Available online at www.sciencedirect.com
ScienceDirect
Journal of Transport Geography 16 (2008) 313–323
JOURNAL OF TRANSPORT GEOGRAPHY
www.elsevier.com/locate/jtrangeo

Efficient spatial interaction: attainable reductions in metropolitan average trip length

Morton E. O'Kelly *, Michael A. Niedzielski

Department of Geography, The Ohio State University, 1036 Derby Hall, 154 N Oval Mall, Columbus, OH 43210, USA

Abstract

We consider the intrinsic levels of spatial interaction necessitated by current North American urban spatial structure and investigate the potential for reducing average commuting trip lengths. In this paper, we use a doubly constrained spatial interaction model to generate critical parametric values which are then used to calculate the *effort* needed to reduce *average trip length* by a fixed target for each urban area. The measure of effort is defined as the degree of difficulty to achieve commuting and fuel consumption goals, and compared across a set of 26 US cities. Since the urban structure varies spatially, it is harder for some cities than others to reach a given level of reduced energy consumption.
© 2007 Elsevier Ltd. All rights reserved.

Keywords: Spatial interaction; Commuting; Entropy maximization; Energy; Fuel consumption

1. Introduction

> "Maybe I'll have to find work closer to where I live already. Change the job." (Daniels, 2006)

These words spoken by a Washington, DC commuter traveling 55 miles to work provide anecdotal evidence of an individual's efforts to adapt behavior when gas prices reach a critical point. The levels of energy prices at which commuting behavior and spatial patterns begin to show signs of adjustment have attracted wide attention in the media. While abundant cheap gasoline has fueled the decentralization of American cities for several decades, the present price surge has caused a good deal of attention to the built-in demand for gasoline at the national level. By built-in demand we mean a necessary demand for movement that is a function of the spatial arrangement of the city. It is much harder to reduce the demand for movement proportionally in those cities where places of work and residence are widely separated than in those cities where the activities are more uniformly intermingled. In this respect it is important to gauge from data the exact level of effort that would be needed to achieve a given improvement, and track the spatial variability in a comparative framework.

The issue at hand becomes one of determining the comparative difficulty of moving an observed trip distribution pattern in a city to a lower average trip length. The challenge is to model the likely trip distribution that is consistent with current residential and workplace distributions and at the same time attain a reduced average trip length. We wish to do this without pushing the trip distribution all the way to the linear programming based minimum average trip length (as advocated in the excess commuting framework). While the minimum is a socially optimal value associated with commuting, it is impossible to force individuals to produce a collectively socially optimal outcome. The entropy maximizing doubly constrained trip distribution model is an ideal vehicle for this analysis because it reflects the *most probable* distribution of trips given the data, which are then modified to reflect the lower target average interaction trip length. The question that is the focus of our research is this: if a city is on average required to travel shorter distances, how difficult is it for the city to reflect that modified pattern while maintaining the origins

* Corresponding author.
E-mail addresses: okelly.1@osu.edu (M.E. O'Kelly), niedzielski.4@osu.edu (M.A. Niedzielski).

0966-6923/$ - see front matter © 2007 Elsevier Ltd. All rights reserved.
doi:10.1016/j.jtrangeo.2007.11.003

and destinations of the trips? It is implicit in our analysis that the basic origin and destination locations supporting residential and commercial land-use patterns are themselves relatively fixed. Long term macro adjustments are beyond the scope of this article (see Horner and Murray, 2003).

This paper develops a partial approach for examining the spatial variability of the effort needed to reduce energy consumption across urban regions in the US. Generally speaking, within a given urban context, reducing trip length produces trip distributions that are less random and more structured.[1] Defining entropy as a measure of disorder or randomness, note that entropy *decreases* when trips become more highly structured. We employ a particular definition of "effort" in terms of the *change in entropy*, which captures the work performed by the spatial interaction model to reduce trip lengths while maintaining consistency with origin and destination totals. Exactly how much change in the aggregate statistic might be needed varies with the structural make-up of the city. The next section reviews and integrates three related strands in the literature on commuting, land-use and energy consumption. This is followed by the presentation of our modeling approach for addressing the degree of difficulty in meeting reduced energy consumption goals. Modeling results for the analysis of 26 metropolitan areas in the US are then given followed by a discussion and conclusions.

2. Background

Our focus is on reduction of intrinsic demand through trip distribution adjustments. By intrinsic demand we mean the locked-in demand for gasoline as a function of the current spatial arrangement of the city. This level is the result of a complex process involving worker locational decisions and auto ownership rates, among other factors. Imagine a family with two cars living in a single family house in the suburbs. Each parent drives to work; one of them to the office in a suburb on the other side of the central city, the other to a factory in an adjacent suburb. (We use an office and a factory as representative of fixed employment destinations – we do not mean to imply that journeys to work are only to professional and industrial work zones.) Given the family's preference for home and job locations, the clear pattern of journey-to-work will show the daily interchange between the origin and the destinations. Not much can change in terms of trip patterns in this situation, *given the fixed trip origins and destinations*.[2]

From the physical layout of the built-environment, we cannot avoid some elements of these costs, given the spatial separation of homes and jobs. In turn, some cities have a comparatively easy time of it, given that the spacing between their homes and factories may be particularly well laid out for the purposes of interaction. But as more and more of our city interaction patterns (Hanson and Giuliano, 2004) involve cross town commuting and movement between decentralized urban sub-centers (Rogerson, 1987) the task ahead is difficult if reduction of travel is to be achieved. It seems to be the case that the long average commutes are an unavoidable by-product of our urban spatial organization. In a sense then, we are indeed hooked on gasoline.

