Valuing the Consumption Benefits of Urban Density

Victor Couture *

University of California, Berkeley

This draft: 31 July 2014

Abstract: Despite growing interest in urban consumption amenities, little is known about their origin and importance. This paper estimates the consumption value of urban density by combining travel microdata with Google's local business data. This dataset allows to integrate travel costs into a discrete choice model for restaurants. I find that in high density areas, consumers enjoy large benefits from visiting places that they prefer, and relatively smaller gains from shorter trip time. These results demonstrate the importance of non-tradable consumption in explaining the value of cities, and represent the first estimates of the gains from variety in the service sector.

Key words: consumer cities, gains from variety, urban density, accessibility, travel demand.

JEL classification: D12, R41

*I am grateful to my advisors Gilles Duranton, Robert McMillan, Matthew Turner and Will Strange for their advice and encouragement throughout this project. The paper also benefited from conversations with Victor Aguirregabiria, Nate Baum-Snow, Ettore Damiano, Vernon Henderson, Maciej Kotowski, Nicholas Li, Martin Osborne, Diego Puga, Olmo Silva, Aloysius Siow and Nathan Yang. I also thank seminar participants at Baruch College-Zicklin, Brown University, HEC Montreal, London School of Economics, New York University-Furman Center, Sciences Po, Universite Laval, U Penn-Wharton, UC Berkeley-Haas, and the University of Toronto, and conference participants at the NARSC 2013 for additional comments. Financial support from the Social Science and Humanities Research Council and from the Royal Bank Graduate Fellowship in Public and Economic Policy is gratefully acknowledged. All remaining errors are my own.

Contact: Haas School of Business, University of California, Berkeley, 2220 Piedmont Avenue, Berkeley, CA 94720, USA (e-mail: couture@haas.berkeley.edu, website: http://faculty.haas.berkeley.edu/couture/)

1. Introduction

In 2010, 80.7% of Americans lived in urban areas, covering only 3.0% of the US land mass. To explain this sharp agglomeration pattern, empirical research has focused on measuring how larger, denser cities make workers more productive.[1] The benefits from density, however, are not limited to workers. Consumers may also gain from density, through better access to a variety of goods and services (Glaeser, Kolko, and Saiz, 2001). Such consumer-based agglomeration forces are now part of the mainstream discussion among urban analysts and policy-makers, but being inherently hard to measure they receive less academic attention than agglomeration forces increasing workers' productivity. So while the potential for consumption amenities to drive urban success is an area of vigorous policy debate,[2] empirical evidence remains scarce as to the origin and importance of the consumption advantage of cities.

This paper sets out an approach to estimating the consumption value of urban density. The estimation uses travel data and exploits the recent availability of detailed online micro-geographic data on local businesses. While it is simple to compute the value of shorter trip times by using estimates of value of travel time, my methodology also allows me to estimate the gains from increased choice in denser areas (so-called 'gains from variety'). I identify an individual's willingness to pay for access to a preferred location from the extra travel costs that she incurs to reach it.

I provide estimates of the gains from density in the US restaurant industry, a prominent part of the urban service sector. This exercise is of interest for three reasons. First, American households spend almost a quarter[3] of their income on non-tradable services such as restaurants, live entertainment, and many professional services (e.g. medical care) requiring face-to-face

[1] See for instance Melo, Graham, and Noland (2009) for a meta-analysis of estimates of agglomeration economies, and Combes, Duranton, and Gobillon (2011) for a survey of key empirical issues.

[2] For instance, Clark (2003) shows that cities providing more natural and constructed amenities experience faster population growth. Carlino and Saiz. (2008) find that metropolitan areas that are attractive to tourists (likely because of consumer amenities) are also growing faster. Moretti (2012) and Diamond (2013), however, argue that local production shocks drive urban success, with higher consumption amenities being a largely endogenous outcome.

[3] I exclude legal services and some other professional services from this computation. I include health care services, which account for 16.7% of total personal consumption expenditures, meals plus drinks with purchased meals, which account for 5.2%, and personal care services, which account for 0.9%. Sports, museums, live entertainment and laundry account together for another 1.4%. These numbers are from the Personal Consumption Expenditure (PCE) 2013 Tables of the Bureau of Economic Analysis.

interactions. The paper is the first to measure the importance of variety in non-tradables, and the results suggest that it represents most of the consumption value of modern cities. This finding stands in contrast to existing urban theories, which only model variety in the tradable sector.[4]

Second, the particular channels through which consumers benefit from density have important implications for urban policy. Intuitively, spatial proximity to restaurants in dense urban areas allows individuals to take very short trips to eat out, and a vast literature studies the impact of higher density on reducing trip length. However, I show that urban density also reduces the travel costs of substituting among destinations to visit places that one prefers. My estimates suggest that such gains from variety are in fact more important than gains from shorter trips. This trade-off between gains from variety and gains from shorter trips allows to anticipate the impact of increasingly popular policy attempts at reducing vehicle travel by encouraging higher density living.

Third, estimates of the gains from variety in the service sector contribute to an emerging literature estimating the gains from variety in consumer goods (Broda and Weinstein, 2006, 2010).

To estimate the welfare gains from urban density, I specify a discrete choice model of demand for travel destinations. I focus the analysis on restaurants because of data availability, but they alone account for more than 5% of household expenditures. In the model, each restaurant receives a logit utility shock, and locations farther away from an individual are more expensive because of travel costs. Individuals face a trade-off between the gains from visiting a preferred restaurant and the costs of a longer trip. The key parameter of the model is an elasticity of substitution between restaurants, which I estimate by maximum likelihood. This estimation exercise does not require data on restaurant choice, only on trip length. Using transport costs to provide variation in restaurant prices solves an endogeneity problem that typically arise when estimating an elasticity of substitution, due to the unobserved relationship between higher prices and better quality. If individuals always travel to

[4]Prominent examples of such models are in the "New Economic Geography" literature (e.g. Helpman, 1998), that builds on Krugman's (1991) seminal contribution.

the closest restaurant, then restaurants must be perfect substitutes and there are no gains from restaurant variety in dense areas, only savings through shorter trips. If individuals take long trips to eat out, then restaurants are imperfect substitutes and gains from variety are correspondingly large. To obtain welfare estimates, I derive a variety-adjusted restaurant price index from the logit model.[5] In the simplest specification of the model, all the price variation among restaurants come from differences in transport costs, and the restaurant price index in a location is low if there are many restaurants nearby, or if travel speed is high.

Estimating a discrete-choice model for travel destinations requires not only data on travel behavior, but also comprehensive microgeographic data on the location of all destinations available to an individual. The travel data come from the National Household Travel Survey (NHTS), which identifies trips to a restaurant and the location of an individual's home at the block group level. Importantly, the data allows to restrict to sample to trips with known origin whose only purpose is visiting a restaurant. The NHTS also allows to estimate car travel speed in different areas. I collect data on each restaurant online, from Google's local business pages. I observe the exact location of almost all restaurants (273,000 units) in 14 states representing 50% of the US population. For robustness tests and extensions of the model, I also collect data on restaurant characteristics. Each restaurant's name and type of cuisine (e.g. 'pizza') comes from its Google page. For a subset of the sample, there is additional information, such as meal price and quality ratings, from the popular review website Yelp.com.

I find substantial variation across areas in the variety-adjusted restaurant price index, generating large spatial welfare differentials. The gains from density are very localized, and much of the variation in the price index occurs within metropolitan areas (MSAs).[6] For a car driver - I also compute indices for pedestrians - in a large MSA, the restaurant price index generally drops by more than 20% from an MSA's outskirts to its downtown, which represents yearly gains of about $600 for an average household. Less than half of these gains from density comes from shorter trip times, with the remainder accruing through gains from variety. In the

[5]This index turns out to be identical to the 'love-of-variety' constant elasticity of substitution (CES) price aggregator. Anderson, de Palma, and Thisse (1992) prove that the under a linear utility specification, which I use, the logit and CES model lead to the same choice probabilities.

