation above the mean. The stochastic kernel plots these conditional distributions for all values of market potential.[17]

B. *Stochastic kernels for city sizes conditional on market potential*

To study how city sizes and wages might be affected by spatial interactions, we report results for several stochastic kernels in the form of three-dimensional figures and contours.[18] Figure 4 reports stochastic kernels $\hat{f}(P_i|mp_i)$, for city size conditional on market potential (cities), Figure 5 on market potential (counties) and Figure 6 on market potential (wages).

From Figures 4, a and b, and 5, a and b, we see that the 1910 kernels are somewhat skewed towards the diagonal. This is perhaps easiest to see in the contour plot. For the low market

[16] It would be possible to allow for the effect of distance to decrease through time. However, the changing composition of consumption from manufacturing to services, means that, at an aggregate level, it is not clear whether general transport costs have been rising or falling. Thus, Hanson (2000) finds that the estimated effects of distance increase between 1970 and 1980, which he interprets as a net increase in effective transport costs.

[17] As before we do not give results for the very largest cities although we do use them to calculate market potentials.

[18] The contours work exactly like the more standard contours on a map. Any one contour connects all the points on the stochastic kernel at a certain height.

potential cities, we can readily identify a peak in the stochastic kernel that lies on the diagonal in the south-west of the diagram.[19] This peak contains most of the mass for the smaller cities. In contrast, the conditional distribution for the largest cities is relatively flat although there is some evidence that the very highest market potential cities do tend to be bigger than average.

This relationship between market potential and city size does not persist as the US urban system evolves. The entire series of snapshots, not reported here, show the stochastic kernels for the each decennial year 1900 – 1990, slowly twisting back until they appear, by 1990, to have become virtually independent of market potential. The peaks become less and less pronounced, as the distribution of city sizes conditional on low market potential shows greater variance. By 1990, Figures 4, c and d, and 5, c and d, suggest that the conditional distributions of city sizes are almost identical across all values of market potential. Only for the very largest cities is city size positively related to market potential. We underscore the importance of this finding. It suggests, at least from a non-parametric vantage point, that the distribution of city sizes conditional on market potential is basically independent of the level of market potential.

This basic result does not change when we incorporate information on the wages paid in different cities. Figure 6 considers the co-evolution of market potential (wages) and city sizes. The stochastic kernels for city size distributions conditional on market potential (wages), for 1910 and 1990, accord with those in Figures 4 and 5. Again, the entire series for the century – not reported here – shows the kernel slowly twisting back until by 1990, the distribution of city size has become virtually independent of market potential. These results thus provide additional support for our earlier comments.

Before proceeding, we summarise what our results so far tell us about the spatial interactions between cities. First, they tell us that this relationship is non-linear — at least to the extent that there may be differences between small and large cities. Second, the nature of the interaction evolves over time. That is, the mapping from market potential (however measured) to population is not stationary. Third, if, as theory suggests, we can capture the second nature features of the system through a reduced-form market potential variable, then the spatial interaction between cities was weak at the beginning of the century and has further weakened over time.

As for the results on city location, these findings are puzzling given our theoretical priors. Again, we can investigate the issue further by considering extensions along a number of dimen-

[19] Around (-1.0,1.0) for figure 4 and around (-0.75,0.75) for figure 5.

sions. We choose to focus on three particular features. First, we consider spatial interactions between neighboring cities. Second, we consider whether spatial interactions might determine population growth rates rather than levels. Finally, we consider the effect of spatial interactions on wages rather than population.

C. *Spatial interactions among neighbors*

Examining the relationship between cities and their neighbors is interesting for two reasons. First, because we saw above that looking at nearest neighbors allowed us to identify a difference between large and small cities. Second, because the new economic geography literature suggests that large cities might cast an agglomeration shadow that affects their immediate neighbors. [20] To get at the first of these two issues, Figure 7, c and d, reports stochastic kernels for the size of nearest neighbors conditional on city size, for 1910 and 1990, respectively. From this, we see that the distribution of nearest neighbor size is practically independent. This finding is consistent with the low correlations between sizes of cities and nearest neighbors reported in the eighth column of table 2. Putting this together with our earlier finding on the location of cities, we see that there are differences in the spacing of large and small cities but *not* in the type of cities that they have as neighbors. Large cities tend to be found close to other cities and their neighbors may be large or small. The same is true for the neighbors of small cities, even though small cities tend to be further away from their neighbors.

Turning briefly to the issue of agglomeration shadows, Figure 7, a and b, reports stochastic kernels for city size distributions conditional on market potential (cities), excluding the component of market potential from the nearest neighbor. If agglomeration shadows mattered, then cities with very high market potential might be small if they fell inside this agglomeration shadow. This is one possible explanation of our finding of only weak evidence for a positive effect of market potential when we considered the stochastic kernels from market potential to city size (Figures 4, 5 and 6). Figure 7, a and b, show that such considerations do not change our overall conclusions with respect to the spatial interactions between cities. Clearly more work remains to be done on this issue. Yet for the moment, we find relatively limited evidence in support of the idea of agglomeration shadows. Overall, our results for nearest neighbors suggest that this dimension

[20] Krugman (1993) suggests that once a particular site is settled, its presence may skew further development in its vicinity in its favor, via its agglomeration shadow.

might matter more for city location than it does for city outcomes. We now turn from nearest neighbors to consider the growth rates of cities.

D. *City growth rates conditional on market potential*

Given that there is only a very weak relationship between city size and market potential, it would be unlikely that we would find one between city growth and market potential.[21] Instead, we look at whether a city's growth relative to its long run average can be explained by deviations of that city's market potential from its long run average. That is, whether spatial interactions between cities can explain periods of 'over' or 'under'-performance relative to long run trends. To do this, we consider the difference between this period's relative growth rates and the time average of growth rates for that city. We also do the same for market potential. Thus for both variables we look at differences from a city specific fixed effect. Figure 8 shows stochastic kernels for the distribution of growth rates conditional on the distribution of market potentials.[22]

Once again, they show clearly, that the relationship changes slowly during the century. At the beginning of the century, cities with (historically) unfavourable levels of market potential were actually seeing (historically) high levels of growth. By 1990 cities with (historically) high market potential were seeing higher growth. Interestingly, this finding is consistent with our findings above on first nature and city size. Our results there suggested that if first nature confers an advantage, the effect was most important at the beginning of the century. Put together with our results here, it seems possible that good first nature sights at the beginning of the century were growing fast even though their second nature position might have been less favourable. By the end of the century, once the system had developed, first nature matters less but cities with historically high market potential do now tend to see higher growth. Thus, by the end of the twentieth century spatial interactions between cities do help to explain good times and bad times even if we find limited evidence for their role in explaining which cities are large or small.

[21] If there was a relationship between growth and market potential we would clearly expect it to manifest in a relationship between size and market potential.

[22] As before, we normalise both variables using the mean and standard deviation when calculating the stochastic kernel. This normalisation helps control for the fact that the *cross-sectional* average market potential and growth rates vary over time.

E. *City wages conditional on market potential*

We have just seen that there is a role for spatial interactions in explaining fluctuations in city growth rates. We now show that there is also a role in explaining outcomes in terms of wages. To do this we use our stochastic kernels to analyse the evolution of the wage distribution rather than city sizes. Again, we capture spatial interactions between cities in the determination of the wage distribution through the use of our market potential measures. In general, we would expect cities with high relative market potential to have high relative wages. This prediction is not confirmed by the 1910 data, reported in Figure 9, a and b. Wages are relatively high for cities with low market potential. However, as before, as the urban system develops the relationship changes. According to Figure 9, c and d, in 1990, the stochastic kernel is slowly twisting towards the diagonal with higher wages associated with larger market potential. Putting this together, we reach a now familiar conclusion. At the beginning of the twentieth century, a city's ability to generate high wages appears to be independent of spatial interactions. However, by the end of the century, spatial interactions do help determine wages even if they have relatively little impact on overall population.