Now, imagine an experiment where we try to judge cities in terms of excessive use of transportation (i.e. aggregate vehicle miles). Cities that have *less* flexibility (because of their layout) should perhaps be recognized as having a lower potential for reducing spatial interaction costs. On the other hand cities that have a compact or "distance friendly" set up, might initially be among the more wasteful consumers of vehicle miles (distance), and so the reduction goals could be rather easily met through a reduction in the number of truly excessive commutes across town. At any rate, this spatially varying effort, degree of difficulty or flexibility is precisely the focus of our research.

3. Integrating commuting, spatial interaction and entropy maximization

By setting a goal of reducing average trip length by say 3% of current levels, it is reasonable to expect that such a goal may be overly ambitious for cities with an unfavorable spatial structure for commute reductions. Given the current snapshot of the layout of cities in the US, we are concerned with charting the spatially varying levels of the effort or degree of difficulty of adjusting travel patterns to high gas prices given the goal of a particular work travel distance reduction.[3]

One such approach to measuring and comparing theoretical and observed average trip lengths is the linear programming based approach to commuting efficiency (White, 1988). In this area of land-use, sustainability and transportation related work, the observed distributions of home and job locations along with an empirically derived distance matrix are used to derive the extreme ends of the trip distribution possible given the actual urban spatial structure (Horner, 2002; White, 1988). These extremes, the minimum and maximum average trip lengths, are then used to infer something about the commuting efficiency of the region, known in the literature as excess commuting (Horner, 2002). The more the home and job locations are

[1] For example, in the extreme case of a minimum cost solution to the related transportation problem, only $2n - 1$ of the possible n^2 cells are active. In other words $n^2 - 2n + 1$ cells are structurally zero.

[2] At very extreme levels of transport costs, the residential rents would at least technically have to be lowered to make the matching of locations feasible (Herbert and Stevens (1960) and as a logical deduction from the Alonso (1964) model of residential land rent).

[3] Three percent is a reasonable goal, because discussion of energy savings has recognized that such a relative saving would have a large aggregate impact on energy demand. "[I]f only people could cut back on 3% of travel" (Healey, 2006). Drivers curb use as gas goes up. *USA Today*.

decentralized and dispersed, the more possibilities exist for very short and very long commutes. Using the observed trip length sandwiched between the extremes, researchers then evaluate the ability of commuters to take advantage of commute minimizing opportunities within an inter-regional comparative framework. However, the current optimization based approach not only lacks the ability to generate theoretical trip distributions other than the extremes, it also cannot provide information on the difficulty involved in moving toward the two extremes. Thus, we must look elsewhere for a modeling approach that accounts for both of these deficiencies in the linear programming approach.

Based on the work of Wilson (1974) and the concepts developed in the excess commuting literature (Horner, 2002; White, 1988), we use the spatial interaction model to investigate the spatially varying possible energy consumption reductions. There are strong mathematical connections between the transportation-problem approach preferred in the excess commuting literature and the doubly constrained spatial interaction model in the investigation of urban form. The doubly constrained model provides beta parameter values corresponding to input values of the average commute distance. Clearly, the model generates a beta parameter associated with the observed commute. It is critical to note that in the present formulation, beta is ordinarily a positive number and enters the model with a negative sign.[4] Thus, as the magnitude of the beta parameter deterrence factor in the model *increases*, the average trip length *decreases*, and travel patterns become more efficient. We can continue to increase beta until the trip length approaches the asymptotic lower bound set by the associated minimum cost transportation problem. Similarly, in the opposite direction, smaller values of beta reduce the deterrent effect of distance allowing for inefficient, and ultimately "distance seeking" commuting patterns. As we continue to decrease beta,[5] even going to negative values, the trip length approaches the asymptotic upper bound trip length obtained by optimization. Thus, we can plot the curve of corresponding beta and average trip length values (see Fig. 1).

As opposed to the usual excess commuting based approach, where only two optimization benchmarks are compared, the doubly constrained spatial interaction model allows us to compare a range of feasible trip length solutions. Though minimization obviously provides the greatest reductions in average trip length, the prescription is impractical. However, noticeable changes in work travel distance can be accomplished by sensitivity analysis of the doubly constrained interaction model. The model accepts as input an attainable target average trip length in order to track the comparable difficulty associated with travel demand reductions given an urban spatial structure.

4. Methodology

We begin the model development by noting that entropy maximization provides trip distributions that are maximally consistent with the known data and at the same time noncommittal with respect to the missing data. The spatial interaction component provides the necessary calculations, while incorporating distance-decay effects (Fotheringham and O'Kelly, 1989; see also Sweeney, 1999). The exogenous input to the model is the known data consisting of the number of workers living in the origin zone (O_i), the number of jobs in the destination zone (D_j), the distance (C_{ij}) and observed flow (T_{ij}) between them. The maximum entropy model seeks to fill in the cells of the modified trip distribution maintaining consistency with the observed flows and origins and destinations of the trips while reflecting the reduced average trip length. The result of the model is the entropy statistic corresponding to the average trip length currently in the model. The analyst is able to generate entropy statistics for the entire range of beta between the maximum and minimum commutes found by optimization. Entropy is at a maximum when beta is zero, and represents a random trip distribution over the region. This random distribution is computed by taking the expected origin to destination flow as a function of their sizes, but independently of any distance effect. It is worth noting that this random trip distribution and associated average trip length is probably a more accurate upper bound on commuting patterns than the maximum average trip length from linear programming since commuters do not deliberately travel to work inefficiently; rather the worst journey-to-work pattern is a result of a random process spreading trips throughout the region.

There are intimate connections between the three statistics presented thus far, the average trip length, beta and entropy. For completeness Fig. 1 shows the full spectrum of beta values *before* they are inserted into the model; thus a "0.1" on the x-axis becomes an exponent of "−0.1" when beta enters the model. This conventional approach of beta values entering the model with a negative sign reflects the distance deterrent effect. For example, a beta value of 0.1 has a negative effect on interaction (i.e. tendency toward shorter trip lengths) because it appears as an exponent of −0.1 in the model, while a beta value of −0.1 has a positive effect on interaction (i.e. tendency toward longer trip lengths) because it appears as an exponent of 0.1 in the model. From Fig. 1, it is clear that there is a negative relationship between the average trip length and beta: as beta increases, the average trip length decreases reflecting the increasing deterrent on longer journeys-to-work.

