Point Demand Forecasting

By

Wilpen Gorr
Michael Johnson
Stephen Roehrig

January 17, 2001

H. John Heinz III School of Public Policy and Management
Carnegie Mellon University
Pittsburgh, PA 15213-3890

Acknowledgements

We wish to acknowledge Nancy Zionts of the Jewish Healthcare Foundation for providing funding for the work reported in this paper and for providing access to data and officials. She provided us with a challenging and interesting research project.

Abstract

This paper provides geographic information system (GIS) methods and empirical models to forecast point demand for home-delivered goods. A point forecast consists of stops on a street network, including demand at each stop. The purpose of the forecast is to support a network optimization model, based on the traveling salesman problem, to locate one or more new facilities in a region. We illustrate our approach with a case study of home-delivered meals (meals on wheels) in Allegheny County, Pennsylvania.

Keywords: Point Forecasting, Demand Forecasting, Network Optimization Models, Geographic Information Systems, Demographic Forecasts

1. Introduction

Network-based optimization models are among operations research's most powerful and widely used methods. Examples include: the shortest-path problem, the transportation problem, the traveling salesman problem, the facility location problem, the vehicle routing problem, the location-routing problem, network design and others. These models are the basis for decision support systems that serve many industries and classes of decision problems. In one case, a large company that provides goods and services to home-based customers uses models and algorithms to generate daily delivery routes that account for road networks and technician skills (Wiegel and Cao, 1999). Another large, vertically integrated company uses a combination of models and algorithms to design and refine the entire supply chain: warehouses, break-bulk centers and delivery routes (Camm et al., 1997). Finally, a case study shows how a county health agency could use models and algorithms to locate fixed and satellite facilities to provide preventive and primary health care to spatially dispersed customers (Clarke, Lapierre, and Ratliff, 1997; Lapierre, Ratliff, Goldsman, 1998).

Demand forecasting is essential for implementing many network optimization models. Standard time series methods are applicable to network models needing demand level forecasts by geographic area. An example is forecasting end user electricity demand in the current deregulated retail market. Hourly market price patterns depend on load, which in turn depends on many factors, including weather, time of year, and so forth. Minimizing power costs and risks requires accurate forecasts combined with sophisticated optimization models.

Many network problems, including the last two applications listed above, require point demand forecasts. The operations research/management science literature, while rich in models and algorithms for solving network optimization problems that rely on point demand forecasts, has relatively little to say about how such demand forecasts are developed, and the integral role geographic information systems (GIS) might play in developing these forecasts. Robusté, Daganzo, and Souleyrette II (1990) present a number of approximate heuristics as well as applications of the simulated annealing algorithm to solve the traveling salesman problem and the vehicle routing problem. Laporte, Louveaux and Mercure (1989) and Berger (1997) have presented innovative optimization planning models for simultaneously locating facilities and generating delivery routes.

Daskin, Hopp, and Medina (1992) modify the fixed-charge facility location problem to account for a multiple-period planning horizon in which the length of the planning horizon is endogenous. Current, Ratick, and ReVelle (1997) modify the P-median facility location problem to account for multiple planning periods and allow the total number of facilities located to be endogenous. Hesse Owen (1999) adapts the P-median, set covering and maximal covering facility location models to account for demand uncertainty using scenario analysis. All of these authors rely on either simulated data or standard datasets not associated with any actual applications.

Other OR/MS research in network optimization models makes limited reference to demand forecasting. Daganzo (1987) argues that discontinuous functions, such as costs based on full-truckload shipments, may be approximated with linear functions that result in smaller forecast errors than the original, exact expression, when there is a moderate amount of uncertainty in input parameters. Daganzo (1991) generalizes this argument and presents a number of exact algorithms based on abstractions of real-life logistics problems and argues that the use of approximate models allows key parameters to be more easily forecast, results in forecast errors that are easier to characterize, and enables stochastic effects to be more easily incorporated than in more exact, but harder-to-solve models. However, Daganzo gives few details on how forecasts of relevant data might be performed. Narasimhan, Pirkul, and Schilling (1992) adapt the P-median model to account for multiple backup facilities with fixed service capacities and service levels; they implement their model using simulated data as well as actual data from a major U.S. city. However, the fact that they use one year of historical demand data implies that they use a forecast in which the future is the same as the past, clearly a naïve assumption. Clarke, Lapierre, and Ratliff (1997) and Lapierre, Ratliff, Goldsman (1998) use Census block group data to approximate demand, but, as do Narasimhan, Pirkul, and Schilling (1992), assume that the future is the same as the past and use no other demand forecasting algorithms. Camm et al. (1997) note that their planning models rely on demand forecasts to design the supply chain but say nothing about how such forecasts are generated.