Overall, we can reach a number of conclusions about the impact of spatial interactions on city size. As a general rule, there is no *simple* relationship governing the spatial interaction of cities. The relationships we can identify are generally non-linear, sometimes non-monotonic and nearly always changing over time. All of this clearly urges caution in application of standard parametric techniques to studying the evolution of the US urban system. As well as urging caution, our non-parametric techniques have allowed us to identify a number of spatial features of city outcomes. There is limited evidence of a weak relationship between good access to markets (high market potential) and city size at the beginning of the century. This relationship has weakened further over time. This result may be driven by the fact that market potential should take into account competition effects that cause agglomeration shadows. However, a preliminary analysis of these phenomena looking at nearest neighbors finds no clear evidence that it is empirically important. Despite these somewhat negative findings, spatial interactions do play a role, both in terms of periods when cities over-perform and in terms of wages. In direct contrast to our findings on spatial interactions and population, both of these relationships are strong towards the end of the century and weak at the beginning.

5. Conclusions

This paper has used a number of different approaches to analyse the spatial evolution of the US urban system over the period 1900 to 1990. Our work is the first to consider a wide range of spatial features of the US urban system as it has evolved during the twentieth century. The techniques that we use are particularly appropriate for studying the complex evolution of the system because they impose so little structure on the analysis. Our results confirm some theoretical insights, but also throw up a number of puzzles.

The first group of findings concern the spatial pattern of the location of cities in the US. By the end of the twentieth century, the location of cities in the US was essentially random. This finding is partly explained by the fact that we pool regions with very different topographies. Breaking down the result by region suggests that cities in the West are too close together, while cities in the South are too far apart. But, this aggregation does not tell the whole story. Cities in the North and Mid-West are still randomly distributed. Our finding of randomness does not totally preclude a role for first and second nature in determining the location of cities. One possibility is that there are many good first nature sites for cities at the end of the twentieth century. This was less true at the beginning of the century as our results on city size and settlement dates have shown. We can also identify a role for second nature, by considering the location of different sizes of cities. Large cities tend to be located close to other cities, while small cities tend to be isolated. Many other spatial features of the location of cities remain to be investigated and we are still surprised at how little work is being (has been) done in this area.

Our second group of findings concern the role of the spatial interactions between cities in explaining city outcomes. Our results suggest that there is no simple positive relationship between city sizes and market potentials. This relationship appears to change substantially over time. There is some evidence of a positive relationship between city sizes and market potential at the start of the century. At the end of the century, such a positive relationship apparently only holds for the largest cities. In fact, an important finding stands out very clearly. By the end of the century the distribution of city sizes conditional on market potential is nearly independent of market potential. Similar results hold if we recalculate market potential weighting by city wages or using data from counties rather than cities. Further evidence suggests that this result would not change significantly if we could take account of the agglomeration shadow effects in models of new economic geography. However, much more work needs to be done before we can conclude that these effects are definitely not present in the data.

There is, however, a role for second nature in determining city outcomes along at least two dimensions. First, spatial interactions do matter for understanding when a city grows fast relative to its historical average. Second, spatial interactions do matter for understanding which cities pay higher wages. Interestingly, these spatial features are only weakly present at the beginning of the period and both slowly emerge during the century.

Taken together, our results provide a rich picture of the spatial evolution of the US urban system. But many key features of that spatial evolution remain to be examined. We leave these issues to further work.

Appendix A. Technical appendix

A. *Estimating the stochastic kernel*

The stochastic kernel shows the distribution of some variable y (e.g. population) conditional on the distribution of another variable x (e.g. market potential). To estimate that stochastic kernel, we first derive a non–parametric estimate of the joint distribution $f(x,y)$. Kernel estimation of this joint density requires picking a kernel and a bandwidth. Results are generally not sensitive to the choice of kernel (here we use a Gaussian kernel). The choice of bandwidth does matter. The bandwidth we use is the optimal bandwidth based on Silverman (1986) and is a function of the range or the variance whichever is the larger.

Once we have the joint distribution $f(x,y)$ we numerically integrate under this joint distribution with respect to y to get $\hat{f}(x)$. (We could also estimate the marginal distribution $f(x)$ using a univariate kernel estimate). The asymptotic statistical properties of both estimators are identical, and in practice tend to produce very similar estimates. Next we estimate the distribution of y conditional on x by dividing through $f(x,y)$ by $f(x)$. Thus we estimate $f(y|x)$ by: $\hat{f}(y|x) = \frac{\hat{f}(x,y)}{\hat{f}(x)}$. Under regularity conditions, this gives us a consistent estimator for the conditional distribution for any value x. The stochastic kernel records this conditional distribution for all values of x.

B. *Treatment of outliers*

For the particular applications in this paper, we sometimes ignore very large cities when calculating both the optimal bandwidth and *plotting* the figures. At points in the sample period, New York is up to 25 times the mean city size (1930). Including these very large cities is conceptually simple, but technically problematic. Very large outliers automatically drive up the optimal bandwidth that we use when calculating the stochastic kernels. When this happens, the detail in the lower end of the distribution (comprising the main body of cities) is obscured, as the estimates are over–smoothed. To get round this problem, we calculate the bandwidth restricting the sample range to the bottom 95% of all cities in any single year. A similar problem occurs with the plots. Including cities 25 times the mean city size would obscure all the detail just to include one additional city. Thus, we also ignore these cities when plotting the stochastic kernel. However, once the bandwidth is calculated *all* cities are used to *calculate* the stochastic kernel.

C. *Sample size*

Stochastic kernels for the beginning of the century are always for 1910 rather than 1900. This is for two reasons. (i) For consistency when we plot the stochastic kernels using growth rates. (ii) To increase the sample size. In 1900 we only have 112 cities, in 1910 we have 139. There are small-sample issues in non-parametric estimation concerning the speed with which the stochastic kernel converges to the true distribution. See Hardle (1994) for details.

D. *Parametric spatial regressions*

The stochastic kernels from market potential to population are quite closely related to the parametric spatial autoregressions suggested by Anselin [Anselin (1988)] and others. In fact, the calculation of market potential uses a spatial weighting matrix with each element (w_{ij}) equal to the inverse of the distance D_{ij} between cities i and j. However, our nonparametric approach does not impose a uniform coefficient on the spatial AR term thus constructed and does not require the mapping from the spatial AR term to population to be linear or one-to-one.

2. References

Anselin, Luc (1988), *Spatial Econometrics, Methods and Models,* Boston, MA. Kluwer Academic.

Black, Duncan, and J. Vernon Henderson (2002), "Spatial Evolution in the USA," *Journal of Economic Geography,* forthcoming.

Bogue, Donald (1953), *Population Growth in Standard Metropolitan Areas 1900 – 1950,* Oxford, Ohio: Scripps Foundation in Research in Population Problems.

Carrol, G. (1982), "National City Size Distributions: What do we know after 67 years of research", *Progress in Human Geography*, 6, 1–43.

Cheshire, Paul (1999), "Trends in Sizes and Structure of Urban Areas in *Handbook of Regional and Urban Economics* Vol 3., P. Cheshire and E. S. Mills (eds.), North Holland.

Clark, Philip J., and Francis C. Evans (1954), "Distance to Nearest Neighbor as a Measure of Spatial Relationships in Populations," *Ecology,* 35, 4, 445–453.

Cliff, A. D., and J. K. Ord (1975), "Model Building and the Analysis of Spatial Pattern in Human Geography," *Journal of the Royal Statistical Society,* B, 39, 297–348.

Combes, P.-P., and M. Lafourcade (2001), "Transportation Costs Decline and Regional Inequalities: Evidence from France", CEPR discussion paper, 2894.

