Describing Urban Sprawl:
Evidence from remote-sensing imagery

Henry G. Overman*‡
London School of Economics

Diego Puga*§
University of Toronto

Matthew A. Turner*¶
University of Toronto

Work in progress
Printed 20 November 2001

ABSTRACT: To study the evolution of US urban land use, we merge high-altitude photographs from the 1970s/80s with satellite images from the 1990s. The resulting data set allows us to quantify and describe in detail the features of urban sprawl. In combination with other data sources, it also enables us to assess empirically the importance of the diverse potential causes of sprawl advanced by economists, journalists, politicians, and environmentalists.

Key words: urban sprawl, remote-sensing, spatial statistics.
JEL classification: H73, D72, R14.

*We are grateful to Marcy Burchfield and Kent Todd for first-class research assistance.

‡Department of Geography and Environment, London School of Economics, Houghton Street, London WC2A 2AE, United Kingdom, h.g.overman@lse.ac.uk, 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.

§Department of Economics, University of Toronto, 150 St. George Street,Toronto, Ontario M5S 3G7, Canada, d.puga@utoronto.ca, http://dpuga.economics.utoronto.ca. Also affiliated with the Canadian Institute for Advanced Research and the Centre for Economic Policy Research. Funding from the Social Sciences and Humanities Research Council of Canada is gratefully acknowledged.

¶Department of Economics, University of Toronto, 150 St. George Street,Toronto, Ontario M5S 3G7, Canada, mturner@chass.utoronto.ca, http://www.economics.utoronto.ca/mturner. Support from the Social Sciences and Humanities Research Council of Canada and from the National Fellows Program at the Hoover Institution is gratefully acknowledged.

1. Introduction

Economists have long been interested in issues relating to the determinants of urban land use, for good reasons. Land is the largest item in most household budgets, as well as the main independent source of revenue of subnational governments. Land use patterns have received renewed attention recently as the popular press and public policy debates focus on the issue of urban and sub-urban sprawl. Although definitions differ, a necessary condition for 'urban sprawl' to occur is an expansion in the amount of land that is either built-up or paved. However, despite the increased interest, we know remarkably little about patterns of sprawl, we know even less about what may be causing it and have no evidence on whether or not it is desirable. In this paper we tackle the first fundamental gap in our knowledge by measuring the extent of urban sprawl and describing where that sprawl occurs.

Describing and explaining urban sprawl has proved difficult for two reasons. First, because we have no consistent definition of what counts as sprawl. Second, because sprawl is both a macro and micro-spatial phenomena. At the macro level, sprawl may reflect rising populations, interregional migrations, increasing incomes, and changes to commuting technology. At the more micro level differences in climate, geography and local public policy may all impact on the way in which expanding cities develop. Thus, to properly study sprawl we need data where sprawl is clearly defined and provides sufficiently detailed information to capture the micro-spatial determinants of sprawl but has exhaustive coverage to also allow us to allow for macro-spatial determinants.

To develop such a dataset, we merge high-altitude photographs of the continental US from the 1970s/80s with corresponding satellite images from the 1990s. Together, these two images allow us to track changes in land use in the continental United States. These data have two important advantages. First, they allow a direct and precise measure of sprawl by tracking over time the amount of land that is developed — instead of proxying for sprawl with building permits issued or with changes in the land area of administrative or statistical units classified as urban. Second, the data has an extremely fine resolution that allows to track development for each 30m^2 pixel in a regular grid covering the *entire* continental United States. This will allow us to look at local as well as city or region-wide determinants of sprawl, and will help us get around several econometric concerns that we address in more detail below.

The rest of the paper proceeds as follows. In section 2 we describe the data. In section 3 we briefly outline the existing literature. We also consider a number of econometric problems, and suggest how our data will allow us to address them. In section 4 we consider the extent of sprawl and where it is occurring. The final section 5 concludes and outlines our future work.

2. Data

Our land cover data is based on two fine resolution land use/land cover data sets for the continental United States that have recently become available to the research community. The first describes land cover in the late 1970s (1970s data), and the second describes land cover in the early 1990s (1990s data).

The 1970s land-use data is derived mainly from aerial photographs that NASA acquired from high-altitude missions and other sources. While these photographs were collected over the period 1971 to 1982, dates in the late 1970s are most common in the data. The photographs were interpreted by specialists who delineated the boundaries of contiguous areas ('polygons') representing a homogeneous element in the mapping scheme (e.g. a contiguous residential area). Secondary sources from earlier land use maps and field surveys were also used as needed. Since the scale of the images is 1:250,000, they allow the identification of land cover features about the size of a city block, although finer features, such as roads, are also identified.[1] The hand drawn and labelled polygons were digitised and stored using the vector data model.[2]

The source 1970s land-use data are stored in an un-projected format, spatially referenced by a latitude and longitude coordinate system. The study area is described by 450 data files, each of which describes land cover in a single US Geological Survey (USGS) quadrangle (henceforth 'quads'). The spatial extent of each quad is approximately 129km x 194km. The data is available electronically from `http://edcwww.cr.usgs.gov/pub/data/LULC/250K` and detailed documentation can be found in United States Geographical Survey (1986).