[6]These findings are consistent with Albouy and Lue's (2011) quality-of-life estimates, which are higher in denser areas and vary almost as much within metropolitan areas as across them.

countryside, individuals generally travel to one of the few restaurants that are closest to home, while in the densest areas travelers often pass by hundreds of restaurants on their way to a favorite destination. These results have important policy implications, because they allow to anticipate the impact of policies promoting denser, walkable, mixed-used neighborhoods, that are arguably the most popular set of urban policies in recent years.[7] For instance, my results suggests that such policies do reduce travel times, but have a larger effects on increasing gains from variety.[8]

A comparison of these results with recent estimates of the gains from variety in tradable goods hint at the primacy of non-tradable variety in explaining the consumption advantage of dense urban areas.[9] Handbury and Weinstein (2012) use precise scanner data to show that residents of larger MSAs face a lower price index for groceries (items with a barcode), controlling for store amenities, individual characteristics and differences in the number of varieties available. I do not compute MSA level indices like Handbury and Weinstein (2012), because my goal is to measure the gains from urban density, which turn out to be very localized. In fact, I show that most restaurants in a large MSA are essentially irrelevant to the welfare of any given resident, because of high travel costs. Keeping in mind these obstacles to a direct comparison, my estimates suggest much larger geographic welfare differentials for the non-tradable sector than for the tradable (barcode items) sector. This result is consistent with the dramatic decline in the cost of shipping goods over the last century, without a reduction of corresponding magnitude in the cost of moving people (Glaeser and Kohlhas, 2004). Urban density facilitates the movement of people, on which much of the non-tradable sector depends.

These welfare estimates are subject to a number of econometric and specification issues.

[7] Since the mid-1990s, the United States department of Housing and Urban Development has invested billions of dollar into such policies, that are usually inspired by New Urbanism, an influential planning movement (see Congress of the New Urbanism (2013)). For instance, New Urbanism theories are a major influence behind HOPE VI, which is HUD's program to revitalize distressed areas (Popkin, Katz, Cunningham, Brown, Gustafson, and Turner, 2004). From 1993 to 2010, HOPE VI made 263 grants worth $6.2 billions (`http://portal.hud.gov/hudportal/HUD?src=/program_offices/public_indian_housing/programs/ph/hope6`, webpage last visited in October 2013).

[8] Note that increasing density has a larger effect on reducing travel distance than travel time, because it lowers travel speed. Individuals, however, take more restaurant trips in dense areas.

[9] Murphy (2013) suggests that access to a high density of non-tradables enables individuals to save on land and durable goods (e.g. a car and a washing machine become unnecessary). According to this theory, the gains from non-tradable density are even larger than what I document here.

First, restaurant characteristics may vary across areas with different density, for instance low density areas feature more pizza, burger and family restaurants. To address these issues, I estimate a nested-logit model in which restaurants serving the same type of cuisine are more substitutable. I also estimate specifications in which restaurants in the same chain are perfect substitutes. Average meal price and quality may also vary systematically with density, but partial price and quality ratings data from Yelp only show a small correlation for meal price and none for ratings. Comparing the price of a McDonald's Big Mac - a price index popularized by the Economist - across areas does not suggest enough spatial variation to qualitatively affect my results.

Second, individuals may sort into areas based on unobservable characteristics affecting their gains from density. For instance, individuals with higher value of travel time or higher taste for variety may prefer to live in denser areas. To assess the strength of this sorting, I compare the effect of restaurant density on trip time in OLS regression to its effect in IV regressions. The instrument for restaurant density is past growth in population density, which has a large effect on current restaurant density after controlling for population in the initial period. The identification strategy is to restrict the sample to individuals with a very low probability of moving in any given year (55 years and older, married homeowner). For these individuals, recent population growth is an almost exogenous event, because a vast majority of them have lived in the same area for many years. OLS regressions predict shorter trip time in denser area than IV regressions, suggesting that individuals with a higher value of travel time, who make shorter trips, sort into dense areas. Based on these results, I estimate a version of the logit model in which high value of time individual sort into high density areas.

Third, individuals may sort into areas featuring their favorite type of restaurants. For instance, Waldfogel (2008) shows that availability of different restaurant types depends on the characteristics of the local population. To address this issue, I specify a restaurant supply model that allows to derive tastes for each type of cuisine in each area. I then use these taste parameters to estimate the nested-logit model.

Finally, the model may be misspecified. In particular, the independence of irrelevant alternative (IIA) property imposes a strong restriction on the logit model. I show how to test the logit model's predictions, and the IIA, using regressions of trip time on measures of restaurant density. The logit model performs remarkably well, but regression analysis nevertheless identifies a discrepancy between the data and the model's prediction. Extensions of the logit model provide predictions closer to the data, and generate similar but slightly larger welfare gains from density.

2. Literature review

The paper relates to an influential literature initiated by Rosen (1979) and Roback (1982), which uses data on wages and house prices to value city amenities. Albouy (2008) is a recent study using this approach. These studies estimate the value of urban amenities indirectly, as part of a larger residual explaining lower wages or higher house prices in a spatial equilibrium model. Therefore, this approach cannot determine willingness to pay for a particular amenity. I solve this measurement problem by using a different methodology, exploiting data on individual travel decisions to obtain direct estimates, at a precise location, of the value of a consumption amenity.

The use of travel cost differentials to value an amenity has a distinguished history in environmental economics, starting with Hotelling's (1947) letter to the National Park Service. Clawson (1959) provides the first of many applications of a travel costs method to the valuation of recreational facilities and environmental resources, which lack variation in prices from which demand curves are usually estimated. Ben-Akiva and Lerman (1985) develop the idea that travel demand can be modeled as a discrete choice problem, and Bockstael, McConnell, and Strand Jr. (1989) study of fishing lakes in Florida is an early example of a discrete-choice model with travel costs as an attribute of an amenity.

Some recent papers combine aggregate travel data with local business data to study spatial competition (Davis, 2006, Houde, 2012) and network effects (Wang, 2010). This paper, however, uses considerably more precise micro-data on travel, and focuses on a different problem: the consumption value of urban density.

Finally, the paper also relates to two major strands of literature in urban planning and transportation. First, the variety-adjusted price index for destinations, that I estimate for restaurants, corresponds to what transportation researchers call a 'travel accessibility index.' Bhat, Handy, Kockelman, Mahmassani, Chen, and Weston (2000) provide a literature review. Unlike available travel accessibility indices, my index has a natural interpretation as a price, and it depends on standard structural preference parameters. Second, reduced-form regressions of trip time on measures of restaurant density, that I run to test the model, belong to a large empirical literature measuring the relationship between travel and the built environment. The motivation behind these studies is generally to test whether higher population density living reduces vehicle travel, as argued by planning theories that influence much of recent urban policy. Ewing and Cervero (2010) provide a meta-analysis. Consistent with my regression results, other studies find a relatively small effect of density on vehicle travel. My paper suggest that this empirical regularity arises from a trade off between gains from variety and gains from shorter trips.

3. A logit model of travel demand

My analysis starts from two assumptions about the demand for travel. The first is that destinations are substitutable, which implies that an individual prefers some destinations to others. The second is that travel is costly, so that the price of visiting a destination farther away is higher. These assumptions imply a trade-off between the gains from going to a preferred destination and the costs of a longer drive.

Now, consider the problem of an individual living at location k and choosing a restaurant. Later, I solve the problem of how much to spend on restaurants versus all other goods. Let i index the number I_k of restaurants available, so that $i \in \{1,2,3,\ldots,I_k\}$. The restaurant with index $i = 1$ is closest from location k, $i = 2$ is second closest and so on. Denote travel time to restaurant i by t_{ki} and fuel cost by f_{ki}. The price of a meal at any restaurant is a constant p. Setting a constant meal price is equivalent to assuming that utility is invariant to variation in meal prices, because higher prices are always exactly compensated by higher quality.[10] Section 8.2 discusses the issue of price and quality variation across areas in more

[10]Such an assumption is reasonable if quality is mostly produced through variable costs, as is likely the case in the restaurant industry (see Berry and Waldfogel (2010) for a discussion of product quality in the restaurant industry).

detail. The total price of eating at restaurant i, including transport costs to and from the restaurant, is $p_{ki} = p + 2(\gamma t_{ki} + f_{ki})$, where γ is the value of travel time. This total price is what should be understood when I refer to restaurant price elsewhere in the paper, unless I mention 'meal' price specifically. Each restaurant receives a random idiosyncratic shock ϵ_{ki}, which captures an individual's preference for restaurant i. ϵ_{ki} is a random draw from a type I extreme value distribution with scale parameter $1/(\sigma - 1)$, where σ will turn out to be the elasticity of substitution between restaurants.[11] Note that parameters of the model do not vary with individual characteristics like income, which I introduce later in robustness check and extensions of the model.