Diggle, Peter J. (1983), *Statistical Analysis of Spatial Point Patterns*, Academic Press, London.

Dobkins, Linda H., and Yannis M. Ioannides, "Dynamic Evolution of the U.S. City Size Distribution," in J. Huriot and J. Thisse, eds., *The Economics of Cities,* Cambridge University Press, 2000, 217–260.

Dobkins, Linda Harris , and Yannis M. Ioannides (2001), "Spatial Interactions among U.S. Cities," *Regional Science and Urban Economics,* 31, 6, 701-731.

Eaton, J. and Z. Eckstein (1997), "Cities and Growth: Theory and Evidence from France and Japan" 27, 443-474.

Fujita, Masahisa, Paul Krugman, and Anthony Venables (1999), *The Spatial Economy,* MIT Press, Cambridge, MA.

Fujita, Masahisa, and Jacques-Francois Thisse (2002), *The Economics of Agglomeration,* Cambridge University Press, Cambridge.

Gabaix, Xavier (1999), "Zipf's Law for Cities: An Explanation," *Quarterly Journal of Economics,* CXIV, August, 739 – 767.

Gabaix and Ioannides (2003), "The evolution of city size distributions", forthcoming J. V. Henderson and J. Thisse (eds.) *Handbook of Regional and Urban Economics*, Vol 4, North Holland

Hanson, Gordon (2000), "Market Potential, Increasing Returns and Geographic Concentration," University of Michigan, working paper.

Hardle, W. (1994), *Applied Nonparametric Regression"*, Cambridge University Press.

Harris, C. D. (1954), "The Market as a Factor in the Localization of Industry in the United States," *Annals of the Association of American Geographers,* 44, 315–348.

Henderson, J. Vernon (1974), "The Types and Size of Cities," *American Economic Review*, 64, 640-656.

Henderson, J. Vernon (1988), *Urban Development: Theory, Fact and Illusion,* Oxford University Press, Oxford.

Ioannides, Yannis M., and Henry G. Overman (2003), "Zipf's Law for Cities: An Empirical Examination," *Regional Science and Urban Economics,* 33, 2, 127–137.

Kim, Sukkoo (1997), "Economic Integration and Convergence: U.S. Regions, 1840-1987," *Journal of Economic History,* 58(3), 659–683.

Krugman, Paul (1991), "Increasing Returns and Economic Geography," *Journal of Political Economy,* 99, 483–499.

Krugman, Paul (1992), "A Dynamic Spatial Model," NBER Working Paper No. 4219, November.

Krugman, Paul (1993), "First Nature, Second Nature and Metropolitan Location," *Journal of Regional Science,* 33, 129–144.

Overman, Henry G., and Yannis M. Ioannides (2001), "Cross-Sectional Evolution of the US City Size Distribution," *Journal of Urban Economics,* 49, 543–566.

Quah, Danny T. (1993), " Empirical Cross-Section Dynamics and Economic Growth," *European Economic Review,* 37, 2/3, 426–434.

Quah, Danny T. (1996a), "Empirics for Economic Growth and Convergence" *European Economic Review,* Vol 40, no.6 pp 1353-1375.

Quah, Danny T. (1996b), "Regional Convergence Clusters across Europe," *European Economic Review,* 40, 951–958.

Quah, Danny T. (1997), "Empirics for Growth and Distribution: Stratification, Polarization and Convergence Clubs, " *Journal of Economic Growth,* 2, 1, 27–59.

Quah, Danny T. (1999), "Regional Convergence from Local Isolated Actions: II Conditioning". Published in Spanish as "Cohesion Regional Mediante Actuaciones Locales Aisladas: Condiciones," in *Dimensiones de la Desigualdad, III Simposio Sobre Igualdad y Distribucion de la Renta y la Riqueza,* Volumen I.

Ripley, B. D. (1979), "Tests of 'Randomness' for Spatial Point Patterns," *Journal of Royal Statistical Society,* B, 41, 3, 368–374.

Silverman, Bernard W. (1986), *Density Estimation for Statistics and Data Analysis,* Chapman and Hall, New York.

Tabuchi, Takatoshi (1998), "Urban Agglomeration and Dispersion: A Synthesis of Alonso and Krugman," *Journal of Urban Economics,* 44, 333–351.

Thomas, Alun (1996), "Increasing Returns, Congestion Costs, and The Geographic Concentration of Firms," I.M.F., March, mimeo.

1	2	3	4	5	6
Year	U.S. Pop. (000)	U.S. Pop. Urban (000)	Number of cities	Mean Size	Median Size
1900	75,995	29,215	112	259952	121830
1910	91,972	39,944	139	286861	121900
1920	105,711	50,444	149	338954	144130
1930	122,775	64,586	157	411641	167140
1940	131,669	70,149	160	432911	181490
1950	150,697	85,572	162	526422	234720
1960	179,323	112,593	210	534936	238340
1970	203,302	139,419	243	574628	259919
1980	226,542	169,429	322	526997	232000
1990	248,710	192,512	334	577359	243000

All figures are taken from *Historical Statistics of the United States from Colonial Times to 1970*, Volumes 1 and 2, and *Statistical Abstract of the United States, 1993*.

Table 1. Descriptive statistics: decennial data, 1900 - 1990

	Bilateral distances			Nearest neighbor distances			Nearest neighbor correlations	
	mean	median	variance	mean	median	variance	sizes	growth rates
1900	802.5	642.5	594.8	70.9	55.7	61.8	-.073	.557
1910	863.8	686.5	623.2	68.3	54.6	58.3	-.059	.256
1920	864.0	697.9	609.6	66.2	51.8	54.5	-.058	.528
1930	876.9	720.1	600.2	64.8	51.8	50.0	-.065	.457
1940	884.9	734.9	596.7	64.4	53.4	46.1	-.062	.674
1950	890.8	745.7	594.0	65.3	53.4	46.6	-.062	.436
1960	940.4	813.8	603.0	56.9	46.3	52.5	.027	.126
1970	981.3	841.3	631.3	52.5	42.0	41.2	.091	.394
1980	998.7	856.9	639.6	45.9	36.9	33.2	.138	.467
1990	1005	868.5	637.1	45.5	37.0	32.3	.172	

Columns two to four provide summary statistics for the matrix of bilateral distances between all cities at time t. Columns five to seven provide summary statistics for the vector of distances from nearest neighbor. Distances are calculated as described in the text. The last two columns give correlations between nearest neighbors in terms of population size and growth rates. Calculations exclude Honolulu and Anchorage.

Table 2. Distances and Nearest neighbor Correlations

Year	Area	Number of cities	Actual distance	Density	Expected distance	Variance	CE [N(0,1)]	R
1900	2969834	112	70.9	3.77E-05	81.41	4.02	-2.61	0.87
1910	2969565	139	68.3	4.68E-05	73.08	3.24	-1.47	0.93
1920	2969451	149	66.2	5.01E-05	70.58	3.02	-1.45	0.93
1930	2977128	157	64.8	5.27E-05	68.85	2.87	-1.41	0.94
1940	2977128	160	64.4	5.37E-05	68.20	2.81	-1.34	0.94
1950	2974726	162	65.3	5.44E-05	67.75	2.78	-0.88	0.96
1960	2968054	209	56.9	7.04E-05	59.58	2.15	-1.24	0.95
1970	2967166	242	52.5	8.15E-05	55.36	1.86	-1.53	0.94
1980	2966432	320	45.9	11.0E-05	48.14	1.40	-1.59	0.95
1990	2963421	332	45.5	11.2E-05	47.23	1.35	-1.28	0.96

Figures for land area exclude Hawaii and Alaska and are taken from *Historical Statistics of the United States from Colonial Times to 1970*, Volumes 1 and 2, and *Statistical Abstract of the United States, 1993*. Note that both sources show the (non-water) land area decreasing from 1950. Number of cities is as per table 1 but excludes Honolulu and Anchorage. Actual distance is from table 2. Density is number of cities divided by area. Expected distance and variance are calculated as per the formulas in the text. The CE column gives the Clark-Evans test while the final column, R, reports the ratio of actual to expected distance.