The 1990s land-use data set derives mainly from satellite imagery. The base, leaves-off, images were acquired circa 1992 using the Landsat Thematic Mapper satellite to record the spectral reflectance of the earth's surface at a 30m^2 resolution. The images were assembled in mosaics. An unsupervised computer clustering algorithm identified areas or clusters of pixels representing a homogeneous element (e.g. a contiguous residential area) on the basis of a range of spectral values. Land use codes were then assigned to these clusters/classes on the basis of aerial photographs as reference data. Corrections and refinements are made on the basis of other ancillary data layers.[3] The source 1990s land-use data were stored as raster data,[4] using the Albers Conical Equal Area projection.[5] The study area is described by 50 data files, each of which describes land cover in a single state. The data is available electronically from `http://edcwww.cr.usgs.gov/pub/edcuser/vogel/states` and detailed information can be found at `http://landcover.usgs.gov/natllandcover.html`

In order to measure and describe urban sprawl, we first process the two land-use data sets so that they are comparable. The first step of the data processing is to convert the 1970s land-use data to the same projection and data model as the 1990s land-use data. This is done using They are using ESRI's ArcInfo GRID software, and a collection of purposely-written scripts using ARC Macro Language.

[1] The guidelines given to the specialists who interpreted the photographs specify minimum area of a polygon at 10 acres or 4 hectares with a minimum width of 200m. An exception to this width restriction is made for roads and streams, which may be recorded in polygons as narrow as 100m (Loelkes, Howard, Schertz, Lampert, and Miller, 1983).

[2] The vector data model stores real world features (in this case land use/land cover attributes) using geometric objects (in this case polygons), storing the attributes of the features assigned to each geometric object in a table and spatially referencing the features with an x-y coordinate system.

[3] These include leaves-on Landsat Thematic Mapper satellite images, USGS 3-arc second Digital Terrain Elevation Data and derived slope, aspect and shaded relief, US Bureau of the Census population and housing density data, USGS land use and land cover, and National Wetlands Inventory data where available.

[4] The raster data model stores data as an image which is made up of pixels that are associated with a cartesian coordinate.

[5] This projection uses NAD83 datum and GRS80 spheroid.

The resulting output provides the following information for each $30m^2$ pixel: (1) year of creation for the 1970s land-use data, (2) x-coordinate, (3) y-coordinate, (4) unique-identifier for $990m^2$ grid (see below), (5) 1980 census tract in which the pixel lies, (6) 1990 census tract in which the pixel lies, (7) 101^{st} congressional district in which the pixel lies, (8) 1970s land use, (9) 1990s land use. The land use classifications differ between the two data sets. The data appendix supplies a full list of land use classifications for both the 1970s and 1990s data and explains how we reconcile the two classifications. The census tract and congressional district information will allow us to subsequently match the land-use data with data from the United States Census of Population (and other sources), and to identify the location of various jurisdictional boundaries.

At the time of writing, we only have finished processing the data for the south-east region of the US. Already, the processed land use data is approximately 60 Gigabytes. Once completed, the data is expected to have a size of about 520 Gigabytes. While it is feasible (barely) to calculate simple summary statistics for this much data, more sophisticated econometric analysis is not practical. To facilitate such analysis we aggregate our $30m^2$ pixels (henceforth 'small' pixels) up to $990m^2$ pixels (henceforth 'big' pixels). For each such big pixel we know how many of the 1089 associated small pixels are in each land cover category in each of the two periods, as well as how many of the 1089 associated small pixels change between each pair of land cover categories. We also know how many small pixels are in each intersecting census tract and congressional district. Our strategy of calculating changes in land use at the level of small pixels and then aggregating them by big pixel, as opposed to calculating changes by comparing the aggregated land uses, is important to preserve the fine resolution of our data.[6] We therefore accomplish a nearly thousand-fold reduction in our storage requirements, while sacrificing little of our information about the micro-spatial structure of land use.