Define the utility from making r_{ki} trips to restaurant i as:

$$u_{ki} = ln(r_{ki}) + \epsilon_{ki}.$$

Let y_k be expenditures on restaurants in location k, so that $p_{ki}r_{ki} = y_k$ is the individual's budget constraint. Note that although quantity r_{ki} can vary, in a discrete-choice model individuals can only choose one restaurant. This specification has the advantage of generating exactly the same aggregate consumption shares and welfare gains as the standard CES love-of-variety model. Substituting r_{ki} from the budget constraint into the utility function leads to the following indirect utility from choosing restaurant i:

$$v_{ki} = ln(y_k) - ln(p_{ki}) + \epsilon_{ki}.$$

The individual's problem is to choose the restaurant i that maximizes her utility:

$$max\{-ln(p_{k1}) + \epsilon_{k1}, \ldots, -ln(p_{ki}) + \epsilon_{ki}, \ldots ln(p_{kI_k}) + \epsilon_{kI_k}\}. \quad (1)$$

Note that expenditures y_k do not affect restaurant choice. The logit choice probability is equal to $\frac{e^{ln(p_{ki})/(\sigma-1)}}{\sum_{i=1}^{I_k} e^{ln(p_{ki})/(\sigma-1)}} = \frac{p_{ki}^{1-\sigma_j}}{\sum_{i=1}^{I_k} p_{ki}^{1-\sigma}}$, for all restaurants i (see Train (2009) for details and a proof). The number of trips to restaurant i is equal to $r_{ki} = y_k/p_{ki}$, so the probability of a trip of length t_{ki} to restaurant i, given the set of travel times to all restaurants $T_k = \{t_{k1}, \ldots, t_{ki}, \ldots, t_{kI_k}\}$ and

[11] The elasticity of substitution, in this setting, is a parameter measuring, for any two restaurants, the ratio of percentage change in relative demand to percentage change in relative prices. For instance, a low elasticity of substitution means that demand is not responsive to price variation.

the set of fuel costs $F_k = \{f_{k1}, \ldots, f_{ki}, \ldots, f_{kI_k}\}$, is:

$$prob_{ki} = prob(t_{ki}|T_k, F_k) = \frac{p_{ki}^{-\sigma_j}}{\sum_{i=1}^{I_k} p_{ki}^{-\sigma}}. \tag{2}$$

An important property of the probability in equation (2) is that its computation uses only data on travel time and fuel cost for restaurant trips, without requiring information on exact restaurant choice. To better understand the workings for the model, consider the probability ratio of trips to restaurants 1 and 2 (the closest and second closest restaurants):

$$\frac{prob_{k1}}{prob_{k2}} = \left(\frac{p_{k1}}{p_{k2}}\right)^{-\sigma_j} = \left(\frac{p + 2(\gamma t_{k1} + f_{k1})}{p + 2(\gamma t_{k2} + f_{k2})}\right)^{-\sigma}. \tag{3}$$

Equation (3) highlights two key features of the model. First, if σ is high, then restaurants are very substitutable and the ratio in equation (3) is large. When σ is high, individuals are sensitive to price differences, so they travel mostly to the closest restaurant, which is cheaper because of lower travel costs. This will be the main intuition behind the strategy for estimating σ using travel data. Second, if the difference between t_{k1} and t_{k2} is large, in a low-density area in which restaurants are far apart, then the proportion of trips to the closest restaurant is also large. Individuals living in low-density areas mostly visit the closest restaurant instead of traveling to places that they prefer, because substituting between restaurants is expensive. Therefore, higher density allows individuals to cheaply substitute between destinations, and to visit places that they prefer.

Equation (3) also shows that σ represents the elasticity of substitution between restaurants. This elasticity has two interpretations. In the first interpretation, individuals have constant tastes and always travel to the same restaurant, as in the model. In this case, a low σ represents heterogeneous preferences for restaurants across many otherwise identical individuals. In the second interpretation, individuals get new idiosyncratic shocks from the same distribution before each restaurant choice. In this case, a low σ represents a taste for variety. The price index that I derive next does not distinguish between these two interpretations, and neither do my empirical results.

3.1 *A price index*

This subsection discusses price indices able to measure the gains from density across locations. It is easy to derive the following - well-known - relative price index from the logit model:

$$R_{k,k'} = \frac{\left(\sum_{i=1}^{I_{k'}} p_{k'i}^{1-\sigma}\right)^{1/(1-\sigma)}}{\left(\sum_{i=1}^{I_k} p_{ki}^{1-\sigma}\right)^{1/(1-\sigma)}}. \tag{4}$$

$R_{k,k'}$ is the factor by which restaurant prices in area k would have to change in order to equalize utility in area k and k'. It is exactly the relative price index that would be derived from CES preferences, a result first shown by Anderson *et al.* (1992).[12]

It is useful to define the numerator in equation (4) as a variety-adjusted price index in area k, denoted by R_k, and the denominator as a variety-adjusted price index in area k', denoted by $R_{k'}$, so that for instance:[13]

$$R_k = \left(\sum_{i=1}^{I_k} (p + 2(\gamma t_{ki} + f_{ki}))^{1-\sigma}\right)^{1/(1-\sigma)}. \tag{5}$$

To compute welfare gains in monetary units, the relative price index must account for the degree of substitution between restaurants and all other goods, and for the possibility that expenditures on restaurants are higher in locations with a lower restaurant price index. To provide such an index, I solve a nested-logit model with one nest for restaurants and one nest for all other goods. Suppose that G_k is a price index for all other goods in location k. Then I show in Appendix A that the aggregate relative price index is:

$$P_{k,k'} = \frac{(R_{k'}^{1-\nu} + G_{k'}^{1-\nu})^{1/(1-\nu)}}{(R_k^{1-\nu} + G_k^{1-\nu})^{1/(1-\nu)}}, \tag{6}$$

[12] I do not use the exact price index proposed by (Sato (1976) and Vartia (1976)) for CES preferences, or the equivalent index that is robust to the introduction of new goods, introduced by Feenstra (1994). These indices are useful because they provide expressions in which the expenditure share on each variety captures an unobserved quality parameter, which disappears from the expression of the exact price index. In my framework, however, the particular quality parameter of any given restaurant is less relevant, because variation in prices come from variation in transport costs, not from unobserved quality. I include such a quality parameter, that I estimate directly for each types of cuisine, in the nested-logit model of section 8.

[13] This price index also has an interesting interpretation as what transportation researchers call a 'travel accessibility index'. Ben-Akiva and Lerman (1985) propose to use the denominator of a logit probability in a travel demand model as a travel accessibility index, and Niemeier (1997) is the first to estimate such an index in a model of mode choice and commute to different types of job. My index is easy to interpret because I use the linear utility specification of Anderson *et al.* (1992), and introduce value of travel time as a structural parameter.

where R_k is given by equation 5 and ν is the elasticity of substitution between restaurants and all other goods. This index obtained from a nested-logit model is exactly that which would be obtained from a nested-CES model, a correspondance first established by Sheu (2011).

I do not model residential choice, and I obtain the parameter ν from a literature review. Specifying a residential choice model could suggest that individuals choosing to live in high restaurant density areas have a preference for restaurant expenditures. Studies measuring ν do not necessarily account for this endogeneity problem.[14] However, I will show that welfare gains in dollars depend little on ν, but a lot on total restaurant expenditures and on the restaurant price index, the measurement of which is the focus of this paper.

3.2 *Maximum likelihood estimator for σ*

The elasticity of substitution σ is an unknown parameter, whose estimation is necessary to compute the gains from restaurant density. With data on multiple trips starting from the same location, one could estimate σ as the coefficient from an OLS regression of difference in log prices on difference in expenditure shares, for pairs of restaurants available from that location. In the dataset that I assemble in section 4, however, almost every trip originates from a different location, and no two locations offer an exactly similar set of available restaurants. I therefore propose a maximum likelihood estimator for σ that accounts for the exact restaurant choice set of each traveler. This simplest version of the estimator does not allow for variation in individual or restaurant characteristics.

Suppose that we have a sample of N trips to restaurants, indexed by n. Consider a trip of length t_{nk}, that originates in location k. From equation (2), one can write the predicted probability of that trip as a function of T_k, F_k and the parameter σ, so that:

$$prob(t_{nk}|\sigma,T_k,F_k) = \frac{(p+2(\gamma t_{nk}+f_{nk}))^{-\sigma}}{\sum_{i=1}^{I_k}(p+2(\gamma t_{ki}+f_{ki}))^{-\sigma}}.$$

The log-likelihood function is therefore:

$$\ell(\sigma,T_N,\mathbb{T}_K,\mathbb{F}_K) = \sum_{n=1}^{N} log(prob(t_{nk}|\sigma,T_k,F_k)),$$

[14] This endogeneity problem resembles that considered in Dubin and McFadden (1984), who note that electricity usage and appliance choice may share a common error term, or Goldberg (1998) who studies car usage and vehicle choice.