Common sense tells us that implementation of the models described above depends on good forecasts of relevant parameters such as demand. Indeed, Camm et al. (1997) note that model inaccuracy is associated with forecast demands and costs and that the model is “only as good as the forecast used” (p. 138). Advances in simulation modeling and GIS have resulted in impressive power at the desktop PC level to generate forecasts of key parameters. Moreover, it is possible that good quality forecasts could enable the design of more precise model solution methods than those that are calibrated using randomly generated data. But how can these forecasts be developed? In this paper, we study the problem of forecasting point demand for home-delivered goods from multiple facilities within a region. The point forecasts are stops (points) on a street network, each with a particular level of demand. Such forecasts are needed to choose locations and capacities for new distribution facilities. We have used these demand forecasts in an interactive network optimization model that constructs the catchment area served by each facility and the associated delivery routes, subject to network, demand and capacity constraints (Gorr, Johnson and Roehrig 2000).

Many Web-based delivery services require similar optimization models supported by point forecasts for facility location. For example, WebVan, NetGrocer, and other competing firms in the so-called instant delivery sector must somehow estimate future demand in order to properly locate delivery depots and warehouses. The retail energy market mentioned above is now discovering that a changing customer base, volatile market prices, and actively modified load shapes make obsolete the traditional customer class aggregate forecast models used by regulated utilities. Forecast of point demand for energy consumption is required.

Other researchers are starting to place emphasis on point data. For example Pace et al. (1999, p. 229) state that "… the increasing capabilities of information systems and especially geographic information systems (GIS) have greatly aided work with disaggregated data having precise spatial and temporal references." Furthermore, they state (p. 230) "The trend towards large data sets with substantial spatial and temporal detail raises the issue of how to forecast such data." Lewis-Beck and Tien (1999) discuss macro and micro forecast models in the context of forecasting presidential elections. Macro modelers use aggregate-level data and national-level time series data. Micro modelers use data from individual polls. The advantages of the latter are more detailed and current information.

Geographic information technology and infrastructure enable point demand forecasting. Geographic information systems (GIS) provide unique spatial processing algorithms for use with demographic data and street networks: spatial join, geocoding, and reverse geocoding. Governments and commercial vendors provide base maps, including street networks, and small-area demographic forecasts. This technology allows us to build demand models that include demographic factors to condition forecasts as well as to remove certain estimation biases. The end result is a simulation that generates point forecasts according to demographic trends and consumption patterns.

The second section of this paper reviews GIS technologies for making point forecasts. The third section provides a brief overview of the case study used in this paper to illustrate the point forecasting problem. The case study is on home-delivered meals (HDM) for the elderly, also known as meals on wheels. The policy problem addressed is to conduct a “gaps analysis” of the 63 HDM distribution facilities in Allegheny County, Pennsylvania; that is, to identify areas of the county in which the elderly population does not have access to home-delivered meals, and to estimate the number and locations of additional facilities needed. The fourth section uses a sample of facilities and their clients to calibrate a demand model. A principal use of the model is to remove an estimation bias caused by overlapping service areas. The fifth section describes our simulation model and forecast results. The final section is a discussion of future research.

2. Geographic Information Systems

GIS has become a mainstream information technology in the past decade. In the U.S., a key GIS infrastructure is the TIGER (Topologically Integrated and Geographically Encoded Referencing system) base maps available for the entire country and its possessions. These maps include street centerlines, political boundaries (states, counties, municipalities), statistical areas (census tracts and block groups), administrative areas (postal areas – zip codes), and other features. Census tracts are neighborhood-like geographic areas of approximately 4,000 population. Census block groups partition census tracts into the smallest census reporting areas, approximately 1,000 persons each. Hence, a geographic region with approximately 1,000,000 population will have 1,000 block groups – a small-grained areal unit suitable for planning and estimation.