After aggregating our land cover data to big pixels we augment our land cover data with the following information about each big pixel: distance to the Atlantic and Pacific oceans, distance to the Gulf of Mexico, distance to the Mississippi river, distance to the nearest of the Great Lakes, distance to the nearest interstate highway, distance to the nearest intersection of interstate highways, and finally, distance to the nearest business loop interstate highway (i.e. any interstate assigned a three digit number). We are also able to calculate the properties of each pixels neighbours, for example, the share of developed area in a ring of radius x kilometres around each big pixel.

The map in Figure 1 illustrates the nature of our data. It represents new development in the South-Eastern United States between the mid-1970s and 1992, as well as previous development, undeveloped area, and water features. Sprawl is particularly noticeable in the map around Atlanta in much of Florida. In the rest of this paper we examine the patterns of urban sprawl in this region in more detail.

[6] To see this, imagine a situation in which one half of a particular big pixel was classified as forest in the first period and as pasture in the second period, and the other half of this big pixel was classified as pasture in the first period and as urban in the second period. Aggregating land use categories to the big pixel level and then time-differentiating them would make it appear as if the forested area had become urbanised. On the other hand, our chosen procedure of time-differentiating land use categories at the small pixel level and then aggregating the changes to the big pixel level would show that the forested area became pasture and what was previously pasture became urbanised. The only information lost in the aggregation is the geographical location within each big pixel of the small pixels experiencing each change. But the resolution of the changes after the aggregation remains $30m^2$.