Figure 1: What identifies σ

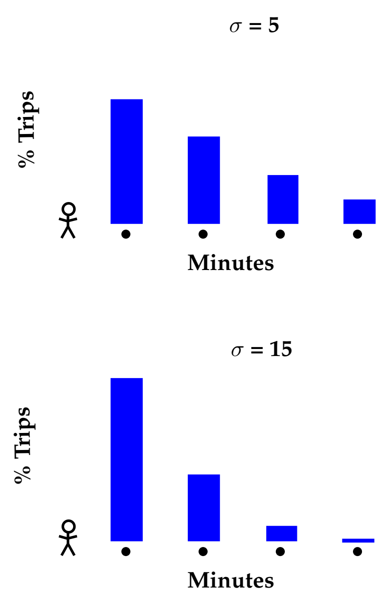

where T_N denotes the set of all trip lengths in the sample, $\mathbb{T}_K$ denotes the set of all sets T_k and $\mathbb{F}_K$ denotes the set of all sets F_k. The maximum likelihood estimate is the value of σ that maximizes the sum of log probability of observing each trip length in the sample:

$$\hat{\sigma} = \underset{\sigma}{\text{argmax}}\ \ell(\sigma, T_N, \mathbb{T}_K, \mathbb{F}_K). \tag{7}$$

Figure 1 illustrates how the distribution of observed trip lengths identifies σ in the model. All distances in the figure are in units of time, and each dot represents a restaurant available to an individual. The blue bar above each restaurant is the model's predicted probability of a trip to that restaurant. In the bottom diagram, σ is high at 15, so demand is very responsive to variation in prices due to transport costs. As a result, the model predicts that the probability of a trip to the closest restaurant is much larger than that to restaurants far away. Therefore, the estimator generates a high σ if the data features a large share of trips to the closest restaurant(s), in all locations. This pattern implies small gains from variety, because individuals are unwilling to incur transport cost to access a place that they prefer.

The elasticity of substitution σ determines the importance of gains from variety, but the model also allows to measure gains from shorter trips. Figure 2 offers a simple example of how gains from shorter trips arise in the model. The same σ is used to obtain predictions

Figure 2: Where do gains from shorter trips come from?

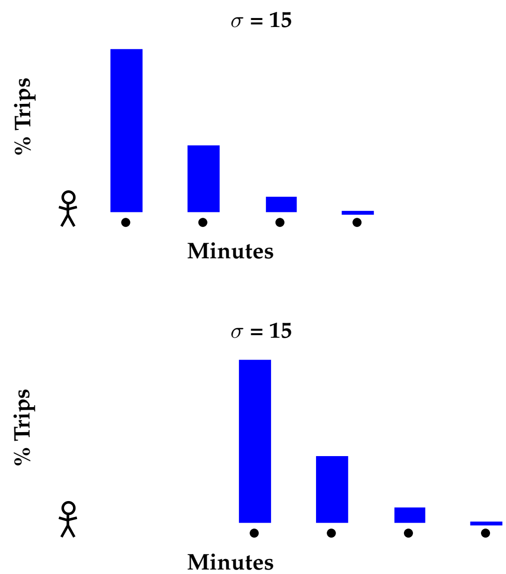

in both diagrams, and the proportion of trips to each restaurant is (almost) the same. The individual in the bottom diagram, however, takes on average longer trips, because she lives farther from the entire set of restaurants available.[15] This prediction is especially relevant given that a majority of Americans live in residential areas, relatively far from commercial zones.

An important restriction on the logit model is the independence of irrelevant alternative property . I now discuss how this property affects estimation results, and I show that it is testable. Intuitively, the IIA means that restaurants are infinitely differentiated, and destinations far from home are never 'similar' to other options closer from home. More precisely, it means that the relative demand for two restaurants depends only on their relative prices, and not on the price of any other restaurants. One important implication of this property is that, given a distribution of restaurants, adding another restaurant exactly at the location of each existing restaurant has no effect on average trip length. This must be the case, because such uniform increase in density maintains the proportion of trips of each length constant. When

[15]The model predicts that an individual living in a suburb five minutes away from the nearest restaurant drives on average approximately five minutes less on his trip than another individual living farther out, 10 minutes away from the same distribution of restaurants. Proposition 3 in Section B of the online appendix provides a formal statement and a proof.

estimating the model, this implies that variation in trip lengths across (uniform) density levels does not contribute to identifying σ. When testing the model, the prediction that a uniform increase in density does not increase trip time means that regressions of density on travel time (Section 7) are tests of the IIA property of the logit model.[16] In Section 8, I suggest extensions of the logit model to relax the IIA.

A number of other econometric and specification issues can confound my estimates. In general, insensitivity to price variation indicates a preference for higher quality products, not a taste for variety. In the model, however, the price variation that identifies σ originates from travel costs that are plausibly unrelated to quality differentials.[17] The same restaurant which is located 5 minutes away from an individual is also 25 minutes away for another individual living elsewhere. The main threats to the identification of welfare gains from density come from the possibility that restaurant or individual characteristics vary systematically with restaurant density. In Section 8, I include some restaurant characteristics in the model, and allow restaurants in the same chain to be perfectly substitutable, or restaurants serving the same type of cuisine to be more substitutable. These extensions also relax the IIA property of the logit model to address potential misspecification issues. I also estimate a model in which local restaurant supply reflects local tastes. This accounts for the possibility that individuals live close to their preferred restaurant types, and take shorter trips as a result. Note that models with restaurant characteristics are hard to estimate without data on restaurant choice, so they do not provide the preferred estimates in this paper. It is nevertheless important to emphasize that they generate results very similar to those obtained from the simpler logit model. I also verify that average restaurant prices and quality ratings vary little across density levels. To investigate heterogeneity along observable individual characteristics, I estimate σ separately by income groups, and I estimate a model in which travel speed and fuels

[16]Proposition 2 in the online appendix shows that this prediction holds using a more standard definition of a uniform increase in density, without fixed restaurant locations. Note that if the IIA does not hold, say because all burger restaurants are similar, then one should observe shorter trip lengths at higher density. The reason is, a high density area almost certainly features a burger place among the dozens of restaurants very close from home, removing the need to travel far for a burger. The IIA property of the logit and CES models is the object of valid criticism, for instance by Ackerberg and Rysman (2005), but ultimately its relevance is an empirical issue.

[17]There could still be small systematic differences between restaurants that are on average close to travelers and restaurants that are on average far, an issue that I address in the extensions of Section 8 by letting meal price vary with travel distance.

cost vary with individual characteristics like age, income, education and vehicle type. In section 7, I propose an instrumental variable strategy to assess the importance of sorting on unobservables (σ and γ) across density levels.[18] I find evidence of sorting by value of time, so I modify the maximum likelihood estimator to let γ vary with restaurant density.

4. Data

Estimating the logit model requires data on the location of a traveler, on the length of her trip to a restaurant, and on travel time and fuel costs to each restaurant available to her. The data on restaurant location come from the Google Places page of each restaurant in the summer of 2011 (these pages are currently called Google+ Local pages). The travel data, which identify trips to a restaurant, are from the 2008–2009 edition of the National Household Travel Survey (NHTS).

4.1 Restaurant data

Data from Google Maps applications offer complete coverage and exact information on restaurant location, both necessary for the innovations of the paper. As an aggregator of local business data, Google Places includes a page for any restaurant with a presence on alternative websites such as Yellow Pages, or an owner willing to create its own page. I collect data on all restaurants in a set of 14 US states containing more than 50% of the US population. I select these states because each of them funded the collection of additional data in my travel database, beyond the federally funded national sample.[19] My restaurant sample consists of 273,000 eating places.[20] My data includes fast food and full-service restaurants, as well as

[18]Endogenous restaurant supply can also lead to a positive relationship between σ and restaurant density. In a model with exogenous, uniform and continuous density of individuals, the density of restaurants in a free-entry equilibrium decreases with σ. I present a richer version of this model in the online appendix.

[19]The states in my sample are Arizona, California, Florida, Georgia, Indiana, Iowa, New York, North Carolina, South Carolina, South Dakota, Tennesse, Texas, Vermont and Virginia. I add Arizona, which purchased two regional-level add-on data, but no state-level add-on, and exclude Wisconsin, that purchased a state-level add-on which lacks geographical coverage. The states that I exclude do not have enough travel data to compute estimates of car speed at the local level, and too few trips to justify restaurant data collection.