The TIGER street centerlines consist primarily of block-long street segments, each with street name, direction, type, and beginning and ending house numbers on the left and right sides of the segment. For example, for the starting node of a street segment (100-block of Oak ST W) there are latitude/longitude coordinates and two house numbers; for example, 100 and 101. For the ending node, there are coordinates and two additional house numbers, 198 and 199. A GIS locates an address, such as 123 Oak ST W by linear interpolation between the starting and ending nodes. Vendors have added value to TIGER street centerline maps by increasing their completeness and accuracy, and by adding attributes useful for network modeling. Among the latter enhancements are 1) average travel time calculated using street segment length and average travel speeds by street type and 2) turn tables that encode allowable turns and average times to make turns. While in many ways crude, such street maps represent the state-of-art infrastructure for network modeling.

A major GIS function is matching street addresses in administrative databases (e.g., facility locations and clients' residences) with TIGER street attributes. The result is that an administrative address can be plotted as an approximate point location on a street centerline segment (123 Oak ST W). This process, known as geocoding or address matching, uses sophisticated address parsing, fuzzy matching using Soundex indices, and a rule base for scoring potential, but imperfect, matches. The candidate street segment that has the highest score above a threshold is taken as the match. In regions that have potential for duplicate street addresses, it is important to have a tie-breaking attribute such as postal code or municipality name.

Reverse geocoding is a GIS function that finds the nearest street address for a given point, whether it is on the street network or not. Another GIS function generates a random point in a given polygon. Together, then, the two functions can be used to generate area-based random points on a street network. An alternative is to randomly select street segments and points on them within a geographic area.

Demographic data, such as census data, are key to demand forecasting. The U.S. census data for 2000 are becoming available, and thus this is a fortunate time window for planners. Commercial vendors provide current-year census estimates and forecasts for census tracts and block groups. Vendors use demographic modeling (birth and death processes, migration) and current year indicator data (e.g., motor vehicle registrations) as the basis for estimation and forecasting. We use current-year and five-year-ahead forecasts by census block group for persons aged 64 and older for the HDM problem.

Another key GIS function is the spatial join or overlay. An area in a vector GIS is represented as a polygon. Input are two, overlapping polygon maps; for example a census block group map and a land use zoning map for Allegheny County, Pennsylvania. The output is the set of all polygons formed by overlaying the two maps. Each resulting polygon has all attributes of both contributing input maps. We use such a map to apportion attributes from one map to the other. For example, we can use block group census data, estimates, or forecasts to estimate corresponding census attributes for zoned areas (residential, commercial, and industrial polygons). The basis of apportionment can be shared area, length of street network, or number of telephones (from digital directories that have been address matched).

3. Home-Delivered Meals Case Study

Home-Delivered Meals (HDM), or Meals on Wheels, is a volunteer-run service with Federal and other funding in the U.S. It has the purpose of providing daily hot meals to home-bound elders, allowing elderly persons to remain in their homes (which is one of the highest priorities of the elderly). For example, Allegheny County, Pennsylvania has a total population of 1.29 million, and 0.23 million of these are 64 or older. The county has 63 HDM distribution centers ("kitchens") which are primarily pre-existing facilities – churches and schools. The average kitchen has 64 clients and four delivery routes. In 1999, 4,017 elders were HDM clients. The net benefits of HDM in Allegheny County are on the order of $50 million per year, in avoided nursing home or other residential facility costs.

In the past, there has been little planning for HDM facility locations. Several organizations have independently opened HDM kitchens, without a central plan. Moreover, the demographics of the elderly in the U.S. have been changing, with increasing numbers in total, but with the increases occurring mostly in suburban communities. At the same time, the elderly are living longer but with the infirmities of increased age, and the availability of caregivers' time is decreasing as the number of dual-income families increase, etc. Hence local government agencies and grant making foundations that provide services for the elderly have turned attention to looking for gaps in the availability of HDM to target populations, as HDM has become increasingly important.

There are several factors that influence the size and performance of an HDM facility's catchment area:

•*Location of facility* – ideally, a facility should be located in densely populated areas, centered among clients. Given a gap in client coverage, it is possible to size a facility and number of delivery routes to meet demand.

•*Capacity of a facility* – the larger the capacity (meals prepared per day), the larger the catchment area.