Table 3. Clark Evans Test - US

Year	Mid West	North East	South	West
1900	1.14	1.11	0.89	0.93
1910	1.15**	1.1	1.13*	0.68**
1920	1.13*	1.1	1.08	0.8
1930	1.12	1.1	1.17**	0.64**
1940	1.12	1.1	1.12*	0.72**
1950	1.12	1.1	1.17**	0.72**
1960	1.04	0.93	1.16**	0.84
1970	1.03	0.92	1.14**	0.83**
1980	1.03	1.08	1.08**	0.83**
1990	1.03	1.08	1.11**	0.83**

The table gives values of R (the actual to expected distance) for each of four census regions. ** indicates significant at the 5% level, * indicates significant at the 10% level. Mid-West comprises East North Central and East South Central; North-East comprises Mid-Atlantic and North East; South comprises South Atlantic; West North Central and West South Central; West comprises Mountain and Pacific (excluding Hawaii and Alaska)

Table 4. Clark Evans Test - Regions

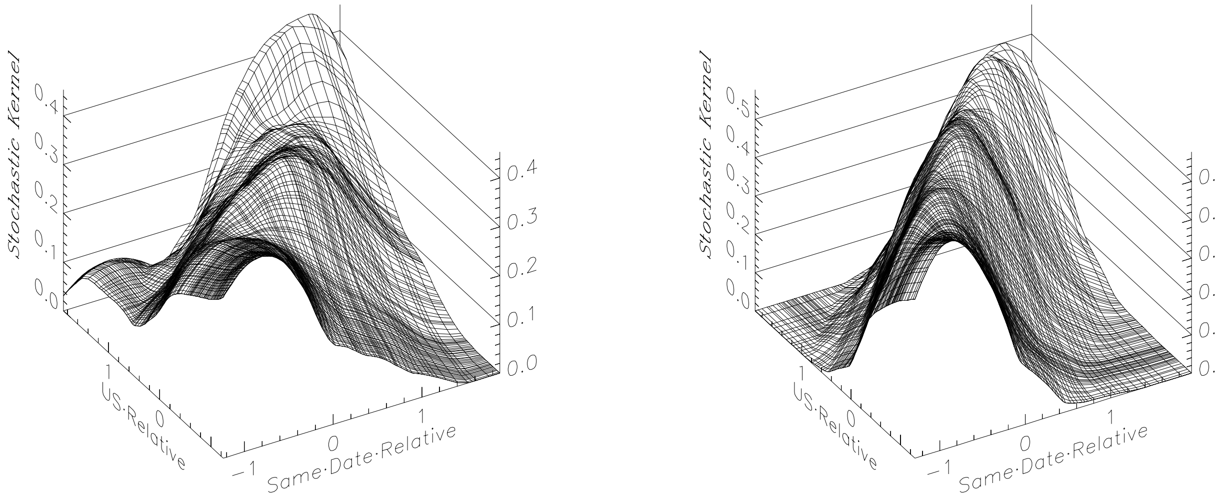

a: 1910 b: 1990

All calculations done using Danny Quah's tSrF econometric shell.
Stochastic kernel from (normalised) population at time t to (normalised) date conditioned population at time t.

Figure 1. Date Conditioning

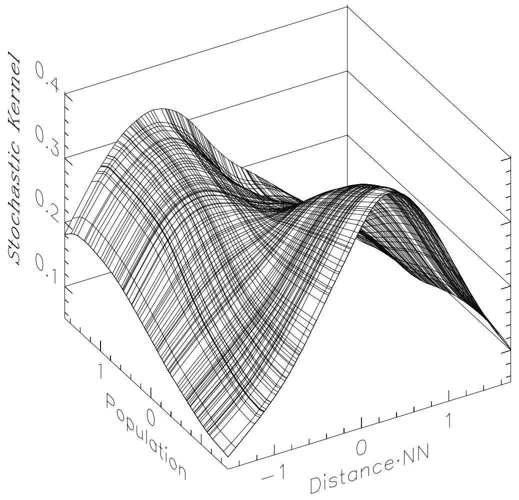

a: 1910

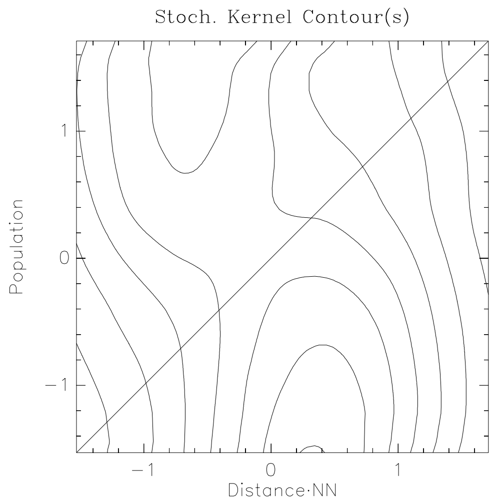

b: 1910

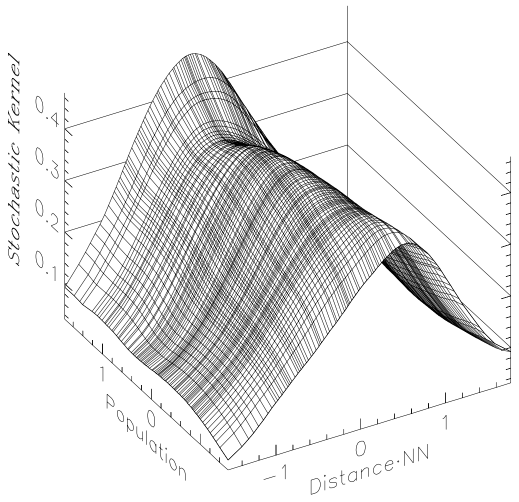

c: 1990

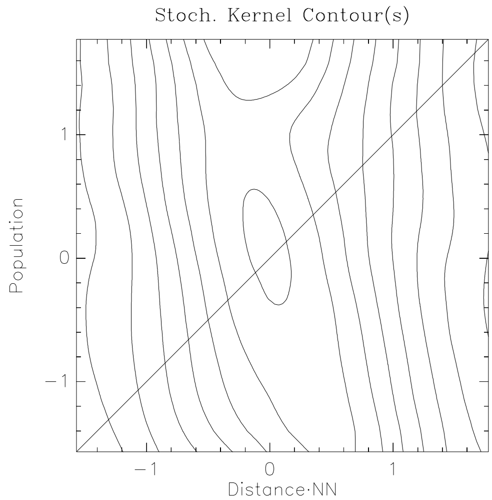

d: 1990

All calculations done using Danny Quah's tSrF econometric shell.
Stochastic kernel from (normalised) population at time *t* to (normalised) distance to nearest neighbor at time *t*.