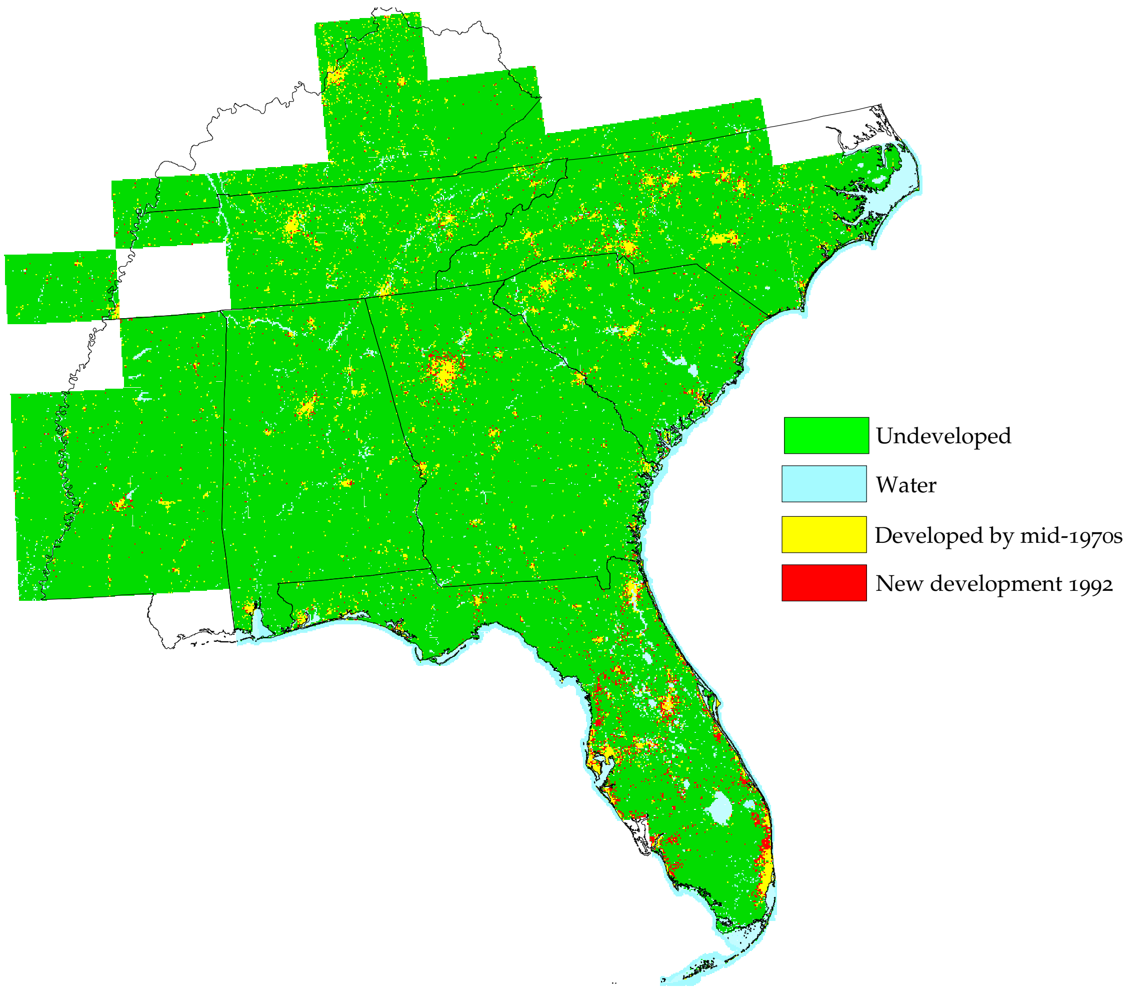

Figure 1. Sprawl in the South-Eastern United States

We note that both land-use data sets have been carefully revised by USGS and US Environmental Protection Agency (USEPA) to eliminate errors. Several studies have found that these data identify asphalt and concrete with an accuracy in excess of 95%. However, there are a few corrections that we have to make due to minor inconsistencies between the two data sets. While most of these problems are minor, and are described in the data appendix, one is important enough to describe in detail.

While well suited describing new development in previously undeveloped areas, the data are not well suited to describing infilling or undevelopment. To illustrate this problem, consider a 10 acre polygon filled with identical small houses, each of which sits exactly in the middle of a $90m^2$ yard (i.e. 3 small pixels by 3 small pixels). Each house occupies 100% of the central pixel in its yard, while the remainder of each yard is undeveloped. The 1970's data is intended to record *land use*, therefore land with more than 10% urban cover is classed as urban (Loelkes *et al.*, 1983). Thus, the whole 10 acre polygon is assigned an urban code. The 1990's data however, is intended to describe *land cover*, consequently it assigns codes on the basis of the preponderance of cover within the pixel. Thus, only one ninth of the area is assigned a developed code, and we record a false undevelopment in our data, even if a fraction of the yards are subdivided and many more houses built. This means that our data is not suitable for tracking either undevelopment or infilling of highly developed areas. On the other hand, if the polygon were initially entirely undeveloped, then the 1990's data would correctly reflect any increase in developed area. For this reason, our analysis addresses only patterns of new development. Since we are interested in explaining the incidence of urban sprawl, which is associated with new development in previously undeveloped areas, we do not regard this as an important handicap.

There is a large literature on the formation and development of cities. This literature is typically concerned with one of two inquiries. The first of these is "Why do cities grow?"(e.g.) In this literature, the unit of study is typically a city, either an MSA or occasionally an Urbanized Area. The objective is to explain the growth of city employment or population. In general, papers in this literature make no effort to explain the arrangement of people and structures within a city. The second sort of inquiry examine the structure of cities. This question requires sub-city level data, and researchers have used many different sorts of data.(e.g.) These studies typically take city level population or employment growth as exogenous (unlike papers interested in the first inquiry) and often make use of small samples of cities.

This literature is far too large to survey exhaustively here, and for the most part has only tangential bearing on our inquiry. However, even the cursory description above make clear the paucity of our knowledge about patterns of urban development.

3. The advantages of remote sensing data

The literature has tried to describe patterns of development by examining city level data in which the area of the city is measured as the area of an administrative unit containing the city. Since the borders of such administrative units are intended to reflect administrative units, rather than the extent of development, and since the borders of these administrative units are often used to determine allocation of state and federal infrastructure funding, such borders provide, at best, an imperfect measure of the amount of development occurring within the administrative area, and no information at all about development beyond the border of the administrative area. Our data resolves both of these problems. We observe developed area directly, for the entire continental US, independent of administrative borders.

The literature has also used a variety of data sources to track development within a city (e.g. housing starts by census block). While these papers do better at measuring the extent of development, the difficulty of collecting such data has limited authors to the examination of a single city or a small region. Our data also overcomes this problem. Our measure of development is at least as accurate as what can be obtained from information about housing starts, and is universal.

Moreover, the fact our data is arranged on a regular and exogenous grid relieves us of two important econometric problems associated with using data collected by administrative units or census tracts. First, the borders of administrative units may not exogenous. They may be drawn where people live or are expected to live. Second, the area of a census tract (or zip code) varies systematically with its population density. Therefore, if we rely on census data the resolution with which we measure population density varies systematically with this density.

The regular-grid nature of our data also helps us overcomes important computational limitations. Computing distances to highways and other features for each census tract or zip code in the whole US as well as distances between every census tract or zip code pair would be an unfeasible task. Because we work with units that are small regular squares instead, we can easily calculate distances just by counting squares. Since we expect proximity to various feature to be essential to understand sprawl, we would not be able to develop this project using administrative data.

4. Some descriptive statistics

Before we can aim to explain urban sprawl, we need to be able to quantify it. Our data allow us to do this directly and precisely, by tracking changes in the development status of every 30m^2 pixel in a regular grid covering the continental United States.