•*Time limit per route* – kitchens use a 45-minute time limit for delivery of the last meal on a route, to maintain the temperature of hot meals. This is a requirement for kitchens that have Federal funding.

•*Number of routes* – the higher the number of routes, the larger the catchment area. The number of routes is a decision variable that needs to be determined to meet demand within the time limit constraint.

•*Density of clients* – the more densely populated the client population, the smaller the catchment area.

•*Street network* – the more dense the street network, the larger the catchment area. A dense street network provides more route alternatives and shorter travel times between stops.

•*Barriers* – elimination of overlapping catchment areas and relocation of kitchens reduces delivery costs by increasing stops density within a catchment area. One approach to building non-overlapping catchment areas is to erect pre-solution barriers between competing facilities.

The street network automatically implements physical barriers, such as lack of bridges across a river or limited access to an interstate highway.

After attempting several area-based approaches to accounting for all of these factors, such as using circular buffers around facilities, we found that the only viable approach was to use point demand forecasts. The problem with circular buffers and similar approaches is that a constant areal density of clients must be assumed; this is hardly reasonable given the nature of street networks and the location of housing in neighborhoods.

Given an estimate of point demand (stops), we consider 1) a proposed facility, 2) a proposed number of routes, and 3) a forecasted demand level at each stop (80 percent of all stops have one client, 15 percent two, and the balance three or more). We then use a routing algorithm based on the traveling salesman problem to maximize the catchment area of a given facility (see Gorr, Johnson, and Roehrig, 2000). Each factor discussed above is handled explicitly, so that no compromises need be made. Natural (e.g., rivers) and man-made (e.g., interstate highway) barriers are encoded in the street network. Additional barriers separating kitchens can be inserted manually or under software control, allowing HDM decision makers to simulate new policies. The result is not only an accurate estimate of the potential size of a catchment area (given the forecast), but a way of actually realizing it (i.e., a concrete set of routes).

4. Point Demand Forecasting

We conducted a census of all 63 kitchens in Summer 1999 to obtain needed data: kitchen name, agency name, street address, zip code, number of routes, number of clients, number of meals per day, and maximum capacity of the kitchen (meals per day). Through repeated attempts, we were also able to obtain a sample of 25 kitchens' client addresses (stops), route number, and clients per stop. While not a random sample, comparisons of income, geographic area, etc. of the sample with the population indicated that we obtained a representative sample.

We address matched all 63 kitchens, and then address matched the 25 sets of stops. Exhibit 1 is a map of Allegheny County, Pennsylvania showing the locations of kitchens. The point markers for kitchens are encoded as to whether or not we had obtained addresses for clients' stops. Also shown are the 25 catchment areas for kitchens that we digitized around client stops (more is said about these catchments below). There appear to be obvious gaps in coverage, even with this limited information.

Exhibit 2 is zoomed-in map, showing the Northside Kitchen, its stops serving 70 clients, and the polygon we digitized representing the Northside catchment area. Also, shown is a second kitchen, one for which we do not have stops data. The coterminous polygons in Exhibit 2 are census block groups, on which we have displayed 1997 estimates for the elderly population using a monochromatic color scale. We used area apportionment to estimate the total elderly population in the catchment to be 2,905. We thus estimate a utilization rate of 0.0241 = 70 clients/2,905 elderly population. Our objective is to use the 25 such ratios to model or estimate utilization rate. Exhibit 3 is the resulting frequency distribution, with a range of 0.003 to 0.042, excluding two high outliers that have retirement communities.

Referring back to Exhibit 1, note that even with only catchments for 25 out of 63 kitchens, there are several kitchens with overlapping catchment areas. Catchment overlapping is a common phenomenon, and one that causes an underestimation bias in utilization rates as calculated above. Note in Exhibit 2 the proximity of the second kitchen to the Northside Kitchen. It is very likely that the two kitchens have overlapping catchment areas. If so, then we have undercounted the number of clients in the Northside Kitchen's catchment area; namely, those clients served by the second kitchen.

An approximate solution for this problem is available using an indicator variable and corresponding utilization rate model. We placed two-mile radius buffers around each kitchen, as seen in Exhibit 4. For each kitchen, we then counted the number of overlapping buffers, giving rise to our indicator variable, which we called NearKitchens, which ranged from 0 to 6. We expected that the higher NearKitchens, the lower the utilization rate as we had estimated it. To correct corresponding underestimation bias, we set NearKitchens to zero in an estimated model for utilization rate.