Figure 2. Population to distance to nearest neighbor

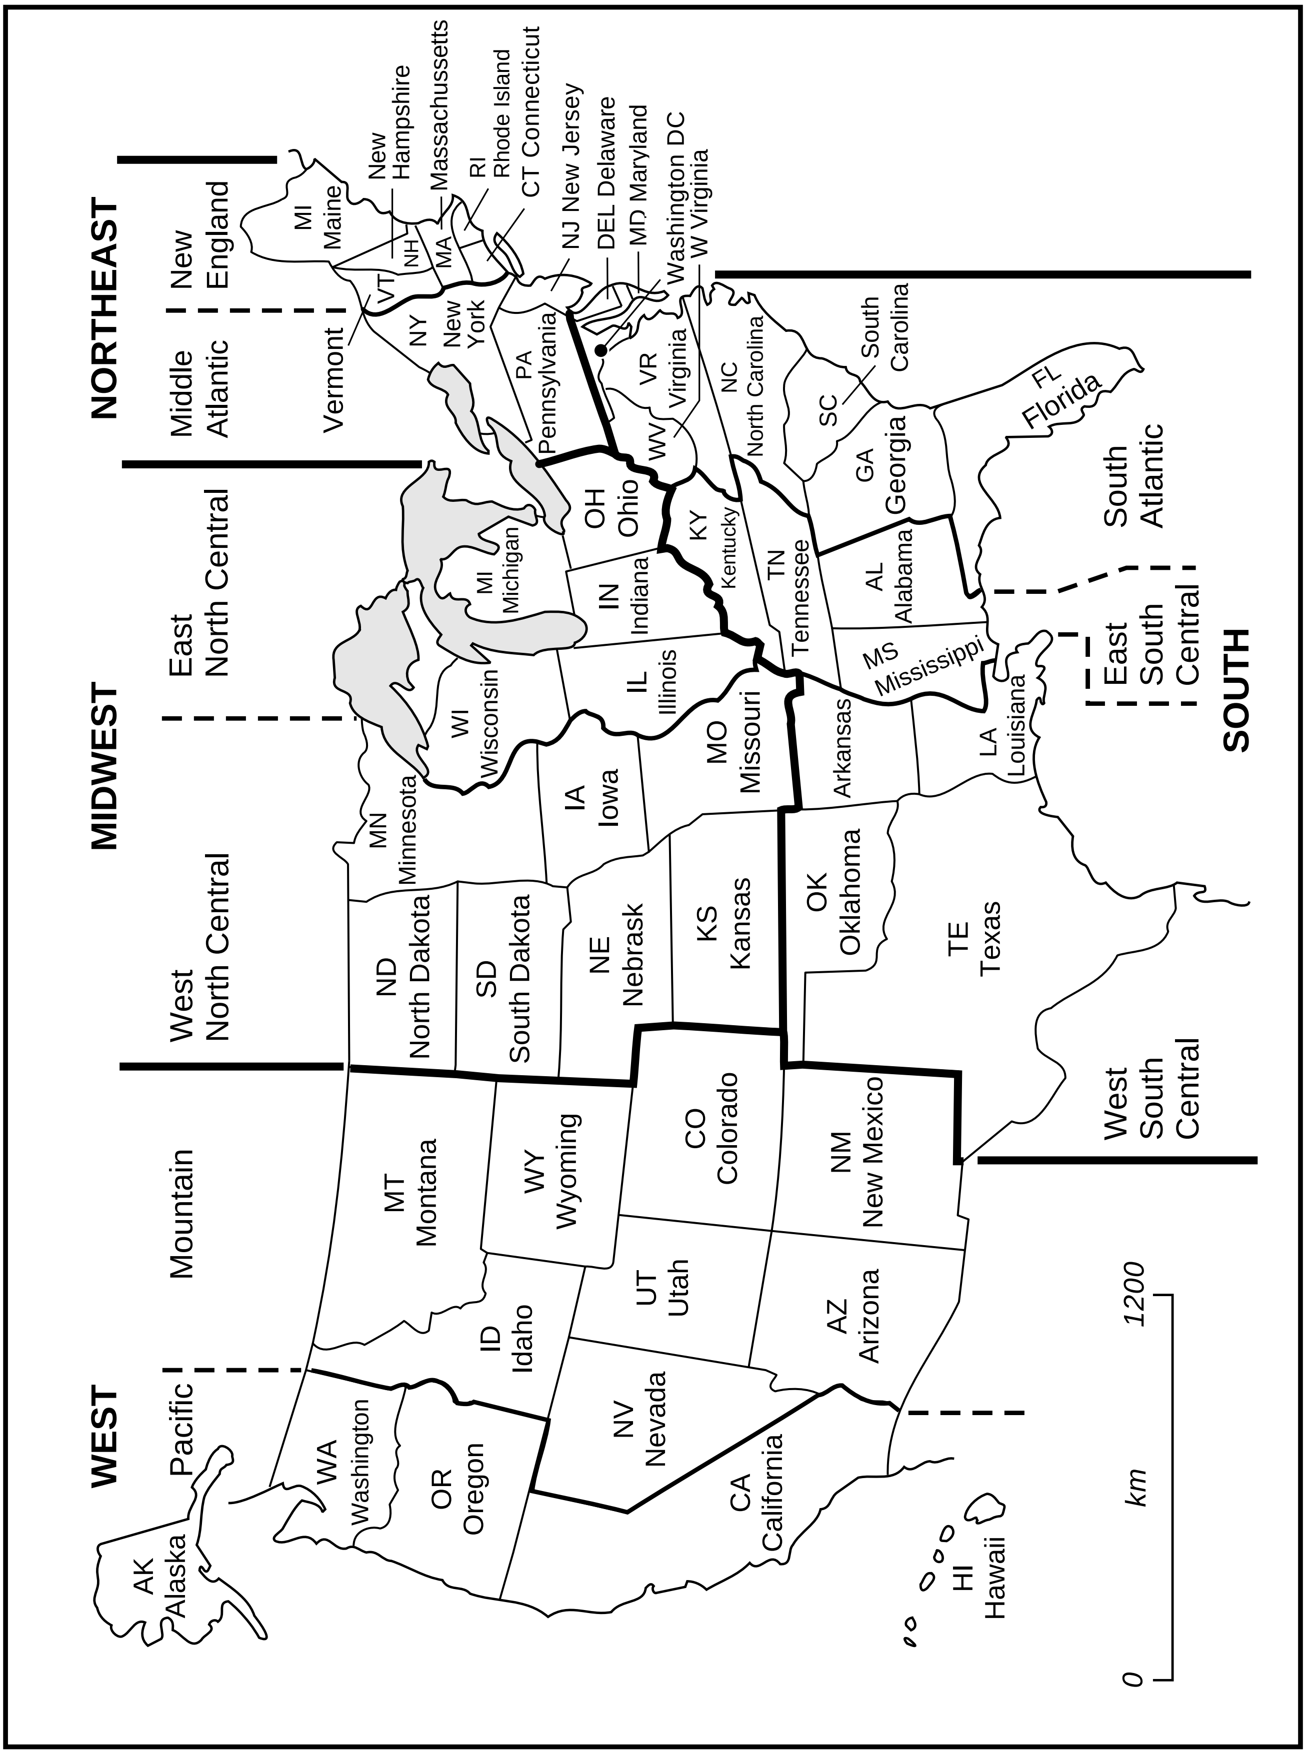

Figure 3. Census regions

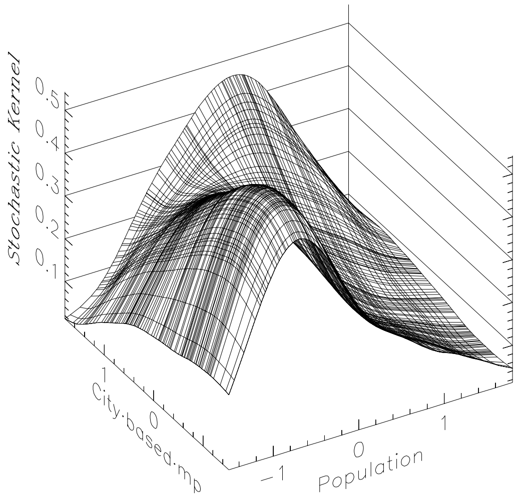

a: 1910

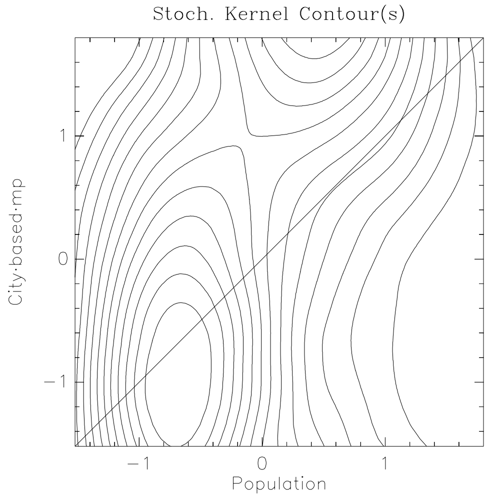

b: 1910

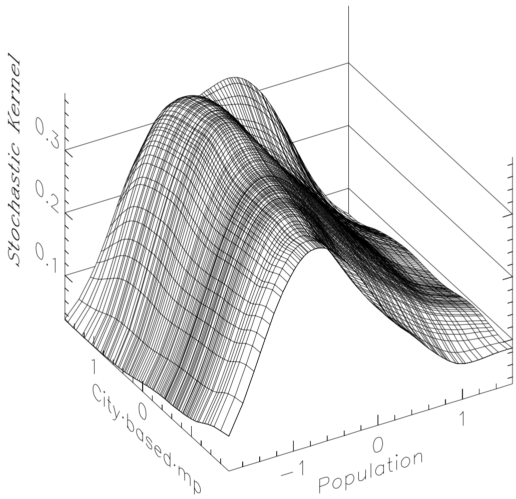

c: 1990

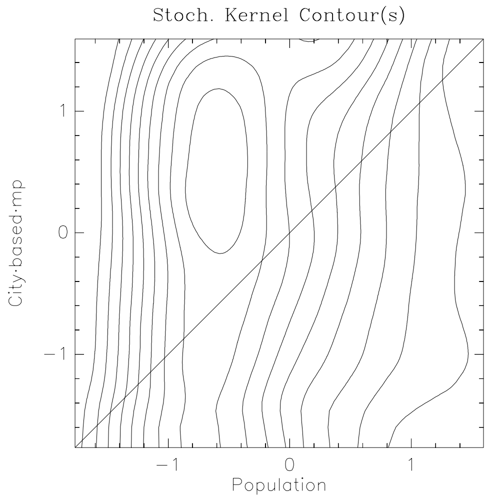

d: 1990

All calculations done using Danny Quah's tSrF econometric shell.
Stochastic kernel from (normalised) city based market potential at time *t* to (normalised) population at time *t*.