[20]Within these states, the National Restaurant Association estimates the number of 'eating and drinking' places at 269,000, suggesting that my sample is comprehensive. The NRA's state-level reports are accessible at `http://www.restaurant.org/research/state/`. I also have a partial sample of 168,000 restaurants in other states of the country (for a total of 440,000 restaurants) to reduce measurement error from trips across state borders.

pubs, delis and other eating places. Coffee shops, such as Starbucks, are almost entirely excluded.

For robustness tests and extensions of the model, I also collect data on restaurant characteristics. The Google Places page of a restaurant provides the name of the restaurant and the type of cuisine that it serves (e.g. Korean, American, chicken, sushi). I code restaurants into 85 such categories, using definitions from Yelp.com, the most popular user review website for restaurants. I also identify restaurants belonging to the 50 largest restaurant chains in my sample, the largest of which is Subway. At the time of data collection, about 50% of Google Places pages contained a hyperlink leading to an alternative restaurant page on Yelp. Yelp contains information on average quality ratings from private reviewers (from 0 to 5 stars in 0.5 increments), prices ($, $$, $$$ and $$$$, that I code as $7, $17, $40 and $80),[21] number of reviews, and sometimes on attire, ambience, parking availability, and whether a reservation is necessary. A conflict between Google and Yelp occurred about a third of the way through data collection, after which Google removed the link to Yelp from its pages. I therefore have Yelp data for 70,000 restaurants, concentrated in the largest metropolitan areas due both to the data collection strategy and to the geographical preferences of Yelp's contributors.

4.2 Travel data

The NHTS is a nationally representative survey of travel behavior conducted about every six years. State transportation agencies can fund the collection of additional (add-on) travel data, which are also publicly available. There are data on 125,000 households in these add-on states that I restrict my sample to, representing 90% of the NHTS total. Each participating household member completes a travel diary on a travel day assigned to the household, recording the purpose, length, duration, start time and mode of every trip undertaken that day. Crucially, the data identify trips to 'get/eat meal', the origin of the trips (e.g. home) and the purpose of the next trip (e.g. return home). The data also contain a rich set of individual, household and trip characteristics, as well as the block group in which an individual resides. Trips to or

[21] On Yelp, the dollar signs represent the 'approximate cost per person for a meal including one drink, tax and tip': $ = under $10, $$ = $11-30, $$$ = $31-60, $$$$ = above $61.

from a restaurant represent about 11% of all trips, and about 25% of households have at least one member going to a restaurant on their travel day. The median trip to a restaurant is about 3 miles and lasts 10 minutes, with higher averages at 6 miles, and 14.5 minutes. About 90% of trips to a restaurant are by privately-operated vehicle ('car', for short) with almost all the remainder by foot.[22]

4.3 Sample selection

The travel data contains enough information to restrict the trip level estimation sample to a set of comparable trips with known origin whose only purpose is traveling to a restaurant. I describe the sample selection process below. In Section C of the online appendix, I estimate the model using different trip samples to show that although sample selection can remove biases in the expected direction, it has no sizable impact on welfare estimates.

I first restrict the sample to trips by car, which are more comparable (there is evidence, for instance, that walkers have a higher value of travel time). However, I will also compute price indices for pedestrians, using the σ that I estimate for car drivers. I eliminate the small percentage of car trips taken in high-density census tracts in which more than 20% of trips are by foot, because individuals in these areas may choose the car only for long trips, and walk for shorter trips.

About 40% of all trips to a restaurant start from home, and for the empirical analysis I restrict the sample to these trips, whose geographical origin is known at the block group level.[23] Importantly, the NHTS also allows to restrict the sample to trips that are immediately followed by a trip back home. When a trip from home is followed by a trip to a destination other than home, the benefits from reaching that next destination (e.g. a movie theater) also affect the travel decision. In these cases, which account for 33% of trips from home,

[22] Results sometimes differ depending on whether NHTS sampling weights are used. These weights account for the oversampling of some categories of individuals (e.g. older) in the NHTS. While most numbers are similar in both cases, the unweighted percentage of car trips is 92.5% while the weighted percentage is 88.3%.

[23] Block groups are small, which alleviates concerns about measurement error on the location of a traveler. The median radius of block groups is approximately 0.4 miles. I assume that each traveler resides at his block group's population-weighted centroid, which I obtain from the Missouri Data Center's MABLE Geocorr2K database. MABLE Geocorr2K computes a block group's centroid from the centroids of each of its constituent census blocks, using census block populations as weights.

observing long trips does not necessarily imply a high willingness to incur travel costs to visit to preferred restaurants.

I eliminate trips longer than 45 minutes, because to estimate the model I also limit the size of the restaurant choice set in each location to all restaurants available within 45 minutes of travel. I verify that cutting the sample at 30 or 60 minutes has little effect on my estimates, as individuals rarely take very long trips to a restaurant (45 minutes corresponds to the 98^{th} percentile of trip time.) Finally, I remove multiple observations of the same trip for different household members, and I keep only the driver's trip. To summarize, the trip level estimation sample consists of 7409 trips to a restaurant shorter than 45 minutes, from home and immediately back, by a driver. Estimation results are robust to using less restrictive samples.

4.4 Assembling the data

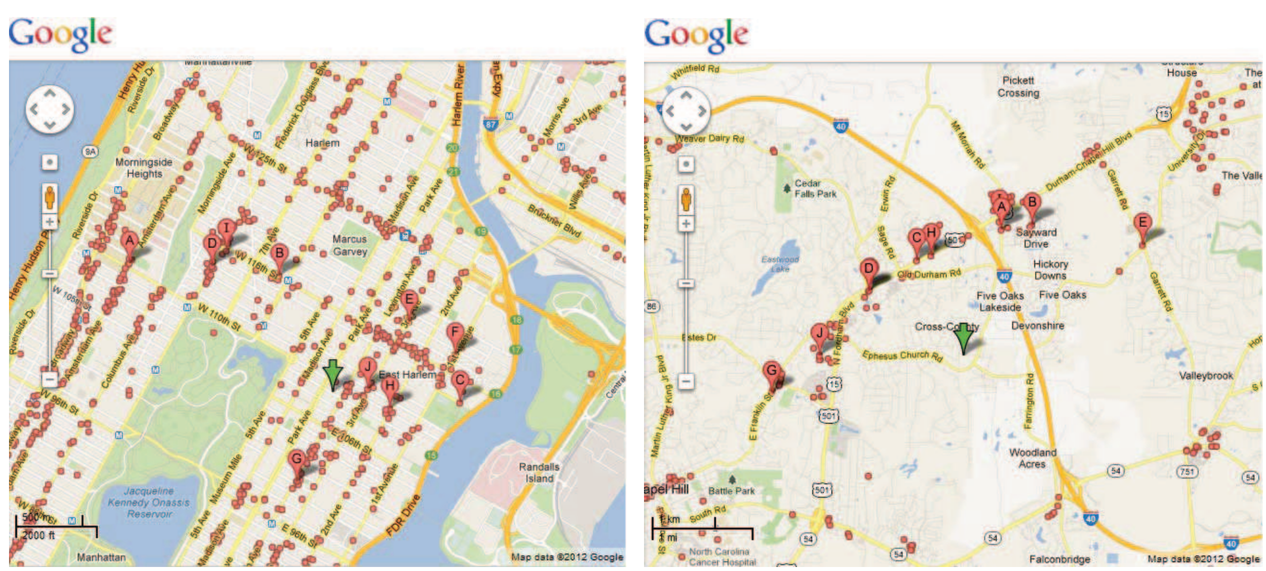

(A) A high-density urban area: East Harlem, Manhattan, New York City, NY (Maps data @2012 Google)

(B) A medium-density suburb in Chapel Hill, NC (Maps data @2012 Google)

Figure 3: Google Maps with restaurants

Notes: Each panel contains a screen-shot from Google Maps resulting from the search command 'Restaurants near [geographical location]'. The downward-pointing arrow indicates the location of an individual's population-weighted block group centroid. Each circle represents the location of a restaurant. The markers from A to G are Google's restaurant recommendations for the search. The scale of the map is at the bottom left. The map in panel B is at twice the scale of that in panel A. Google Maps is available at: `http://maps.google.com/`

Estimating the model requires computing, for each traveler, the travel time and fuel cost to each restaurant in her choice set. The two Google Maps in Figure 3 illustrate the matching between a traveler and her restaurant choice set. On each map, a downward-pointing arrow indicates an individual's residence, at her block group population-weighted centroid. Each circle represents the location of a restaurant (the alphabetical markers are Google's recommendations). The choice set of a traveler consists of all restaurants available within 45 minutes of travel.