[4] By convention in the literature beta is a positive number that enters with a negative sign, indicating that the distance deterrence term is effectively part of the denominator in the model. Thus higher distances serve to impede interaction. The larger the magnitude of beta, the greater this impedance effect.

[5] Note that decreasing beta through to negative numbers, given that beta enters the model with a negative sign, produces a counter-intuitive effect of travel attracted to longer distance.

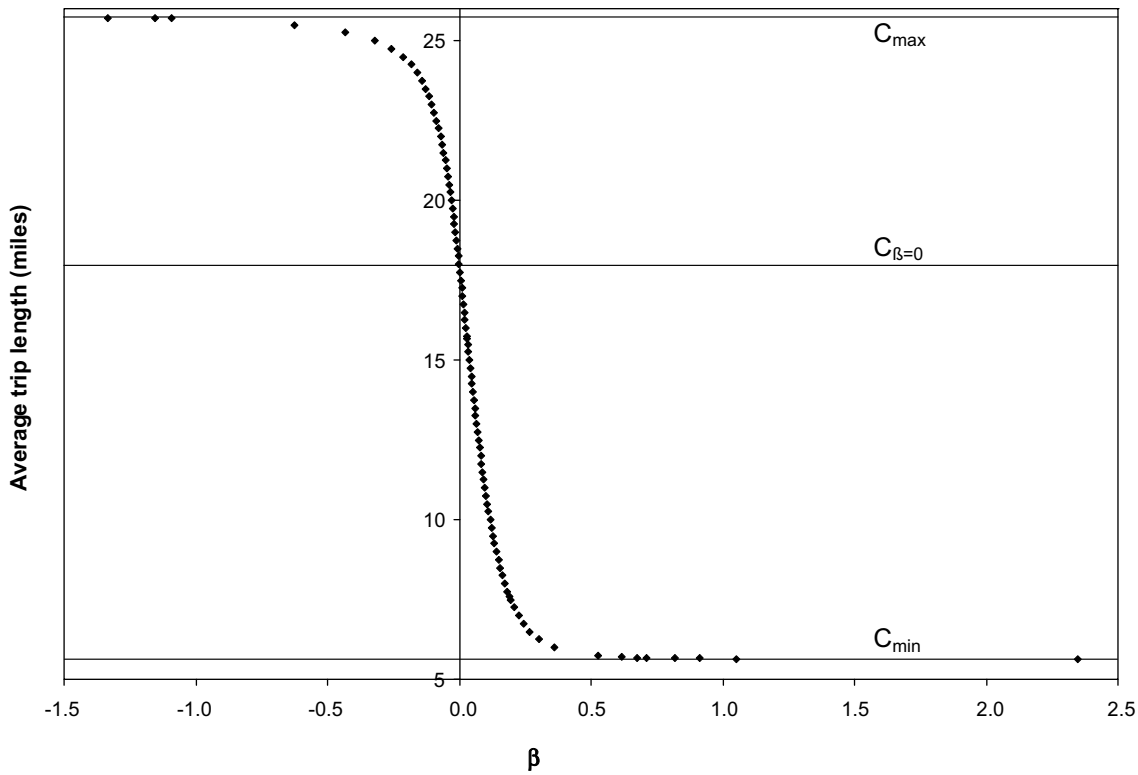

Fig. 1. Curve of β versus the average trip length.

Entropy values trace out a different path (Fig. 2). By logical deduction, as the average trip length is reduced from its maximum toward randomness, entropy increases, reaching a maximum at $\beta = 0$. Conversely, entropy (i.e. uncertainty) decreases as the trip length continues its slide toward the minimum. More importantly, as we move away from a random trip distribution, the system has to work harder to maintain the origin and destination outflow