Other factors thought to influence utilization rates included per capita income (PCIncome) and density of elderly population per square mile (EldPopDensity). We expected that the higher per capita income, the lower utilization rate as wealthier elders may have other alternatives to HDM. Also, we expected that the higher elderly density population, the higher the utilization rate. Diffusion of information about the availability and access to HDM would be more rapid in such areas. Of the three independent variables, the pair highly correlated was NearKitchens and EldPopDen, with a simple correlation of 0.89.

Exhibit 5 provides results of regressing utilization rate on the independent variables, including an intercept term. Only the intercept is significant, with the coefficient for NearKitchens nearly so, but diminished given the collinearity with EldPopDensity. The important result from Exhibit 5 is that PCIncome is not significant in explaining variations in the utilization rate.

Our next step was to re-estimate the model including only NearKitchens as the independent variable, shown in Exhibit 6. Now both the intercept and coefficient for NearKitchens are highly significant. Each nearby kitchen accounts for an underestimate of –0.0025, or 11 percent, in the utilization rate of kitchens with no near kitchens as estimated by the intercept, 0.0228. Hence, our estimate of the utilization rate is 0.0228.

To forecast the number of clients per block group we multiply 0.0228 times the demographic forecasts of the elderly population by census block group. We then draw random points within a block group, reverse geocode to the nearest street address, and take a random draw from our distribution of clients per stop. The result is a point forecast. We continue drawing point samples for a block group until we meet the forecasted number of clients. Exhibit 7 is a map showing one such realization for all block groups in Allegheny County. Also shown on he map are estimated catchment areas of the existing 63 kitchens, redeployed with non-overlapping catchment areas and full capacity used. The resultant uncovered areas are forecasted gaps in coverage, some of which were obvious beforehand and others not.

5. Discussion and Conclusion

GIS provides a useful technology and infrastructure for forecasting spatial demand in a region. Furthermore, street networks, as represented in GIS, provide the means to generate point demand forecasts. In turn, point forecasts can be input into network optimization models to forecast performance of a proposed facility, distribution networks, etc.

We three authors all have backgrounds in operations research and have done applied OR work, including the work reported on here. Thus, the question that we are forced to address and continue to address is “Where do the coefficients come from for calibrating OR models?” There is surprisingly little research in the OR or any other literature on forecasting demand for network optimization models. This, we take as an opportunity for forecast researchers to fill a gap in the literature.

Point demand forecasting, as illustrated in this paper, appears to be unavoidable for some network optimization models. There is no other way to account for variation in delivery networks (e.g., the street network) and demand patterns. Additional work needs to be done on the number of realizations needed to account for variation across realizations, and methods to summarize the resultant collections of catchments. Should we average boundaries? How can that be done? Should we use the median or 95th percentile boundary in terms of extent? Realizations may be time dependent; for example, delivery times may vary considerably depending on time of day. Should these be averaged? Should additional forecasting variables be introduced to account for variations, and then rely on a multi-period location model? How do these techniques "scale" across the time dimension? For example, our work with HDM assumed a five-year time frame. Would there be difficulties in "compressing" the time frame to months or days, as might be required by those in the "instant delivery" sector?

Another area of research needed by forecasters is an assessment of forecast accuracy obtained by commercial vendors of small-scale demographic forecasts. We know of no such papers. Municipalities sometimes conduct a census during non-census years, and the resulting data could provide a means of validation.

In summary, we encourage forecasters who also have backgrounds in operations research and management science to consider the forecasting needs of models in these fields. There are no doubt other difficult forecasting problems that require innovative solutions. Also, we expect that additional comparative forecast studies are needed for important classes of OR models and applications.

References

Berger, R.T. (1997). *Location-routing models for distribution system design*, Unpublished dissertation, Evanston, Ill.: Northwestern University, Department of Industrial Engineering and Management Sciences.

Camm, J.D., Chorman, T.E., Dill, F.A., Evans, J.R., Sweeney, D.J. and G.W. Wegryn (1997). Blending OR/MS, judgment, and gis: restructuring P&G's supply chain, *Interfaces* **27**(1): 128 – 142.