To study changes in land use we need to select pixels that will form our base coverage. These pixels should meet two criteria. First, land use needs to be classified in both time periods. This is necessary if we want to accurately measure changes in land use.. Second, given our emphasis on urban sprawl, it should be feasible to develop these pixels. We deal with each of these issues in turn.

The regular grid that we use to cover the South East region comprises $1,347,831,259$ small pixels (30m^2). Of these $199,069,944$ are small pixels that fall outside the US. This leaves us with $1,148,761,315$ small pixel covering the South east region. Of these 1,121,075,631, or 97.6% were assigned a land use code in the earlier time period. Nearly all the pixels that cannot be classified in the earlier time period are water pixels, reflecting the fact that the satellite data allows for a slightly larger band around the US coastline.

To meet the second criteria, we need to rule out pixels that could not feasibly be developed. At this stage, we will rule out pixels where there are strong physical constraints that prevent development. Given our interest in sprawl, issues of economic feasibility are something that we wish to deal with endogenously. Most important here is to rule out water pixels that are far from land and thus unlikely to be developed. To do this, we use our 990m^2 big pixel constructs. We will assume that it is infeasible to develop any small pixel that is part of a big pixel wholly comprised of water. This should still allow development of ocean and riversides, but tends to rule out large scale 'island like' constructions in the middle of the ocean. Applying this restriction leaves us with a base coverage of 1,100,961,835 small pixels.

For the moment, we will ignore the fact that the year of observation in the earlier period can vary by quad. For the base coverage, 4.0% of the land was developed in 1970 leaving 96% undeveloped[7]. Of those 1,045,499,384 small pix that were undeveloped in the earlier period, 14,762,485 or 1.4% had been developed by 1990. Our 1990s land use codes allow us to classify this urban development according to three different criteria — commercial, low and high density residential[8] When we break

[7]Of the undeveloped land, 35.1% was classified as Agricultural; 53.7% as Forest; 2.1% as Water; 8% as Swamp and 1% as Barren.

[8]Where commercial includes land used for transportation.

down development in to these three categories we see that 63.6% of development has involved the construction of low intensity residential, 12.7% high intensity residential and 23.7% commercial. Comparing these to the total land use in these three categories is informative. In the 1990s total urban land use comprised 58.7% low intensity residential, 16.2% high intensity residential and 25.0% commercial. Two stylised facts emerge from this. First, new urban development is slightly biased towards residential. Second, conditional on that, new urban development is biased towards low intensity residential at the expense of high intensity. To see this, note that the ratio of low intensity to high intensity *development* is 5.0 while the ratio in terms of overall urban land use is 3.6. The contrast would, of course, be even greater if we could classify residential type in the earlier time period. Our initial evidence confirms perceptions of sprawl as reflecting both increasing urban land usage and the density of new development.

So far, we have avoided referring to the rate of development, because we have not yet accounted for the fact that earlier land use codes are assigned quad by quad in different years. (See the data description above for why this is the case). If initial land use codes had all been assigned in 1972, then the aggregate development of 1.4% of undeveloped land would correspond to a development rate of 0.069% per year. Allowing for the variation in the year in which the earlier land use was classified gives us a development rate of 0.076% per year, equivalent to a twenty year development rate of 1.5%.[9]

As we might expect, these aggregate numbers disguise large variations across quads. The lowest annual rate of development is found in Andulusia, Alabama with an annual rate of 0.005% per year (0.1% over twenty years). The highest annual rate is found in Plant City, Florida with an annual rate of 0.62% (13.2% over twenty years). Table 13 shows that differences can also be found across states. The high growth rates for Florida are reflected in the fact that 8 out of the 10 fastest developing quads are in Florida. The other two fast developing quads are James Island, South Carolina (6th fastest) and Atlanta, Georgia (10th fastest). Interestingly these two quads show much higher ratios of low intensity to high intensity residential development, suggesting that Atlanta's position as a classic example of a city suffering from urban sprawl is driven partly by the magnitude of urban development and partly by the type of low density development that is taking place.

5. The next step: explaining sprawl

Having described urban sprawl in the US since the 1970s, we would like to go one step further by trying to better understand what are the driving forces behind sprawl.

Urban economists characterise the spatial growth of cities as being determined mainly by population and income growth and by reductions in commuting costs, as well as by the price of land in alternative uses such as agriculture. The popular press and environmental groups, while paying some attention to those same factors, tend to emphasise public policies as more important determinants of sprawl. In this view, the federal and local governments are subsidising new development in the form of new roads, water and sewer lines, schools, and other public goods

[9]We calculate the development rate by annualising the development rate quad by quad and then weighting by the share of each quad in the total number of undeveloped pixels.

state	% developed 1970s	% undeveloped 1970s	annual development rate
Alabama	2.4	97.6	0.021
Florida	6.0	94.0	0.266
Georgia	3.4	96.6	0.061
Kentucky	5.0	95.0	0.037
Mississippi	1.4	98.6	0.023
North Carolina	6.3	93.7	0.086
South Carolina	4.2	95.8	0.064
Tennessee	4.5	95.5	0.040

Table 1. Annual development rates

and services. This research will study empirically whether each of these and other factors favour or deter sprawl, to what extent, and how different factors interact with each other.