Figure 4. Market potential to population

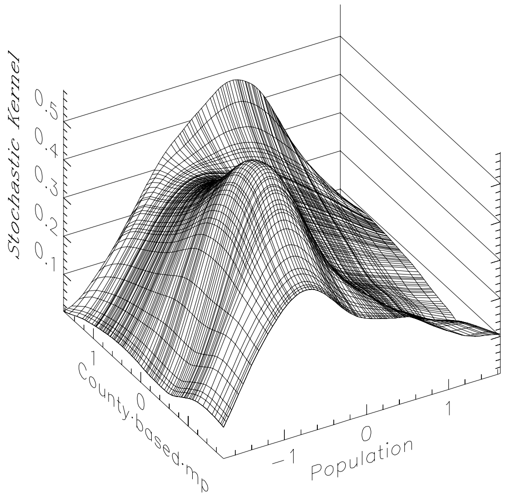

a: 1910

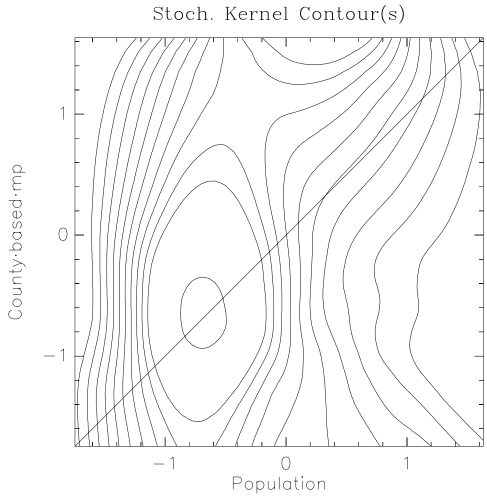

b: 1910

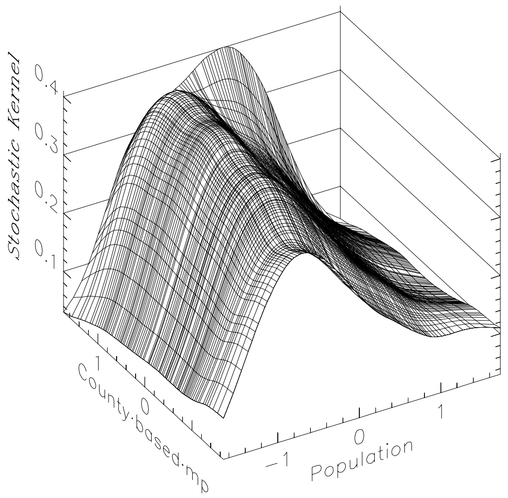

c: 1990

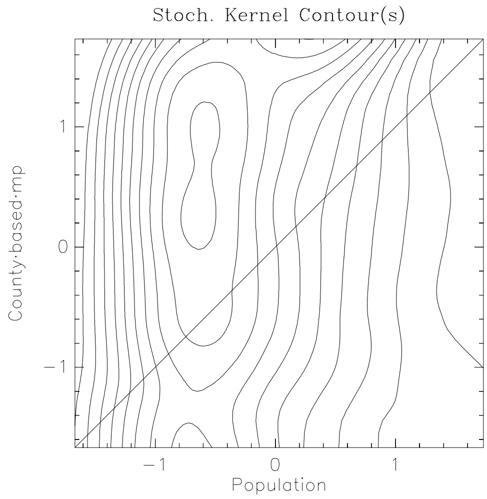

d: 1990

All calculations done using Danny Quah's tSrF econometric shell.
Stochastic kernel from (normalised) county based market potential at time *t* to (normalised) population at time *t*.