To compute travel times from an individual to a restaurant, I first calculate the linear distance between the geographical coordinates of an individual's home and that of the restaurant. I then multiply this linear distance by a correction factor of 1.67, because the driving distance between any two points is longer than the length of the shortest path connecting these points.[24] Travel time, in minutes, is equal to distance times speed. I obtain measures of car travel speed for each trip using fitted values from regressions on the entire NHTS sample of car trips. In these regressions, speed varies with the census tract an individual lives in, with travel distance to the restaurant, and in some estimations with the characteristics of an individual (age, education, income, etc). The details of these regressions are in Appendix B.

Fuel cost depend on travel distance and speed, and in some estimations on the vehicle type of a traveler and on the price of gasoline in a location on a given day. The details of fuel cost construction is in Appendix B.

To assemble the final dataset used to estimate the model, I match each trip in the sample of 7409 trips to a restaurant with information (travel times and fuel costs) on all restaurants available within 45 minutes of the location from where the trip originates.

4.5 Spatial distribution of restaurants

Restaurants are far from being uniformly distributed in space, and a brief discussion of their spatial distribution illuminates many results in the paper. From the perspective of an

[24] I use a Google Maps application programming interface called Google Distance Matrix to obtain actual driving distance for a representative sample of individual/restaurant pairs (using only the 20 restaurants closest to an individual, which are most relevant). 1.67 is the average difference between the linear distance between two points and the driving distance from Google Distance Matrix. It would be possible, but prohibitively costly, to use the application to compute driving distances (or time) from all individuals to all restaurants in my sample.

individual traveler, the distribution of restaurants has two major characteristics. To varying degrees, these characteristics are apparent in the two maps from Figure 3. The individual in panel A lives in East Harlem, a high-density area in New York City. The individual in panel B lives in a medium-density suburban area of Chapel Hill, North Carolina. First, individuals live relatively far from the closest restaurant(s). Most Americans, like the individual in panel B, live in a residential suburb, at some distance from the nearest commercial outlets. Second, the number of restaurants available increases more than proportionally with distance (and time), and this increase is faster in denser areas. Both panels suggest that there are less restaurants available between 0–5 minutes of travel from home than between 5–10 minutes, and that this interval in turn contains less restaurant than that between 10–15 minutes, and so on. This result is a geometrical consequence of individuals living on a plane: the area accessible at any given travel distance increases with the square of that distance. To see why this increase is stronger in a high-density area, note that in panel A restaurants locate on a dense network of major urban roads crossing each other on the plane, and the geometrical argument above fully applies. However, low-density area are closer to a one-dimensional world, in which restaurants locate on the town's sole major road. Therefore, within a given travel distance (or time) interval, in dense areas a larger proportion of the mass of restaurants is located far from an individual. This feature of the restaurant distribution is central to the interpretation of regressions on measures of restaurant density, that I use to test the model.

5. Estimation of the logit model of travel demand

The estimation sample consists of all trips to a restaurant by a driver that are shorter than 45 minutes, start from home and are followed by a return trip home. Each of these trips is matched to the set of trip times and fuel costs to all restaurants that the traveler can visit within 45 minutes of travel. Data on observed trip time to restaurants is sufficient to estimate equation (7), as restaurants are only differentiated by transport costs from home, and by a random utility shock.[25] I set meal price at a constant value of p =$13. To set a value of travel time, I refer to Small and Verhoeff (2007), who review estimates of the value of driving time from a large literature, and suggest a value equal to 50% of a person's average hourly wage. I set γ = 0.2, which corresponds to $12 per hour, or about 50% of the average hourly wage in the United States.

I find the maximum likelihood estimate from equation (7) by grid-search and obtain $\hat{\sigma} = 8.8$ (in column 1 of table 1). A plot of the log-likelihood function suggests that it is concave for any reasonable values of σ, and therefore that $\hat{\sigma}$ is a global maximizer. Estimation results are robust to cutting or expanding the number of minutes in the choice set of individuals, for instance keeping 30 or 60 minutes worth of restaurants leads to $\hat{\sigma} = 8.8$ and $\hat{\sigma} = 9.0$. In column (2), I re-estimate equation (7), but letting travel speed to each restaurant vary with individual characteristics like age and income, and fuel costs vary with vehicle type. I find the same elasticity at $\hat{\sigma} = 8.8$. Letting speed vary with time of day (e.g. evening congestion) complicates data assembly but also leads to similar results.[26] This elasticity of substitution between restaurants is relatively large compared to existing estimates for consumer goods, but it is low enough to generate much extra travel beyond the closest restaurant, and as shown in section 6, substantial welfare gains. I am not aware of other estimates of the elasticity of substitution for services and non-tradables like restaurants.

One concern is that estimates of σ depend on the choice of value of travel time γ, a well-studied parameter, but that I do not estimate.[27] In Section C of the online appendix, I provide a sensitivity analysis showing how changes in parameters' value affect the welfare estimates. I find that welfare gains remain within the same order of magnitude even following large

[25]For about 15% of trips, observed trip time is shorter than my estimate of travel time to the closest restaurant, because of measurement error. In these cases, I assume that a traveler's location is such that trip time t_{nk} is exactly equal to travel time to the closest restaurant t_{k1}, and I add $t_{k1} - t_{nk}$ to the travel time of each restaurant in T_k. Such mistakes are generally small and occur for short trips to restaurants which are almost closest (e.g. someone enters a 5 minute trip - a round-up value - and I estimate that the closest restaurant is 5.5 minutes away). There are some larger discrepancies in low-density areas with large block groups and imprecise measurement of t_{k1}.

[26]I also estimate the model by weighted maximum likelihood, using NHTS sampling weight. These weights are equivalent to the number of individuals that each observation in the sample represents in the population, and account for over-sampling of some areas or type of households. I find $\hat{\sigma} = 9.2$.

[27]The source of variation that identifies σ (variation in trip length) is similar to that which which could identify γ. So I focus on providing the first estimate of σ in the service sector, and rely on a large literature estimating γ, generally from data particularly suited to this task.

Table 1: Maximum likelihood estimation of logit model of travel demand

	(1)	(2)	(3)	(4)	(5)	(7)
$\hat{\sigma}$	8.8	8.8	8.4	8.4	8.4	9.2
	(0.06)	(0.06)	(0.07)	(0.06)	(0.06)	(0.08)
$\hat{\beta}$			0.38			
			(0.02)			
$\hat{\mu}$						3.6
						(0.25)
Speed and fuel costs vary with individual char.		X				
Sorting by value of travel time			X			
Meal price varies with distance				X		
Perfect substitutability within chain					X	
Nested-logit (types of cuisine)						X
Observations	7409	6800	7409	7409	7409	7409

Notes: σ is the elasticity of substitution between restaurants, β captures the strength of sorting across density levels by value of travel time, and μ is the elasticity of substitution between different types of restaurant cuisine. Estimates obtained by grid-search in all columns. Standard errors in parentheses computed using the outer-product-of-the-gradient estimator, as suggested in Berndt, Hall, Hall and Hausman (1974).

perturbation like doubling γ or reducing it by half.[28]

5.1 *Can the logit model match the distribution of trip time in the data?*

As explained before, σ is identified from the distribution of trip time in the data. One can wonder how well the logit model performs in matching this distribution, after estimating only one parameter to fit the data. It is therefore instructive to visually compare the distribution of trip times in the data with a distribution of simulated trip times generated by the model at $\hat{\sigma} = 8.8$. To obtain a simulated trip dataset, I draw a trip time for each driver in the

[28] Another source of bias comes from drivers who undertake long trips to meet friends in a given restaurant, or to accommodate family members on the trip who have different preferences. First note that if the main purpose of the trip is to socialize with friends, then the trip should not be in the NHTS restaurant trip sample. In the reduced-form analysis of section 7, I control for many trips and individual characteristics and find that each additional adult on a trip adds 8% to trip length, and each children about 3%. Whether the passenger is a household member or not does not matter. To the extent that families eating out together are less likely to join friends at destination, then travel to meet friends does not explain the long trips observed in the data. The relatively small effect of additional passengers on trip time - some of which can be due to sharing fuel costs - also suggest that the logit model, in which one decision maker ignores the preferences of other passengers, provides a reasonable approximation of reality. Moreover, estimates of σ vary little when estimating the model under the assumption that two or more decisions-makers with exactly the same preferences make a joint travel decision.