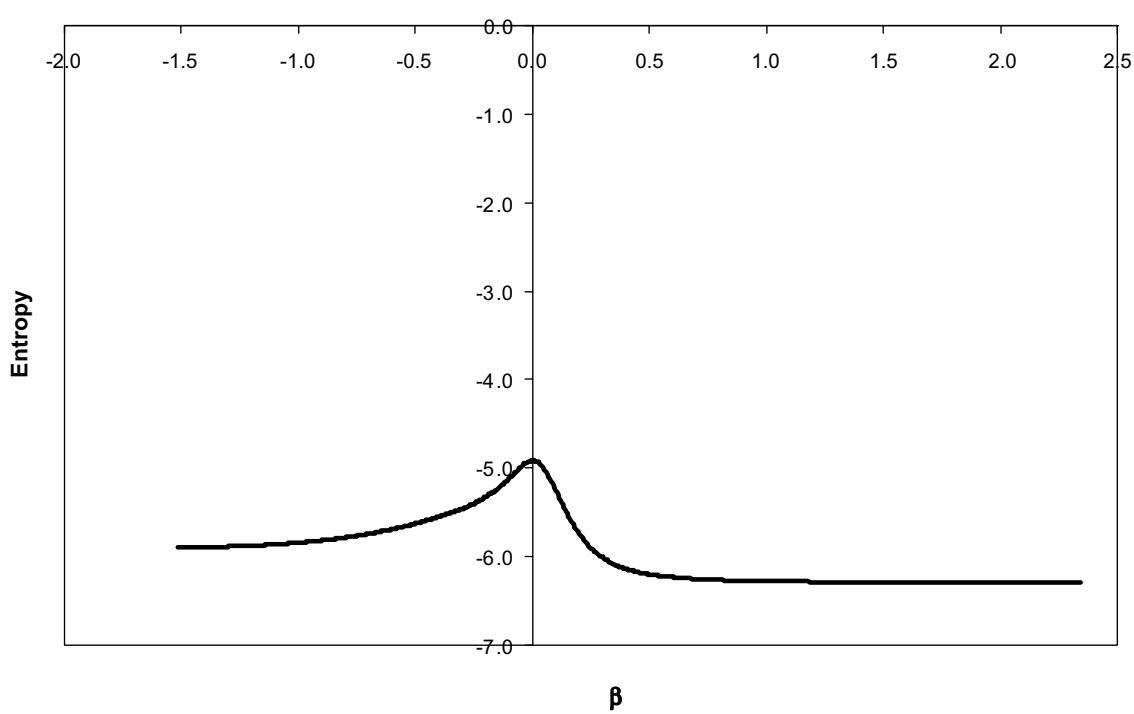

Fig. 2. Curve of β versus entropy.

and inflow totals, respectively. This effort in maintaining a balanced system is what we are calling the degree of difficulty or flexibility depending on the point of view.

Our approach for the estimation of entropy uses the maximum entropy derivation of the doubly constrained trip distribution model which is found from:

$$\begin{array}{ll} \text{MAX} & H = -\Sigma_i \Sigma_j T_{ij} \ln T_{ij} \\ \text{Subject to} & \Sigma_j T_{ij} = O_i \quad \lambda_i \text{ is the associated unknown Lagrangean multiplier} \\ & \Sigma_i T_{ij} = D_j \quad \mu_j \text{ is the associated unknown Lagrangean multiplier} \\ & \Sigma_i \Sigma_j T_{ij} C_{ij} = C \quad \beta \text{ is the associated unknown Lagrangean multiplier} \end{array}$$

It can be shown that

$$T_{ij} = \exp(-\lambda_i - \mu_j - \beta c_{ij})$$

which is the first order condition for the objective function to reach a maximum. Then following the common derivation of the spatial interaction model (Fotheringham and O'Kelly, 1989):

$$T_{ij} = A_i O_i B_j D_j \exp(-\beta c_{ij})$$

where A_i and B_j are the conventional balancing factors and are adjusted by the model to fit the modified trip distribution (T_{ij}) to the reduced average trip, C. Thus, the magnitude of the A_i's and B_j's is a function of the effort of the system to maintain balance with the outflows and inflows, respectively.

Substituting into H,

$$H = -\Sigma_i \Sigma_j T_{ij}(-\lambda_i - \mu_j - \beta c_{ij})$$

Therefore,

$$\begin{aligned} H &= \Sigma_i \Sigma_j T_{ij} \lambda_i + \Sigma_i \Sigma_j T_{ij} \mu_j + \Sigma_i \Sigma_j T_{ij} \beta c_{ij} \\ &= \Sigma_i \lambda_i \Sigma_j T_{ij} + \Sigma_j \mu_j \Sigma_i T_{ij} + \beta \Sigma_i \Sigma_j T_{ij} c_{ij} \\ &= \Sigma_i \lambda_i O_i + \Sigma_j \mu_j D_j + \beta C \end{aligned}$$

Relating λ_i and μ_j to the conventional balancing factors A_i and B_j, we can assume that A_i and B_j are defined as follows (by conventional derivation):

$$A_i = 1/\Sigma_j B_j D_j \exp(-\beta c_{ij})$$
$$B_j = 1/\Sigma_i A_i O_i \exp(-\beta c_{ij})$$

which means that

$$\lambda_i = -\ln(A_i O_i)$$
$$\mu_j = -\ln(B_j D_j)$$

We can now use these insights to solve the maximum entropy derived doubly constrained spatial interaction model for a given value of the average trip length. For the observed average trip length (C_1), this run results in:

$$H_1 = \Sigma_i \lambda^*_{i\,1} O_i + \Sigma_j \mu^*_{j\,1} D_j + \beta^*_1 C_1$$

where the asterisk is designed to emphasize that these are the solution values and all depend on the exogenous data: O_i, D_j, and C. Recall that in the excess commuting literature the O_i and D_j are typically held constant, and the trip length is varied (towards a minimum or maximum).

Notice that the entropy statistic, H, has three components: $\Sigma_i \lambda_i\ O_i$, $\Sigma_j \mu_j D_j$, and $\beta^* C$, which are generated from the known data and the results from the entropy maximizing model. Therefore, we can separately compute each component as follows:

$\Sigma_i \lambda^*_i O_i$ where $\lambda^*_i = -\ln(A_i O_i)$ from the balancing factors at the end of the run

$\Sigma_j \mu^*_j D_j$ where $\mu^*_j = -\ln(B_j D_j)$ from the balancing factors at the end of the run

$\beta^* C$ where β^* is the trip length parameter

The process is rerun for a different, *reduced*, average trip length (C_2)

$$H_2 = \Sigma_i \lambda^*_{i\,2} O_i + \Sigma_j \mu^*_{j\,2} D_j + \beta^*_2 C_2$$

Then after a small amount of algebra,

$$H_1 - H_2 = \Sigma_i O_i(\lambda^*_{i\,1} - \lambda^*_{i\,2}) + \Sigma_j D_j(\mu^*_{j\,1} - \mu^*_{j\,2}) + \beta^*_1 C_1 - \beta^*_2 C_2$$

The difference, $H_1 - H_2$, is the effort or degree of difficulty expressed in terms of the change in entropy, which is precisely the focus of our investigation.

Whereas H is a measure of the overall effort that the system needs to maintain balance while matching the reduced average trip length, each component of H measures the work needed to balance the origins, destinations and the trip length respectively. Thus, this approach generates aggregate and disaggregate measures of the system effort. Since these facts are encoded in the multipliers, we are able to interpret all the measures comparatively.