Clarke, L.W., Lapierre, S.D. and H.D. Ratliff (1997). "Improvement heuristics to solve service location problems." Atlanta: Georgia Institute of Technology, School of Industrial and Systems Engineering.

Current, J., Ratick, S. and C. ReVelle (1997). Dynamic facility location when the total number of facilities is uncertain: a decision analysis approach, *European Journal of Operational Research* **110**: 597 – 609.

Daganzo, C.F. (1991). *Logistics systems analysis*, Springer-Verlag: Berlin.

Daganzo, C.F. (1987). Increasing model precision and reduce accuracy, *Transportation Science* **21**(2): 100 – 105.

Daskin, M.S., Hopp, W.J. and B. Medina (1992). Forecast horizons and dynamic facility location planning, *Annals of Operations Research* **40**: 125 – 151.

Hesse Owen, S. (1999). *Scenario planning approaches to facility location: models and solution methods*, Unpublished Dissertation. Evanston, IL: Northwestern University, Department of Industrial Engineering and Management Sciences.

Lapierre, S.D., Ratliff, H.D. and D. Goldsman (1998). Models for the delivery of preventive health services and application to Fulton County, Montreal: University of Montreal, Centre for Research on Transportation.

Laporte, G., Louveaux, F. and H. Mercure (1989). Models and exact solutions for a class of stochastic location-routing problems, *European Journal of Operational Research* **39**: 71 – 78.

Lewis-Beck, M. S. and C. Tien, Voters (1999). Voters as forecasters: a micromodel of election prediction, *International Journal of Forecasting* **15**, 175-184.

Narasimhan, S., Pirkul, H. and D.A. Schilling (1992). Capacitated emergency facility siting with multiple levels of backup, *Annals of Operations Research* **40**: 323 – 337.

Pace, R. K., R. Barry, O. W. Gilley, and C.F. Sirmans (2000). A method for spatial-temporalf with an application to real estate prices, *International Journal of Forecasting* **16**, 229-246.

Robusté, F., Daganzo, C.F. and R.R. Souleyrette II (1990). *Transportation Research B* **24B**(4): 263-286.

Wiegel, D. and B. Cao (1999). Applying GIS and OR Techniques to Solve Sears Technician-Dispatching and Home-Delivery Problems, *Interfaces* **29**(1): 112 – 130.