This exercise will provide several important benefits. First, the ability to quantify the extent of urban sprawl. Second, a careful assessment of the importance of the conventional determinants of urban growth studied in theoretical urban economics. Third, a contribution to the empirical literature on urban growth, by paying greater attention to the spatial dimension of the urban growth process. Finally, and perhaps most importantly, our results will help inform policy debate. Over 200 sprawl-related smart-growth initiatives were put forward in the 1998 United States elections. Sprawl is no less important in other countries. In Canada, for instance, the population of the Greater Toronto Area is growing faster than that of Los Angeles and most of that growth is taking place in the suburbs, which are growing about twice as fast as the city centre. As this happens 'sprawl' and 'smart-growth' are becoming equally pervasive buzzwords amongst Canadian politicians as amongst US politicians. To have meaningful discussions and an adequate policy design, it is crucial that we know what really affects sprawl.

We intend to examine the economic and political determinants of sprawl in two related but distinct exercises. In the first exercise we will regress our measure of sprawl on a number of explanatory variables. This exercise will look mainly at the economic determinants of sprawl, although some political considerations will be introduced. We have a number of particular issues that we are interested in — the role of socioeconomic characteristics of nearby areas, the relationship between new and old development and whether there is leapfrogging, the role of proximity to roads and sewers, land type and physical geography, agricultural use, amenities, and club goods. Each of these can be studied using our pixel data with much more care than has been possible before.

In the second exercise, we will study jurisdictional boundaries to identify the political determinants of sprawl. These have been hard to identify because it is difficult to separate them from various unobserved characteristics. This is where the fine resolution of our data becomes crucial. Pixels that are on different sides of a jurisdictional boundary will be nearly identical in every respect except in political variables. Thus, studying pixels near jurisdictional boundaries will help us identify the effects of political variables in a way that would not be viable otherwise.

5.1 The economic determinants of sprawl

In studying the economic determinants of sprawl, the general form of the regression that we would like to conduct is

$$\text{Sprawl} = f\,(\text{Explanatory variables}) + \varepsilon.$$

While we do not intend to derive a structural specification for this equation from a theoretical model, key elements of the estimating equation will have much in common with elements of the standard Alonso-Muth framework (Alonso, 1964, Mills, 1967, Muth, 1969, Wheaton, 1974, Brueckner, 1987), and with more recent models of urban growth (Black and Henderson, 1998).

We have tried to anticipate difficult econometric issues in our construction of the data. In particular, having defined our unit of observation to be regular pixels of 990m^2 rather than irregularly-sized administrative units cuts down our endogeneity and heteroskedasticity problems. We also intend to focus on explanatory variables that we can observe before, or early in our study period. Given this, we will begin with OLS regressions. As we refine our list of variables we will consider more sophisticated error structures. For example, we can increase the accuracy of our predictions by using methods specifically developed to predict continuous spatial patterns on the basis of a set of discrete observations (these techniques are described in Cressie, 1993), or by allowing for spatial correlation of errors.

In terms of the variables to be considered as potential determinants of sprawl, urban economists characterise the spatial growth of cities as being determined mainly by population and income growth and by reductions in commuting costs, as well as by the price of land in alternative uses (Brueckner and Fansler, 1983, Mieszkowski and Mills, 1993, Brueckner, 2000). Thus, we will consider **demographic and socioeconomic variables** from the Census of Population and Housing at the Census Tract level. Regarding the **value of land in non-urban uses**, the Census of Agriculture provides a number of relevant variables, such as the average value of farm land and buildings per acre of farmland; and our land-use data provides pixel-level detail distinguishing between different types of agricultural land. Commuting costs are more difficult to measure directly. Yet by superimposing transport maps on our pixel grid, we can calculate at the pixel level measures of **proximity to transport infrastructure**, such as distance to the nearest interstate highway, to the nearest intersection between interstates, to the nearest intersection between interstates and ring roads, and to public transit stations. We will also consider the availability of other infrastructure, and in particular **proximity to sewers and public water**. More generally, **proximity to economic activity and existing development** is likely to matter. Thus we want to measure the distance to, and magnitude of, nearby economic activity. To do this we will describe a series of concentric rings around each pixel and aggregate economic activity in each ring. For example we will consider the relationship between sprawl in a given pixel and 1980 population or employment in different activities within walking distance, within a 20 minute commute, and within a day's drive. By doing the same exercise for developed pixels within each ring, we will be able to answer the question **does new development leapfrog existing development?** (The leapfrogging hypothesis is simply that the closest ring(s) should have little development in them and those further away should have plenty). We will also look at whether the degree of sectoral and functional diversity or specialisation in nearby areas affects growth (Glaeser, Kallal, Scheinkman, and Schleifer, 1992, Henderson, Kuncoro, and Turner, 1995, Quigley, 1998, Duranton and Puga, 2000, 2001).