The relationship between the three statistics should now be clear. Assume C_0 is the average trip length associated with β equal to zero, and C_1 is the observed average trip length with the appropriate corresponding β. What happens when we lower the trip length to C_2? We move away from the entropy maximizing trip length, C_0, which occurs when β equals zero. Thus, we should see that $\beta_1 < \beta_2$ and $H_1 > H_2$ as we reduce trip length from C_1 to C_2.

5. Application results

For the purpose of this analysis, the entropy maximizing version of the spatial interaction model was coded in C++ using origin and destination totals obtained from the Census Transportation Planning Package (CTPP) part 3. The cost matrix is calculated using Euclidean distances between zonal centroids with intra-zonal trip lengths estimated by the radius of a circle circumscribing each zone (Frost et al., 1998). The model was run three times for each of the 26 cities in the study, given the input of the average trip

length associated with maximum entropy ($\beta = 0$), the observed average trip length and the average trip length corresponding to 3% reduction of the observed commuting pattern. The final output from the model contains the beta, aggregate and disaggregate entropy statistics for the given average trip length. In addition, the output file contains a measure of the effort in reducing the city wide trip length from the observed to the trip length at the 97% level of the observed.

Table 1 shows the measures across 26 cities, a sample of urban areas used in previous research (Horner, 2002), depicting the statistics before and after the 3% average trip length reduction. As a check of the theory, we can see that the changes in the statistics are as expected. It is clear that a decrease in the average trip length is associated with an increase in beta and a decrease in the entropy statistic. We can now as well see the relationship with the components of entropy, such that as the average trip length is

Table 1
Comparative regional spatial interaction and entropy statistics

MSA	Commuting level	Average trip length	β	Origin component	Destination component	Trip length component	Entropy
Atlanta	Observed	11.11	0.172	4.24	6.85171	1.91	13.00
	3% reduction	10.78	0.181	4.15	6.85018	1.95	12.94
Baltimore	Observed	8.57	0.199	4.07	6.63149	1.71	12.40
	3% reduction	8.31	0.208	3.99	6.63056	1.73	12.35
Boise	Observed	7.53	0.175	3.96	5.42067	1.32	10.70
	3% reduction	7.30	0.187	3.87	5.41923	1.37	10.66
Boston	Observed	6.57	0.242	5.31	6.60077	1.59	13.49
	3% reduction	6.37	0.254	5.23	6.60005	1.62	13.45
Charlotte	Observed	7.32	0.188	4.84	6.06462	1.38	12.28
	3% reduction	7.10	0.202	4.74	6.06385	1.44	12.24
Cincinnati	Observed	7.98	0.209	4.51	7.09556	1.67	13.28
	3% reduction	7.74	0.220	4.43	7.09505	1.70	13.23
Cleveland	Observed	8.20	0.203	4.25	7.19287	1.67	13.11
	3% reduction	7.95	0.212	4.18	7.19200	1.69	13.06
Columbus	Observed	7.94	0.192	4.71	5.96894	1.53	12.20
	3% reduction	7.70	0.203	4.62	5.96809	1.56	12.16
Denver	Observed	8.15	0.189	4.87	7.01659	1.54	13.43
	3% reduction	7.91	0.199	4.79	7.01591	1.58	13.38
Las Vegas	Observed	7.22	0.140	4.53	4.67013	1.01	10.21
	3% reduction	7.00	0.157	4.41	4.66779	1.10	10.18
Memphis	Observed	7.78	0.171	4.48	5.47578	1.33	11.28
	3% reduction	7.55	0.185	4.37	5.47514	1.40	11.24
Miami	Observed	7.48	0.173	5.16	5.97096	1.29	12.42
	3% reduction	7.26	0.187	5.06	5.96983	1.35	12.38
Milwaukee	Observed	7.68	0.195	4.80	6.36701	1.49	12.66
	3% reduction	7.45	0.203	4.73	6.36657	1.51	12.62
Minneapolis	Observed	9.66	0.196	4.02	6.60234	1.90	12.52
	3% reduction	9.37	0.207	3.92	6.60011	1.94	12.46
Omaha	Observed	5.77	0.238	3.92	5.33339	1.37	10.62
	3% reduction	5.60	0.254	3.83	5.33275	1.42	10.58
Philadelphia	Observed	8.46	0.208	4.15	7.65472	1.76	13.56
	3% reduction	8.21	0.216	4.08	7.65407	1.78	13.51
Phoenix	Observed	11.23	0.155	4.10	6.35675	1.74	12.20
	3% reduction	10.89	0.167	3.97	6.35594	1.82	12.14
Pittsburgh	Observed	8.28	0.211	4.59	6.58151	1.75	12.92
	3% reduction	8.03	0.220	4.52	6.58071	1.77	12.86
Portland	Observed	6.35	0.242	4.57	5.70416	1.54	11.81
	3% reduction	6.16	0.257	4.48	5.70331	1.58	11.76
Rochester	Observed	8.73	0.185	3.21	5.52884	1.62	10.35
	3% reduction	8.47	0.196	3.12	5.52763	1.66	10.30
Sacramento	Observed	8.53	0.166	4.35	6.33852	1.42	12.10
	3% reduction	8.27	0.175	4.27	6.33755	1.45	12.06
St. Louis	Observed	9.32	0.185	4.36	6.56971	1.73	12.65
	3% reduction	9.04	0.195	4.27	6.56849	1.76	12.60
San Antonio	Observed	7.99	0.173	4.68	6.05716	1.38	12.12
	3% reduction	7.75	0.185	4.59	6.05678	1.43	12.07
San Diego	Observed	9.49	0.138	4.17	5.53837	1.31	11.03
	3% reduction	9.21	0.146	4.10	5.53799	1.35	10.99
Seattle	Observed	9.24	0.185	3.98	6.35848	1.71	12.05
	3% reduction	8.96	0.194	3.90	6.35720	1.74	11.99
Wichita	Observed	7.37	0.181	4.02	5.24932	1.33	10.60
	3% reduction	7.15	0.192	3.93	5.24842	1.37	10.56

reduced, the origin and destination components are reduced but the third component increases. Notice that the increase in C is less than the combined decrease in O and D, thus, we see that taken as a whole, entropy decreases, which is what we expect since with tighter trip length constraints, we are moving further away from the maximum entropy. Moreover, the disaggregate statistics clearly indicate that the system has to work harder to balance the destinations than the origins. That is, it is easier to achieve system balance by adjusting the origins than the destination most likely because homes are much more widely distributed than jobs (which tend to be clustered).