Figure 5. Market potential to population

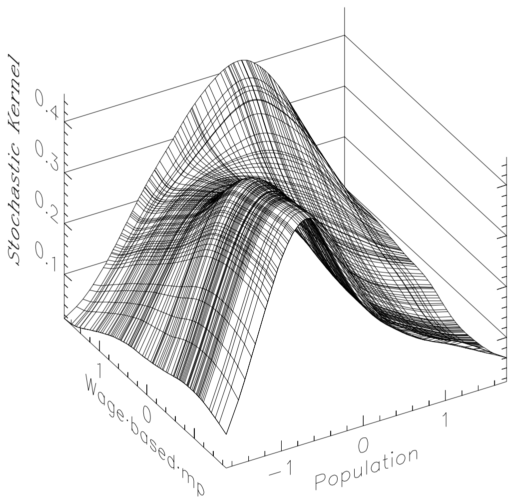

a: 1910

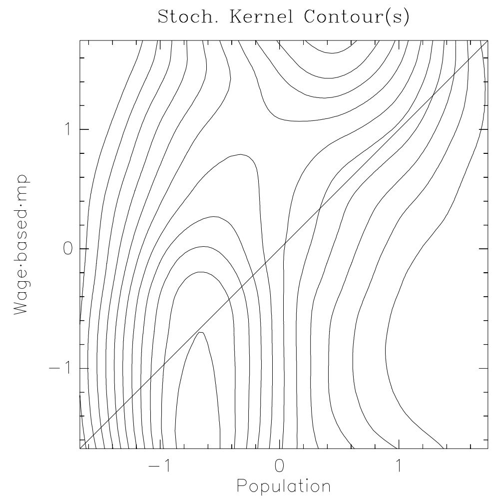

b: 1910

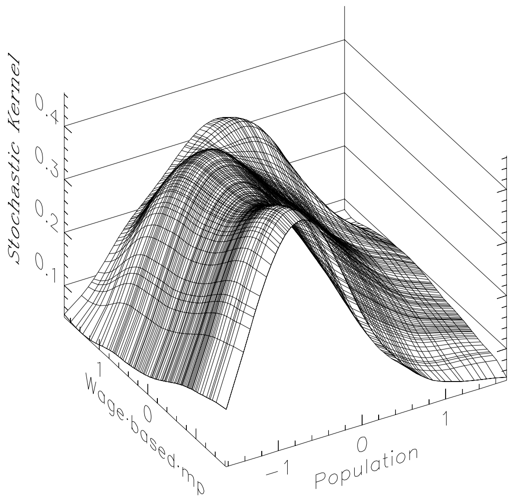

c: 1990

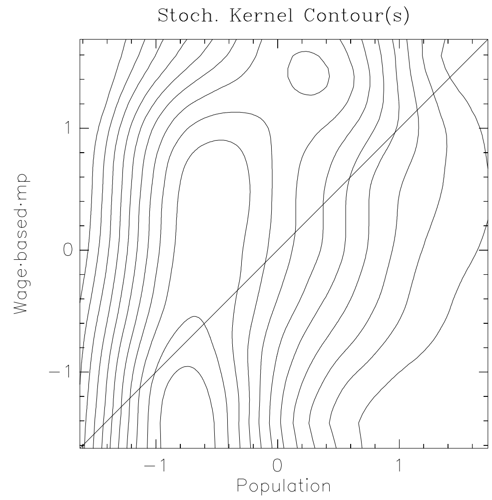

d: 1990

All calculations done using Danny Quah's tSrF econometric shell.
Stochastic kernel from (normalised) wage weighted city based market potential at time *t* to (normalised) population at time *t*.

Figure 6. Market potential to population

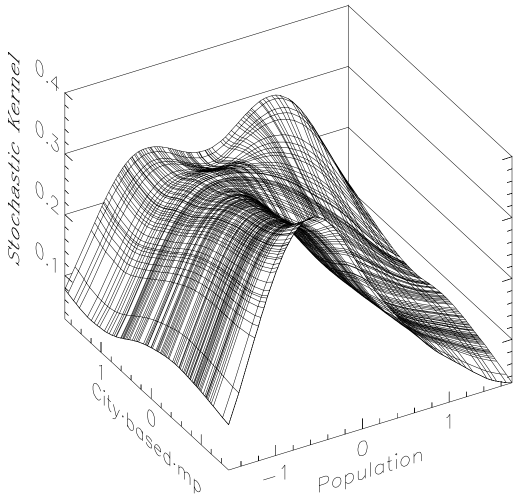

a: 1910

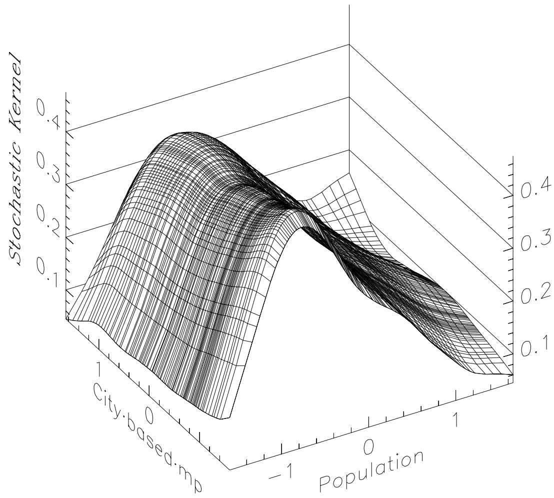

b: 1990

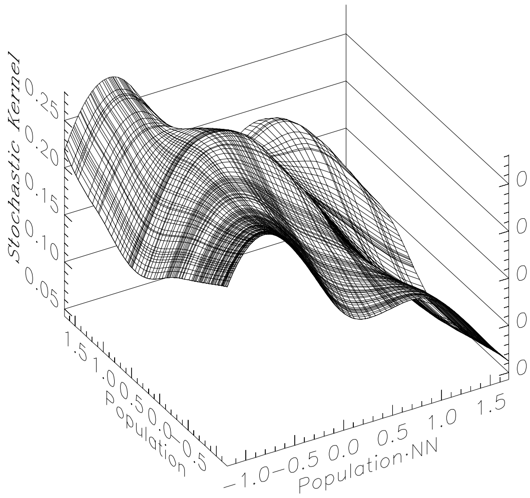

a: 1910

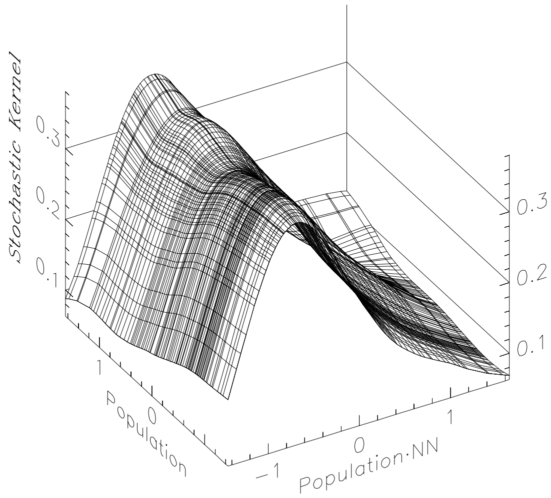

b: 1990

All calculations done using Danny Quah's tSrF econometric shell.
a & b report stochastic kernels from (normalised) city based market potential excluding the market potential component from the nearest neighbor at time *t* to (normalised) population at time *t*.
c & d report stochastic kernels from (normalised) population at time *t* to (normalised) population of nearest neighbor at time *t*.