Whether development happens will depend crucially on how easy it is to build in each pixel, which varies both with physical geography and with policy variables. Similar variables also affect the attractiveness of development. We discuss policy issues below. In terms of **physical characteristics affecting the ease of development and the attractiveness of a location**, we will include data on terrain slope, type of land (from our main land use data set), proximity to the Coast, the Mississippi, and the Great Lakes, as well as weather variables. Local building costs and rental prospects are also likely to be important (Arnott and Lewis, 1979).

Various **amenities** should also affect the rate at which a pixel is developed. Urban regeneration programmes aiming to help people rediscover the city centre (Whyte, 1988), are one of the main tools being proposed to limit sprawl. An attractive city centre, with symphonies, museums, and restaurants, low crime rates and good air quality, is likely to encourage more people to live downtown. In this view, sprawl may be largely the result of a 'flight from blight' (Mieszkowski and Mills, 1993). However, the general equilibrium the effect of such amenities may be to promote rather than deter sprawl: a more attractive city centre will also attract people from other localities, increasing total population and the spatial extent of the city. We would thus like to answer the question **is sprawl partly the result of a 'flight from blight' or of migration into healthier cities?**

We would also like to know whether '**club goods**' like schools, zoning, and racial composition matter. We suspect disentangling these effects will be difficult since they are all related to the same handful of demographic characteristics. A look at the literature confirms this suspicion (Hoxby, 2002, illustrates the difficulty of isolating school district effects, and Bogart, 2000, argues that several different motivations for local policies are observationally equivalent). At a minimum we can take advantage of the fine spatial detail of our data and look at the importance of school district and municipal dummy variables.

5.2 The political determinants of sprawl: Evidence from jurisdictional boundaries

A pervasive view amongst anti-sprawl campaigners is that '[s]prawl is the result of over five decades of subsidies paid for by the American taxpayer' (Sierra Club, 2000, p. 2). The most emphasised are **federal subsidies to infrastructure**. While the locus of housing subsidies has changed from the dwelling unit to the household occupying the dwelling unit (Quigley, 2000), highways, sewers and other infrastructures continue to absorb large sums of federal spending. Also, employment has become increasingly suburbanised (Glaeser, Kolko, and Saiz, 2001), and federal and local subsidies to enterprises are often blamed for this (Sierra Club, 2000). The amount of subsidies and their allocation varies widely across political jurisdictions. Thus, a crucial issue we are interested in is the **link between pork-barrel and patterns of urbanisation**. Politicians differ in their incentives and abilities to provide infrastructure that is enjoyed but not necessarily financed by their constituents. We will look at this issue by relating local development to characteristics of Congress Representatives (obtained from ICPRS and McKibbin, 1997).

More generally, we would like to look at **boundary effects**. Pixels that are on different sides of a jurisdictional boundary will be nearly identical in every respect except in political variables

(Holmes, 1998). Thus, by studying pixels near jurisdictional borders we can identify the effects of political variables without having to worry about controlling for every other characteristic that may matter. As part of this exercise, we will compare gaps in development as we cross the borders for several levels of government. Similarly, we can look at the gap from incorporated to unincorporated land, and at the role of transportation districts (and the importance big city/small city tension). **Zoning and property taxation** are two policy variables of particular interest in this exercise because of their potentially ambiguous effects on sprawl (Pendall, 1999, Brueckner and Kim, 2000). A related issue is the issuance of **permits to dredge or fill wetlands**, which allows developers to assemble large contiguous plots of land where it would otherwise be impossible. Under Section 404 of the federal Clean Water Act, the US Army Corps of Engineers can delegate to a state agency the relevant authority, and states can differ widely in their use of that authority.

Overall, the unique detail of our data will allow us to characterise urban sprawl and its determinants with a detail and rigour that has simply not been possible before.

6. Data appendix

For instance, in the 1970s/80s land-use data, large military installations are classified as fully urbanised even when they contain large undeveloped areas. In the 1990s land-use data the actual land use is captured instead. While these issues are rare in the data, they can account for large areas and need to be taken care of. Thus, the former Badlands Bombing Range, in the Pine Ridge Reservation, South Dakota appears in the source 1970s land-use data as a solidly urbanised rectangle of 2,500 acres.