In addition to summarizing the changes in statistics, Table 2 provides the measures of effort across the urban regions. The fifth column reports the raw value of the degree of difficulty, whereas the sixth column provides a standardized measure, the degree of difficulty as a percentage of the maximum entropy. The spatial variability in the degree of difficulty is apparent in Fig. 3. A general northeast–southwest trend emerges with larger values (more effort) in the northeast and lower values (less effort) in the southwest, though Atlanta and Phoenix are exceptions to this pattern. The interpretation of this pattern is that it is more difficult to reduce the average trip length in older, somewhat more compact, and somewhat more monocentric cities in the rust-belt than it is in the newer, more sprawl-like and multi-centered cities in the sunbelt. Atlanta and Phoenix are exceptions to the general pattern because of the particular spatial distribution of homes and jobs in these sprawling cities resulting in the two longest observed average trip lengths (11.10 and 11.23 miles, respectively), thus necessitating significantly more effort in commute reduction than their sun-belt neighbors.

The particular combination of the three entropy components is different for each city, but generally speaking based on Fig. 4, the trip length component (C) makes up a larger share of the entropy change in newer, larger and more sprawling cities than it does for older, smaller and less sprawling cities. This is not surprising given that sun-belt cities are larger in area than rust-belt cities, thus, it is easier to achieve balance (less work by the system) in sun-belt cities simply by reducing extreme length commutes. Because there are much fewer extreme commutes possible in rust-belt cities (fewer geometric possibilities), then the system focuses on the origin component to move the trip distribution in line with the target trip length. Thus, although it may seem counter-intuitive, it takes less effort to reduce commuting as measured by the average trip length in sprawling urban regions than it does in more compact cities, assuming the job types can be matched. Referring to the two exceptions identified in Fig. 3, Fig. 4 also shows that the change in entropy components for Phoenix falls in line with other sun-belt cities. Like other cities in this region, it is easier to reduce extreme trip lengths in Phoenix than the origin component. Although a sun-belt city, the focus on reducing commutes in Atlanta is reminiscent of that in northeastern cities.

While, in general it is true that the standardized effort metric is lower for the sun-belt and higher for rust-belt, it is also apparent from Table 2 that the effort metric does

Table 2
Comparative regional average trip length reduction effort statistics

MSA	ΔOrigin (%)	ΔDestination (%)	ΔTrip length (%)	ΔEntropy	Effort	Effort as % of maximum entropy
Atlanta	−2.24	−0.022	1.98	−0.451	0.05875	0.416
Baltimore	−1.82	−0.014	1.34	−0.421	0.05227	0.380
Boise	−2.28	−0.027	3.85	−0.382	0.04089	0.364
Boston	−1.51	−0.011	2.03	−0.362	0.04886	0.341
Charlotte	−2.07	−0.013	4.19	−0.349	0.04287	0.338
Cincinnati	−1.87	−0.007	2.00	−0.387	0.05144	0.354
Cleveland	−1.63	−0.012	1.16	−0.390	0.05107	0.348
Columbus	−1.75	−0.014	2.38	−0.386	0.04707	0.365
Denver	−1.64	−0.010	2.14	−0.354	0.04750	0.333
Las Vegas	−2.69	−0.050	9.12	−0.316	0.03224	0.310
Memphis	−2.44	−0.012	5.15	−0.368	0.04145	0.359
Miami	−1.96	−0.019	4.79	−0.324	0.04028	0.316
Milwaukee	−1.38	−0.007	1.40	−0.362	0.04587	0.336
Minneapolis	−2.52	−0.034	2.36	−0.467	0.05848	0.429
Omaha	−2.30	−0.012	3.53	−0.400	0.04249	0.385
Philadelphia	−1.64	−0.008	0.85	−0.397	0.05386	0.343
Phoenix	−3.36	−0.013	4.86	−0.445	0.05423	0.421
Pittsburgh	−1.59	−0.012	1.15	−0.414	0.05350	0.378
Portland	−1.92	−0.015	2.67	−0.403	0.04750	0.384
Rochester	−2.88	−0.022	2.70	−0.483	0.04995	0.442
Sacramento	−1.68	−0.015	2.15	−0.360	0.04359	0.337
St. Louis	−2.05	−0.019	2.17	−0.420	0.05311	0.395
San Antonio	−1.94	−0.006	3.51	−0.353	0.04281	0.341
San Diego	−1.81	−0.007	2.68	−0.368	0.04057	0.351
Seattle	−2.12	−0.020	1.95	−0.435	0.05248	0.395
Wichita	−2.09	−0.017	3.28	−0.389	0.04124	0.360

Fig. 3. The effort needed to reduce the average trip length as a percentage of maximum entropy.

Fig. 4. Regional variability in the change in entropy components. The change in the destination component is excluded due to very small values which occur because it is easier to adjust origins than destinations.

not tell the whole story. While Fig. 4 suggests that the C component is larger than the O component of entropy, this also means that the O component is larger for sun-belt cities than rust-belt cities, because otherwise the effort metric would be positive for southwestern regions. Since C is larger for the southwest, O has to be larger as well to compensate for C in the estimation of entropy.