Figure 4: Distribution of trip times, actual and simulated data

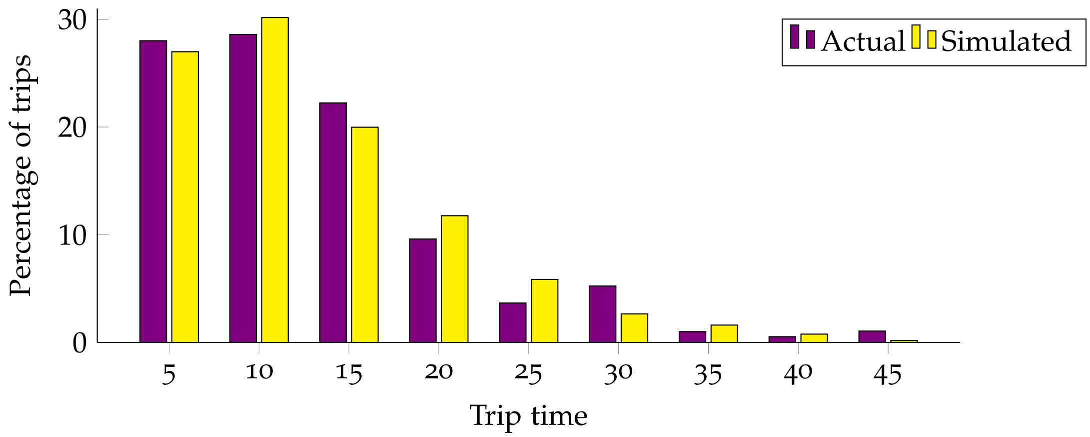

Notes: Sample of 7409 restaurant trips by a driver, starting from home and followed by a return trip home. The y-axis represents a time interval around the values indicated. This interval is 0 to 7.5 minutes for 5 minutes trip, 42.5 to 45 minutes for 45 minutes trip, and otherwise $x - 2.5$ to $x + 2.5$ for a x minutes trip.

sample using the probability distribution given by the model. Recall that the model predicts the probability of traveling to any given restaurant from each location, given the set of travel times and fuel costs to all restaurants and an estimate of σ . The results are in Figure 4, which shows that the proportion of trips within any given time interval looks remarkably similar in simulated and actual data.

6. Welfare gains from density

I now turn to estimating the welfare gains from restaurant density. I first compute the variety-adjusted restaurant price index in each block-group centroid in the sample, for both drivers and pedestrians. I then convert variation in the index into an average willingness to pay for restaurant density.

6.1 Welfare differences across areas

For each block group centroid in the sample, I use equation (5) to compute the variety-adjusted restaurant price index R_k at the maximum likelihood estimate of σ from column (1) in table 1. I compute the index both for a driver and for a pedestrian. Public transportation is not an option is most of America, and its use to eat out is negligible. I assume a constant walking speed of 3.5 miles per hour.[29] The value of the price index at any given location decreases with the number of nearby restaurants, and, for a driver, with car travel speed. Table 2 contains percentiles of the price index for a driver and for a pedestrian, with examples of locations at each percentile. For a car traveler, the median restaurant price index is equal to 9.3, which is lower than the average price of a restaurant meal ($13) before including transport costs. There is wide variation in the price index across areas. The index ranges from less than 6.5 in Manhattan and a few dense parts of San Francisco with faster car travel, to values around 10 in many of America's suburbs and small towns, and to values above 16 at the 99^{th} percentile of the index, in non-metropolitan areas with little gains from variety and hefty transport costs to reach even the nearest restaurant.

Shoup (2005) estimates that 99% of trips in the US end in a free parking, but in high density downtown areas like Manhattan, accounting for parking fees would no doubt increase travel costs for car drivers.[30] Table 2, however, shows that in Manhattan even pedestrians face a low index of about 7.3, which is lower than the index for a car driver almost anywhere else in the US where parking is free.

An important result not apparent in Table 2 is that much of the variation in the index occurs within metropolitan areas. Over the entire sample of block groups, the price index for a car driver decreases by 38% from its 90^{th} to its 10^{th} percentile. Within the ten largest MSAs in the sample, the average decrease from the 90^{th} to the 10^{th} percentile of the index is 23%. This decrease is lowest in Miami at 14% and highest in New York, Houston and Atlanta at close to 30%. These large within-MSA variations in the index reflect the highly localized nature of the gains from restaurant density. Remote restaurants, that are expensive because of travel costs, have little impact on welfare. For instance, removing access to all restaurants

[29]Walking speed in Google Maps applications is about 3 miles per hour. Average reported walking speed in the NHTS is 4.4 mile per hour, and even faster for non leisure trips like walking to a restaurant, which appears implausibly fast. So I choose mid-range value of 3.5 miles per hour.

[30]The time costs of cruising for parking and the walk between a parking spot and a restaurant is already included in trip time (if survey respondent filled their diary properly), and just translate into lower speed. Meters in lower Manhattan charge $3 per hour, and the average restaurant meal lasts 30 minutes, meaning that accounting for parking could increase the price index in Manhattan by as much as 1.50, probably less if some free parking is also available.

Table 2: Percentiles of the restaurant price index

Percentile	Index by car	Index by foot	Example of location
Minimum	6.0	7.3	Manhattan
1^{st}	6.8	8.8	San Francisco
5^{th}	7.6	10.1	Downtown or near in most large cities.
10^{th}	7.9	10.7	Median location in Los Angeles County
25^{th}	8.5	11.7	Suburb/outskirt big central city, Downtown medium-sized city
50^{th}	9.3	13.4	Suburb
75^{th}	10.7	17.5	Remote suburb or small town
90^{th}	12.7	28.9	Country-side
95^{th}	14.0	38.7	Country-side
99^{th}	16.6	57.6	Country-side

Notes: The price index is computed using the estimate $\hat{\sigma} = 8.8$ from column (1) of Table 1. The percentiles are computed over all 51641 block groups in which there is at least one individual in the NHTS sample. The first row contains the lowest values in the sample.

lying between 30 and 45 minutes of travel reduces the price index by on average only about 2%.[31] The conclusion that the gains from restaurant density are localized may generalize to much of the consumption benefits of density, given relatively short trip times for most types of non-work trips in the NHTS.

A decline in the variety-adjusted restaurant price index translates into sizable welfare gains for an average household. For these welfare computations, I set the price elasticity of demand for restaurants at -1, consistent with a literature review by Okrent and Alston (2010) who find an average value of -1.02 for the price elasticity of demand for food away from home. I take expenditure shares from the Consumer Expenditure Survey (CEX) 2009, in which food away from home represents on average 5.3% of household expenditure, or \$2619 out of average expenditures of \$49,067.[32] Using these numbers, and assuming that the price index for all

[31]This implies that individuals do not need perfect information on thousands of remote restaurants for my estimates to be valid. I also estimate a version of the model in which restaurants farther away from an individual are known (i.e. part of her choice set) with a probability parameterized such that 100% of restaurants at 0 minute from home are known, and the probability of knowing restaurants farther away decreases with travel time. I find that individuals know about 69% of restaurants 45 minutes away, which barely affects the welfare estimates, but the results are not precise. I also estimate, by simulated maximum likelihood, a model in which the scale of the type I extreme value distribution of the error term decreases with distance, and obtain similar results.

[32]The CEX considers transport costs separately from other expenditures, and it includes coffee, ice cream and snacks in food away from home, which my index does not take into account. Also, the index is only accurate for trips starting exactly from home, while some restaurant expenditures are incurred away from home.

other goods does not vary across areas, one can compute the aggregate relative price index in equation (6). In reality, the price index for all other goods *does* vary across areas, and in particular we expect higher house prices to cancel out the welfare gains from density. The goal of this exercise is to compute willingness to pay for spatial variation in the availability of a consumption amenity, holding everything else constant. I find that an average household's willingness to pay to enjoy a 20% decrease in the restaurant price index, which is equivalent to moving from a low to a high density part of a large MSA, is about $576 annually (as intuition suggests, $\$576 \approx 0.2 * \$2619$, but see Appendix A for the details). The welfare estimates are not very sensitive to changing the price elasticity of demand for restaurants. Using a high elasticity of -2 instead of -1, I find that for an average household, the gains from a decrease in the price index are larger by about 7%.

In section 7, I use regression analysis to show that most of the welfare gains from an increase in restaurant density are gains from variety, as opposed to gains from travel time savings through shorter trips.