Figure 7. Nearest neighbor

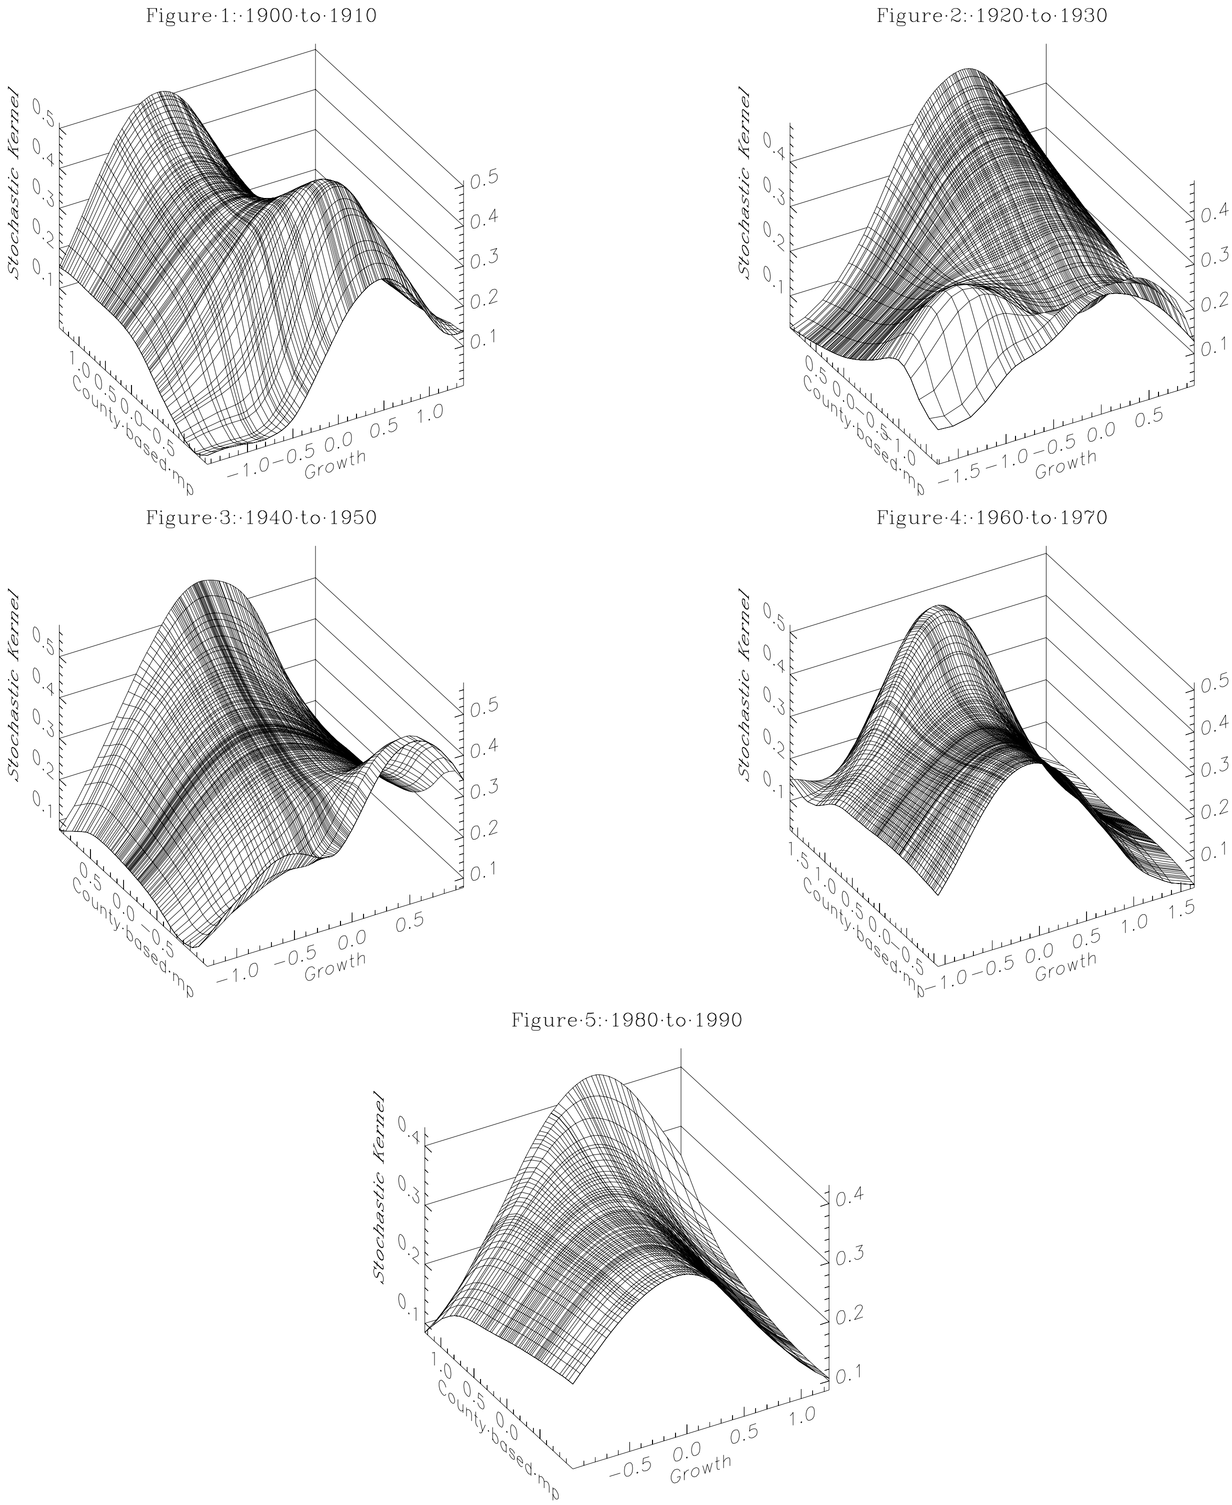

All calculations done using Danny Quah's tSrF econometric shell.
Stochastic kernel from (normalised) county based market potential at time *t* to rate of city growth over the period *t* to *t+1*.

Figure 8. Market potential to growth rates

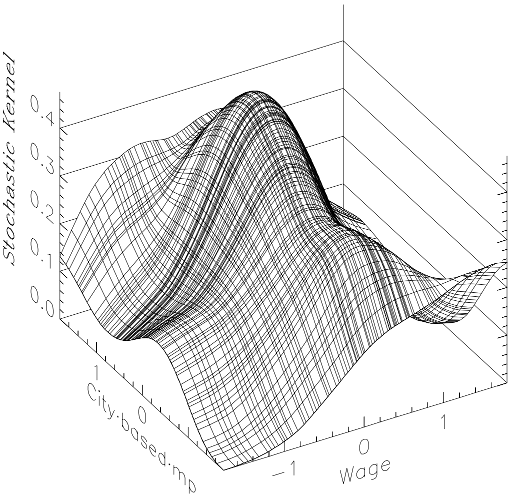

a: 1910

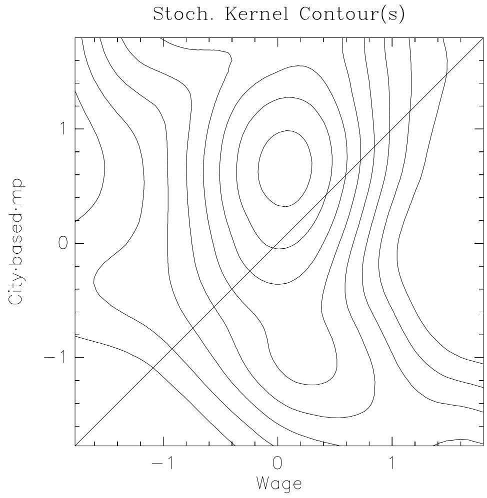

b: 1910

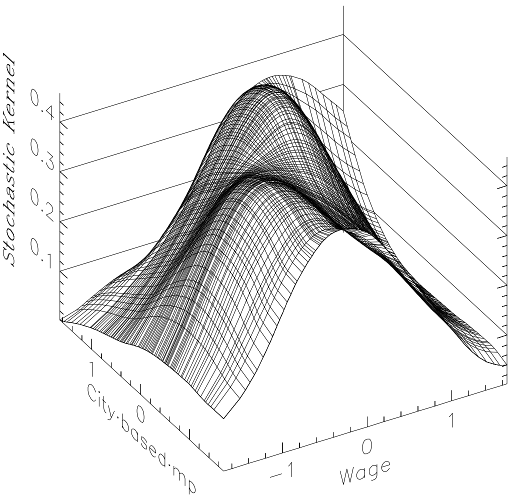

c: 1990

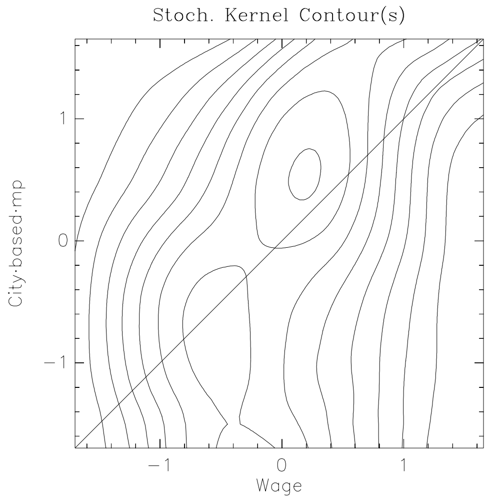

d: 1990

All calculations done using Danny Quah's tSrF econometric shell.
Stochastic kernel from (normalised) city based market potential at time *t* to (normalised) wage at time *t*.

Figure 9. Market potential to wage