Exhibit 1.
Home Delivered Meals Kitchens and Sample
Of 25 Catchment Areas

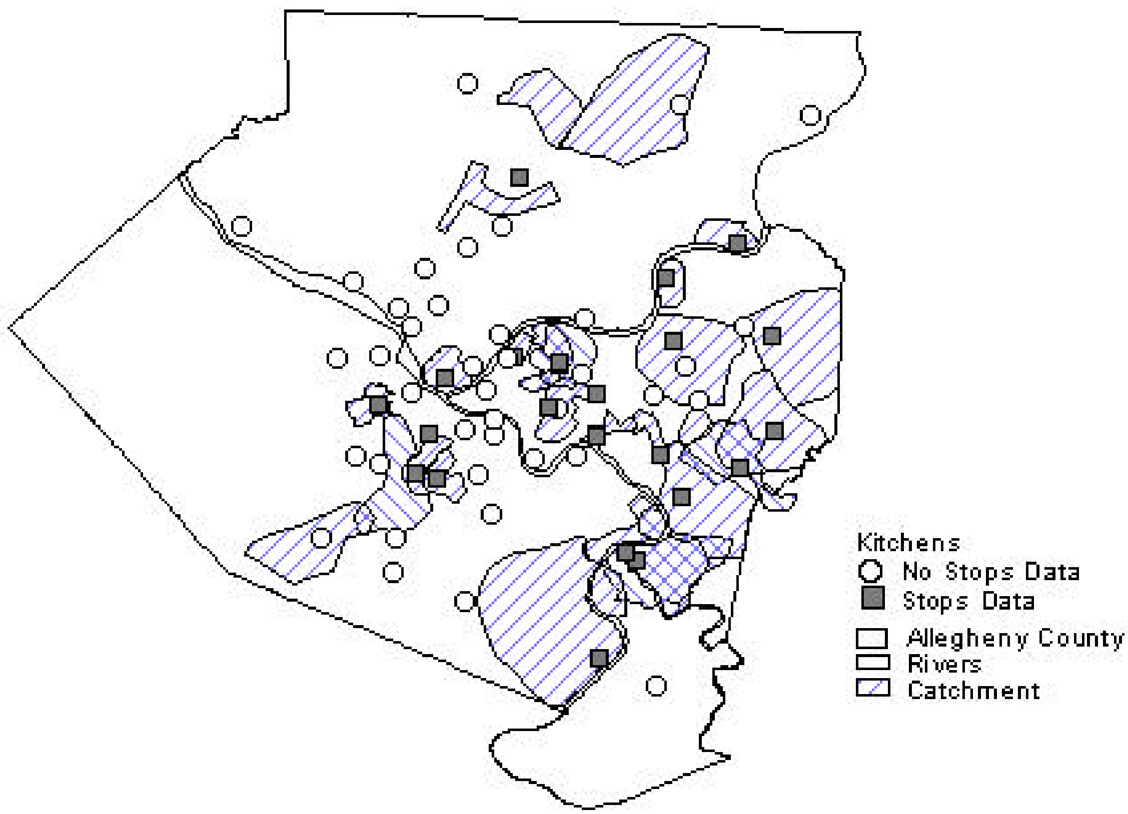

Exhibit 2.
Northside Kitchen, Client Stops, and Catchment Area with Elderly Population by Block Group

Exhibit 3.
Distribution of Utilization Rates.