Insights into the workings of the system are gained by taking a detailed look at the trip distributions before and after the reduction. What does the model do to reduce the average trip length in the easiest way possible? Which trip lengths are reshuffled to bring the distribution in line with the new target commute? Fig. 5 illustrates the trip distribution for Pittsburgh before and after the reduction, and the difference between them. For the observed average trip length of 8.28 miles, the trip distribution is the familiar curve tracing out the distance-decay effect. The trip distribution for the reduced average trip length of 8.03 miles also

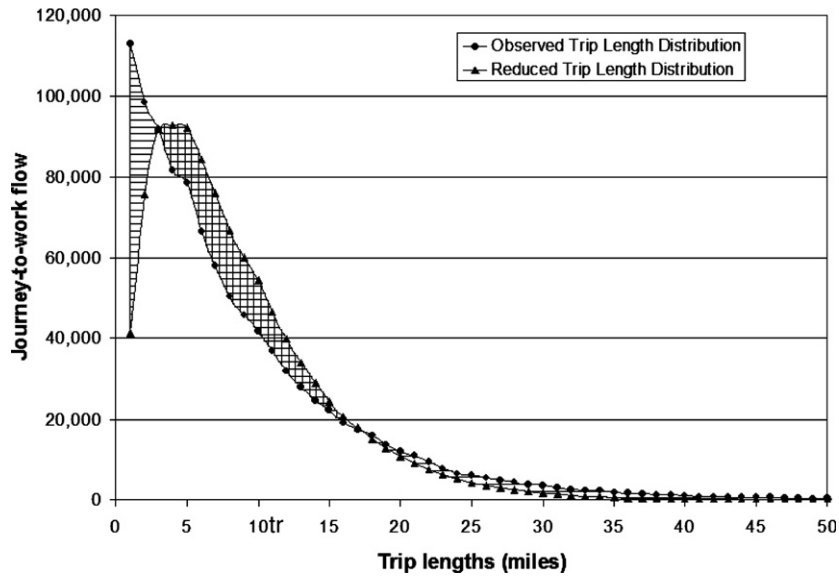

Fig. 5. Comparison of observed and reduced trip length distributions for Pittsburgh. Horizontal lines indicate fewer commuters and crosshatches indicate more commuters in the reduced work travel pattern.

encapsulates the deterrent effect but now extremely short trips are not the dominant type. Further, the curve of the reduced trip distribution is steeper and plateaus earlier than the more gradual decline of the observed trip distribution. Fig. 5 also suggests that in order to reduce the commute in line with the given target, the model reduces the number of extremely short trips and extremely long trips. In general, there is a preponderance of shorter trips, where the small reductions of longer trip lengths taken over a longer range compensates for the decrease in extremely short trips to result in lower overall commuting for the given region. The result of this is the tightening of the trip

Table 3
Comparison of longest average trip length for cumulative percentages of commuters

MSA	Cumulative % of commuters: 100		% change	Cumulative % of commuters: 95		% change
	Observed longest trip length (miles)	Reduced longest trip length (miles)		Observed longest trip length (miles)	Observed longest trip length (miles)	
Atlanta	121	82	−32.23	29.5	26	−11.86
Baltimore	90	82	−8.89	24.5	21	−14.29
Boise	110	87	−20.91	23	21	−8.70
Boston	56	57	1.79	19.5	15.5	−20.51
Charlotte	51	55	7.84	18	16	−11.11
Cincinnati	105	82	−21.90	22	19	−13.64
Cleveland	107	82	−23.36	23.5	20.5	−12.77
Columbus	60	70	16.67	22.5	19.5	−13.33
Denver	102	92	−9.80	23	19	−17.39
Las Vegas	123	115	−6.50	16	15.5	−3.13
Memphis	51	59	15.69	18	16.5	−8.33
Miami	53	59	11.32	18.5	16.5	−10.81
Milwaukee	74	71	−4.05	22	19	−13.64
Minneapolis	128	96	−25.00	28.5	23.5	−17.54
Omaha	53	57	7.55	14.5	13.5	−6.90
Philadelphia	129	93	−27.91	25.5	20.5	−19.61
Phoenix	121	110	−9.09	46	29.5	−35.87
Pittsburgh	87	77	−11.49	24.5	20	−18.37
Portland	71	62	−12.68	16.5	14.5	−12.12
Rochester	86	80	−6.98	27	22.5	−16.67
Sacramento	103	101	−1.94	24.5	21	−14.29
St. Louis	79	81	2.53	25	21.5	−14.00
San Antonio	57	67	17.54	21	18.5	−11.90
San Diego	73	73	0.00	27.5	21.5	−21.82
Seattle	132	100	−24.24	27	22.5	−16.67
Wichita	87	79	−9.20	23.5	20	−14.89

distribution around the reduced average trip length made possible by decreasing both tails within the imposed outflow and inflow constraints.

The trip length distributions provide another way to investigate the connections between spatial structure (land-use), commuting and energy consumption. As opposed to the *O–D* matrix which aggregates work trips by origins or destinations, the trip length distributions tally commutes by their distances. Thus, by comparing the distances at which 95% and 100% of the commuters[6] move through a given region for the observed and reduced work travel patterns, we can identify those places where reduction of extreme commutes is easier and by extension where it is easier to meet energy consumption reduction goals. Cities where a reduction in the longest trip length for all commuters is larger than the reduction in the longest trip length for 95% of travelers are more flexible in their ability to change their home/work dyad. Conversely, if the reduction of the longest trip length for 95% of workers is larger than the decrease for the trip length for all commuters then the distribution of home and work places in these city types is less friendly toward extreme commute reduction (akin to the "long valley" effect). Table 3 provides the trip length frequency for all 26 cities. As an example consider the MSA of St. Louis, Pittsburgh, Minneapolis and San Diego. In Pittsburgh for the observed case, it takes 87 miles for all commuters to move through the city, and only 24.5 miles for 95% of them. After the reduction, these distances decrease to 77 and 20 miles, respectively. In the case of all commuters, the longest distance drops by 11.5%, and the longest distance for 95% of the commuters decreases by 18.4%. For St. Louis, the intrinsic spatial structure is even more pronounced, because the longest distance actually increases by 2.5% from 79 to 81 miles, although the longest trip made by 95% of the commuters decreases by 14% from 25 to 21.5 miles. The long valley effect is evident when these cities are compared to Minneapolis and San Diego. The decrease in the longest distance for all commuters is larger than the decrease for trips made by 95% of the commuters. Thus, Minneapolis and San Diego have more freedom to attain the target because their spatial structure lacks the "long valley" effect.