References

Alonso, William. 1964. *Location and Land Use; Toward a General Theory of Land Rent*. Cambridge, MA: Harvard University Press.

Arnott, Richard J. and Frank D. Lewis. 1979. The transition of land to urban use. *Journal of Political Economy* 87(1): 161–169.

Black, Duncan and J. Vernon Henderson. 1998. Urban evolution in the US. Working Paper 98–21, Brown University.

Bogart, William T. 2000. 'What big teeth you have!': Identifying the motivations for exclusionary zoning. *Urban Studies* 30(10): 1669–1681.

Brueckner, Jan K. 1987. The structure of urban equilibria: A unified treatment of the Muth-Mills model. In Edwin S. Mills (ed.), *Handbook of Regional and Urban Economics*, volume 2. Amsterdam: North-Holland, 821–845.

Brueckner, Jan K. 2000. Urban sprawl: Lessons from Urban Economics. Processed, University of Illinois at Urbana-Champaign.

Brueckner, Jan K. and David A. Fansler. 1983. The economics of urban sprawl: Theory and evidence on the spatial sizes of cities. *Review of Economics and Statistics* 65(3): 479–482.

Brueckner, Jan K. and Hyun-A Kim. 2000. Urban sprawl and the property tax. Processed, University of Illinois at Urbana-Champaign.

Cressie, Noel. 1993. *Statistics for Spatial Data*. Revised edition. New York: John Wiley and Sons.

Duranton, Gilles and Diego Puga. 2000. Diversity and specialisation in cities: Why, where and when does it matter? *Urban Studies* 37(3): 533–555.

Duranton, Gilles and Diego Puga. 2001. From sectoral to functional urban specialisation. Discussion Paper 2971, Centre for Economic Polic Research.

Glaeser, Edward L., Heidi Kallal, José A. Scheinkman, and Andrei Schleifer. 1992. Growth in cities. *Journal of Political Economy* 100(6): 1126–1152.

Glaeser, Edward L., Jed Kolko, and Albert Saiz. 2001. Consumer city. *Journal of Political Economy* 1(1): 27–50.

Henderson, J. Vernon, Ari Kuncoro, and Matt Turner. 1995. Industrial development in cities. *Journal of Political Economy* 103(5): 1067–1090.

Holmes, Thomas J. 1998. The effect of State policies on the location of manufacturing: Evidence from State borders. *Journal of Political Economy* 106(4): 667–705.

Hoxby, Caroline M. 2002. All school finance equalizations are not created equal. *Quarterly Journal of Economics* (forthcoming).

Inter-university Consortium for Political and Social Research and Carroll McKibbin. 1997. *ROSter of United States Congressional Officeholders and Biographical Characteristics of Members of The United States Congress, 1789-1996: Merged Data*. 10th edition. Ann Arbor, MI: Inter-university Consortium for Political and Social Research.

Loelkes, George L., Jr., Gordon E. Howard, Jr., Eddie L. Schertz, Jr., Phillip D. Lampert, and Stephan W. Miller. 1983. Land use/landcover and environmental photointerpretation keys. Bulletin 1600, United States Geological Survey.

Mieszkowski, Peter and Edwin S. Mills. 1993. The causes of metropolitan suburbanization. *Journal of Economic Perspectives* 7(3): 135–147.

Mills, Edwin S. 1967. An aggregative model of resource allocation in a metropolitan area. *American Economic Review Papers and Proceedings* 57(2): 197–210.

Muth, Richard F. 1969. *Cities and Housing*. Chicago: University of Chicago Press.

Pendall, Rolf. 1999. Do land-use controls cause sprawl? *Environment and Planning B* 26(4): 555–571.

Quigley, John M. 1998. Urban diversity and economic growth. *Journal of Economic Perspectives* 12(2): 127–138.

Quigley, John M. 2000. A decent home: Housing policy in perspective. *Brookings Wharton Papers on Urban Affairs* : 53–88.

Sierra Club. 2000. Sprawl costs us all: How your taxes fuel suburban sprawl. Spring 2000 Report.

United States Geographical Survey. 1986. *Land Use Land Cover Digital Data from 1:250,000 and 1:100,000-Scale Maps: Data User Guide 4*. Reston VA: United States Geographical Survey.

Wheaton, William C. 1974. A comparative static analysis of urban spatial structure. *Journal of Economic Theory* 9(2): 223–37.

Whyte, William H. 1988. *City: Rediscovering the Center*. New York: Doubleday.