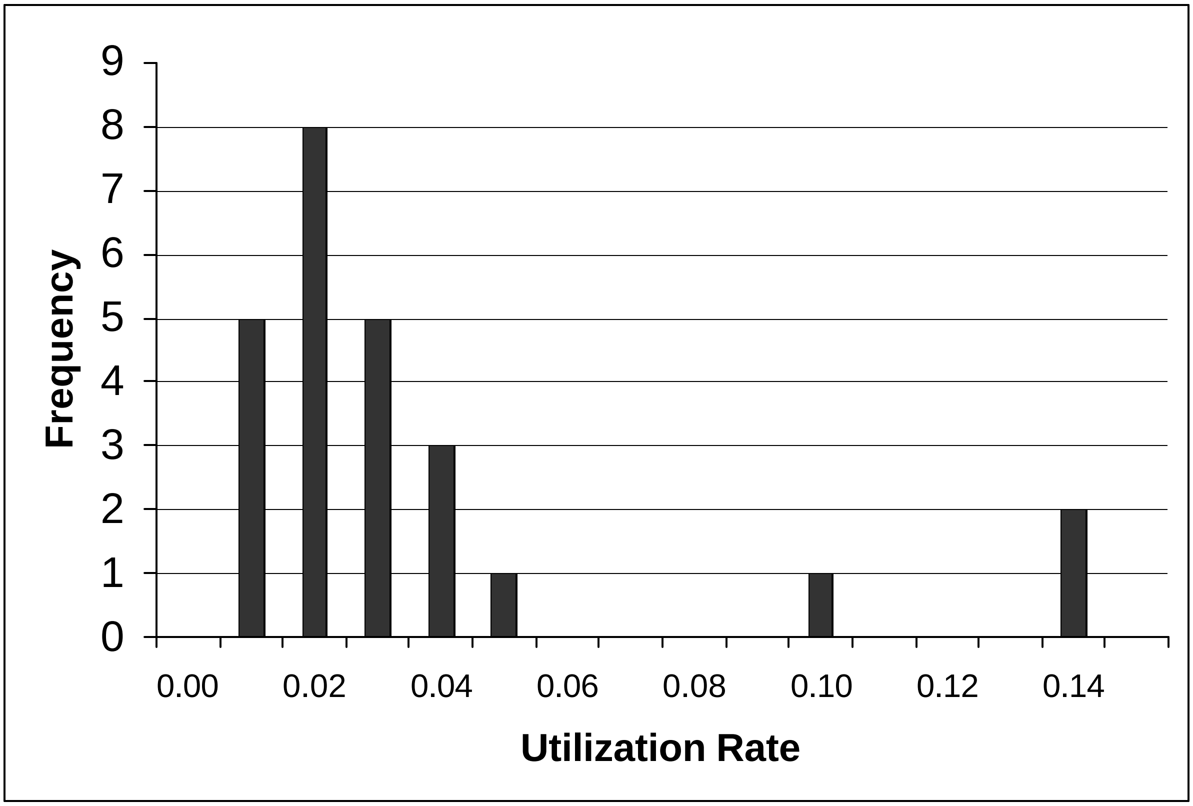

Exhibit 4.
Two Mile Buffers for Kitchens

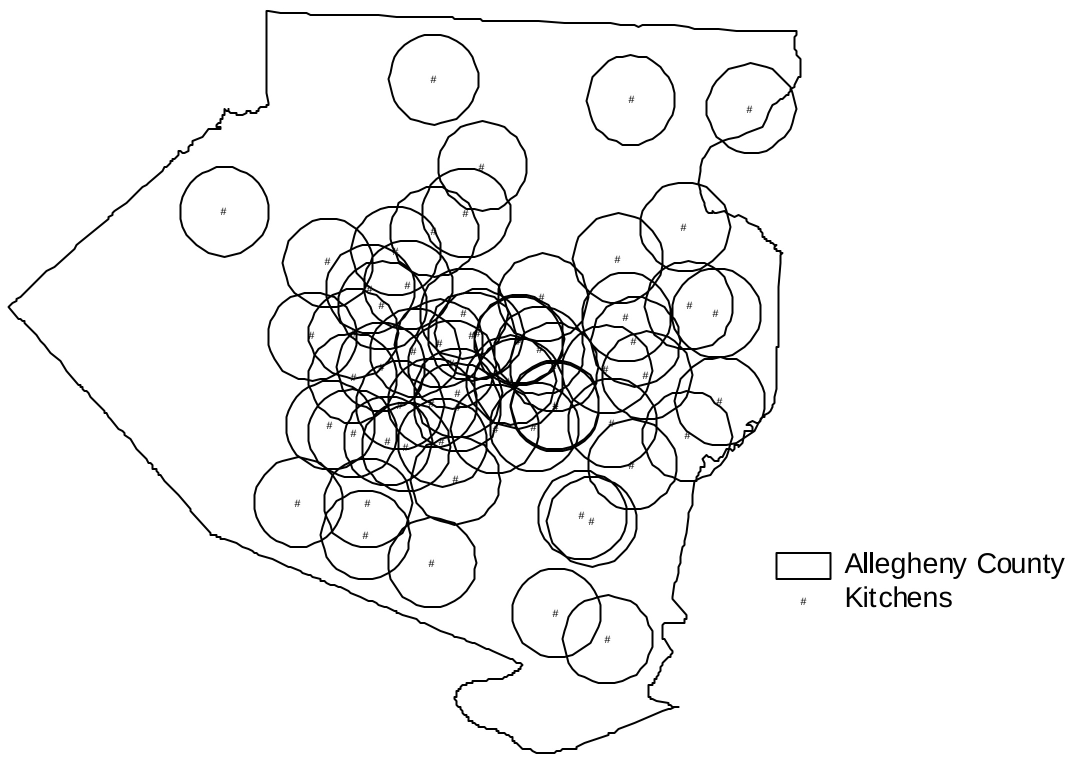

Exhibit 5.
Regression Model Results for Utilization Rate.

Regression Statistics	
Multiple R	0.511
R Square	0.261
Adjusted R Square	0.138
Standard Error	0.010
Observations	22

	Coefficients	*Standard Error*	*t Stat*	*P-value*
Intercept	0.02089	0.00904	2.31	0.03
NearKitchens	-0.00329	0.00222	-1.48	0.16
PCIncome	0.00003	0.00038	0.09	0.93
EldPopDensity	0.00000	0.00001	0.40	0.70

Exhibit 6.
Utilization Rate as a Function of Near Kitchens

Regression Statistics	
Multiple R	0.505
R Square	0.255
Adjusted R Square	0.217
Standard Error	0.010
Observations	22

	Coefficients	*Standard Error*	*t Stat*	*P-value*
Intercept	0.02277	0.00267	8.52	0.00
NearKitchens	-0.00250	0.00096	-2.61	0.02

Exhibit 7
Uncovered 2002 Forecasted Stops and
Estimated Maximum Catchment Areas for
Existing Kitchens.