6. Discussion and conclusions

If the goal of land-use related modeling is to examine the relationships between spatial structure, transportation, and energy consumption and evaluate changes in commuting behavior as a result of the realignment of the work/residence dyad, then previously developed "excess commuting" approaches may not address these concerns adequately. Analyzing the range of *most probable* realignments of the work/home dyad, instead of simply the extremes, is essential. Evaluations of attainable reductions in metropolitan average trip lengths are now possible. The proposed entropy maximization version of the spatial interaction model provides analysts with the ability to simulate various policy scenarios given land-use and energy consumption goals as inputs.

The current approach exploits the analytical features of the trip distribution model to determine the amount of effort needed to reduce average trip lengths across a sample of US cities. While our derivation and discussion has made a good deal of sense in terms of the logic of the model and arguably greatly improves over the linear programming solution, we are well aware of the needs for further refinement. The range of extensions to the current model should incorporate behavioral adjustments such as possible changes to car ownership (more fuel efficient and hybrid vehicles), lower overall discretionary travel levels, and modal shift. The popular view that a modest across the board reduction in vehicle miles (and consequently lower gas usage) could potentially reduce excess demand to the extent that gas would again come back to moderate price levels raises many more questions about sustainability that are beyond the immediate scope of this paper. Additionally, the model should disaggregate workers by job types. The model presented in this paper assumes that work trips can be rearranged without regard to workers job skills. Clearly, the affect on work trip lengths will be different when construction workers cannot change locations with university professors, for example. We are well aware that an aggregate analysis misses some of the finer details of commute realignment, but since our focus is on developing a framework for reductions in *average* trip lengths, the assumption of job homogeneity is reasonable. Furthermore, intra-metropolitan variations in the effort needed to reduce work trip lengths may be uncovered by spatially disaggregating the entropy statistics (refer to Niedzielski, 2006 for the approach).

While we do not yet have these factors embedded in a comprehensive way, we can at least lay out the steps that we need to perform to push this analysis to further levels of realism. A more refined and accurate model needs to address network link performance, equilibrium network assignment and model fuel consumption as a function of various trip length shares (Naess and Sandberg, 1996; Scott et al., 1997; Southworth, 2001) as well as disaggregate workers by their occupation types (see O'Kelly and Lee, 2005, for employment disaggregated excess commuting metrics). We are actively performing these computations in detail for particular cases. Still, the contribution of this paper is the quantification of the spatial variation in the spatial structure based effort needed to reduce work travel distance, which is the first step toward a deeper understanding of the complex processes involved in meeting lower energy goals.

[6] The 5% difference tries to capture the increase in the number of extreme commuters. For some urban areas, reduction of extreme commuting may be easier than for others.

References

Alonso, W.A., 1964. Location and Land Use. Harvard University Press, Cambridge, MA.

Daniels, L., 2006. Gas crisis hitting commuters hard, NBC Nightly News, April 29.

Fotheringham, A.S., O'Kelly, M.E., 1989. Spatial Interaction Models: Formulations and Applications. Kluwer Academic, New York.

Frost, M., Linneker, B., Spence, N., 1998. Excess or wasteful commuting in a selection of British cities. Transportation Research Part A 32 (7), 529–538.

Hanson, S., Giuliano, G. (Eds.), 2004. The Geography of Urban Transportation. The Guilford Press, New York.

Herbert, J., Stevens, B., 1960. A model for the distribution of residential activity in urban areas. Journal of Regional Science 2 (2), 21–36.

Healey, J.R., 2006. Drivers curb use as gas goes up. USA Today (April 20), A1.

Horner, M.W., 2002. Extensions to the concept of excess commuting. Environment and Planning A 34 (3), 543–566.

Horner, M.W., Murray, A.T., 2003. A multi-objective approach to improving regional jobs–housing balance. Regional Studies 37 (2), 135–146.

Naess, P., Sandberg, S.L., 1996. Workplace location, modal split and energy use for commuting trips. Urban Studies 33 (3), 557–580.

Niedzielski, M.A., 2006. A spatially disaggregated approach to commuting efficiency. Urban Studies 43 (13), 2485–2502.

O'Kelly, M.E., Lee, W., 2005. Disaggregate journey-to-work data: implications for excess commuting and jobs–housing balance. Environment and Planning A 37 (12), 2233–2252.

Rogerson, P.A., 1987. Changes in U.S. national mobility levels. Professional Geographer 39 (3), 344–351.

Scott, D.M., Kanaroglou, P.S., Anderson, W.P., 1997. Impacts of commuting efficiency on congestion and emissions: case of the Hamilton CMA, Canada. Transportation Research D 2 (4), 245–257.

Southworth, F., 2001. On the potential impacts of land use change policies on automobile vehicle miles of travel. Energy Policy 29 (14), 1271–1283.

Sweeney, S., 1999. Model-based incomplete data analysis with an application to occupational mobility and migration accounts. Mathematical Population Studies 7 (3), 279–305.

White, M., 1988. Urban commuting journeys are not 'wasteful'. The Journal of Political Economy 96 (5), 1097–1110.

Wilson, A.G., 1974. Urban and Regional Models in Geography and Planning. Wiley, New York.