In Section D of the online appendix, I estimate welfare gains separately by income group, taking into account higher meal price, higher value of time and faster travel speed for richer individuals. Richer household do enjoy considerably larger gains from restaurant density, but simply because of higher expenditures on restaurants. Higher value of time and a slightly lower elasticity of substitution make small contributions towards increasing gains from density for high income individuals, but the relative prices computed from Table 2 are overall quite similar across income group.[33]

It is interesting to compare these results with those of Handbury and Weinstein (2012), who estimate a variety-adjusted price index for tradable consumer goods (groceries with barcodes) in different MSAs. They find that residents of larger cities, controlling for store amenities, individual characteristics and differences in the number of varieties available, face a lower price index for groceries, a result entirely due to the availability of more varieties in larger cities. This price index drops by 5% from New York City, whose residents have access to

[33] Existing paper on how gains from variety vary by income group include Lee (2010), Li (2012) and Handbury (2012).

110,000 types of groceries, to Des Moines, the smallest city in their sample, whose residents have access to 24,000 grocery types. While these numbers are estimated for large areas and do not take transport costs into account, simple comparisons with my results for the restaurant industry suggest much larger spatial welfare differentials in the non-tradable service sector. Some residents of the New York or Los Angeles metropolitan areas have access to more than 20,000 restaurants within 45 minutes of car travel, versus only 800 in Des Moines, a city with faster travel speed. Moving from the densest part of New York City to the densest part of Des Moines leads to a 30% reduction in the price index (some of which is due to lower travel costs).[34]

Future research may demonstrate that individuals derive similar benefits from the higher density of health providers, entertainment options and other services in the downtown cores of large metropolitan areas. In part, the dominance of non-tradables over tradables in explaining the consumption advantage of cities depends on the highly developed supply chains of the major consumer goods retailers. The low cost of moving goods enables the provision of an impressive array of consumer goods to America's suburbs and smaller towns. Such feats of logistics are not easily replicated in the non-tradable service sector, which depends to a larger extent on the movement of people. Dense urban areas still have a unique advantage in reducing transport costs between individuals.

6.2 Discussion: Accessibility

The relatively low median value of the index for car travelers supports the argument in Glaeser and Kahn (2004) that the suburban lifestyle shared by a majority of Americans offers good accessibility through fast car travel. Unsurprisingly, in most of America gains from variety

[34]Handbury and Weinstein (2012) note that additional varieties in larger cities account for a relatively low share of expenditures, which explains the relatively small gains from variety that they measure. In Section 7, I run regressions of density on trip time to test whether the large number of additional restaurants in dense areas are superfluous (if they are, then travelers in dense areas should make very short trip). The finding that additional restaurants in dense areas are not superfluous perhaps reflects how markets in non-tradables respond to local tastes. This is indeed the assumption that I make when I estimate tastes for different cuisines in a nested-logit model (Section 8). If, similar to Handbury and Weinstein (2012), I assumed instead that tastes were constant across the United States, then percentile differences in the index across and within cities would stay relatively similar, but, for instance, some areas in Texas with a vast majority of Mexican restaurants would wrongly receive a high price index, because individuals in these areas are not allowed to have a special taste for Mexican food.

are much higher when traveling by car. Driving only looses its attractiveness in very high density areas, in which low travel speed brings the price index for walker and drivers much closer together, as shown in Table 2. This explains why a sizable majority of Manhattanites walk to a restaurant.

In fact, the lowest price index, or equivalently the best accessibility, belongs to areas with the *slowest* car travel. Faster travel speed mechanically decreases the price index, but the correlation between speed and the index is still positive. Raising the density of destinations, of population and of the street network reduces travel speed, but not enough to annihilate the benefits from greater access to destinations. Clearly, then, it is possible to reach a level of density and "walkability" such that vehicle use is greatly reduced, as argued by New Urbanism proponents. This only happens, however, at very high levels of density that are rarely seen in the modern American landscape.

7. Reduced-form analysis

This section uses regression analysis first to test the logit model, second to investigate the possibility that individuals sort by density levels according to unobservable characteristics, and third to break down the gains from density into gains from variety and gains from shorter trips.[35]

7.1 Variable construction for reduced-form analysis

The reduced-form analysis requires measures of restaurant density. I therefore propose four variables, which captures the features of the restaurant distributions highlighted in section 4. I compute these measures using average speed in a location, so they do not depend on individual characteristics.

1. Travel time to the restaurant closest from home.

[35] Regressions on the determinant of trip time are interesting in their own rights. For instance, the relationship between travel and the built environment is crucial when evaluating urban development schemes designed to reduce vehicle travel. The online appendix presents complementary results on the determinants of the probability of making a restaurant trip.

2. Local density: a measure of local restaurant density passed the closest restaurant. Local density, in restaurants per minute, is equal to travel time to the 20^{th} closest restaurant minus travel time to the closest restaurant, divided by 19.

3. Global density: a measure of restaurant density for an area wide enough to encompass most trips, but not so large as to be irrelevant to a traveler. Global density, in restaurants per minute, is equal to the number of restaurants available within 45 minutes of travel, divided by 45.[36]

4. Skewness: a measure of whether most of the restaurant mass is distributed close to or far from an individual. This ratio is equal to the density of restaurants from 22.5 to 45 minutes of travel over the density of restaurants from 0 to 22.5 minutes of travel.

7.2 *Testing the model: regressions on measures of restaurant density*

A key prediction of a travel demand model with gains from variety is that increasing restaurant density has little effect on trip time, because density makes it cheaper to substitute between restaurants and to visit a location that one prefers. The first step in testing this prediction is to state it precisely by finding the exact impact that measures of restaurant density would have on trip time if the logit model were true. To do so, I create a variable equal to predicted average trip time for each traveler in the sample, and I use this variable in regressions on measures of restaurant density.[37] Then, I run the same regressions, but using actual trip time as a dependent variable. If regressions on actual and predicted trip times generate similar coefficients, then the model accurately predicts the actual impact of restaurant density on trip time.

Regressions on predicted trip time To obtain predicted average trip time, note that the model predicts the probability of traveling to any given restaurant from each location, given T_k, F_k and $\hat{\sigma}$. Using the predicted probability of a trip of each length in an area k, I compute $\bar{t}_{nk}$,

[36] I could define global density starting starting from the closest restaurant, to be consistent with the measure of local density. Such a definition would lead to similar regression results, but they would be harder to interpret.

[37] Note that if the distribution of restaurant was always uniform, the predicted impact of an increase in density on trip time would be 0 because of the IIA, and this step would not be necessary.

Table 3: The determinants of trip times, predictions from the logit model

	(1)	(2)	(3)	(4)	(5)
log Predicted average trip time ($\bar{t}$)					
log Global density	0.089^a		0.138^a	0.210^a	0.207^a
	(0.004)		(0.004)	(0.004)	(0.004)
log Skewness		0.062^a	-0.025^a	-0.089^a	-0.079^a
		(0.005)	(0.004)	(0.004)	(0.004)
log Time to closest rest.			0.518^a	0.588^a	0.577^a
			(0.009)	(0.008)	(0.006)
log Local density				-0.210^a	-0.207^a
				(0.006)	(0.005)
Observations	7407	7405	7405	7405	7405
R^2	0.09	0.05	0.59	0.73	0.78

Notes: OLS regressions with a constant in all columns. Robust standard errors, clustered at the county level, in parentheses. a, b, c: significant at 1%, 5%, 10%.

the model's prediction of expected trip time for each trip n in the sample. It is important to emphasize that the regression coefficients that I report are not sensitive to the particular value of $\hat{\sigma}$ that I use to obtain predictions. That is, I am not testing a fitted model, but rather a general property of logit models. I now run OLS regressions using $\bar{t}_{nk}$ as a dependent variable. The estimating equation is:

$$log(\bar{t}_{nk}) = \alpha + \beta_1 density_k + \epsilon_{kn}, \tag{8}$$

where $density_k$ represents characteristics of the restaurant distribution in area k (travel time to closest restaurant, global density, local density, and skewness). I estimate this equation on the same sample of restaurant trips that I used to estimate the model (trips shorter than 45 minutes by a driver, from home and back). Table 3 reports the regression results.

All the coefficients in table 3 are elasticities. The four measures of restaurant density are correlated with one another, so the coefficient on each variable is sensitive to the inclusion of the others. The predicted effect of global density on trip time is positive, with an elasticity ranging from 0.09 in column 1 when the variable enters alone, to 0.21 in column 5 when all measures of density are included. This corresponds to trips about 2 to 4 minutes longer than average in the densest decile of global density. Recall that a uniform increase in restaurant density does not affect travel time, so the effect of global density on trip time would be about zero if restaurants were uniformly distributed. Therefore, this positive coefficient is due to the larger share of the restaurant mass located far from home in dense areas (Secton 4 provides a geometrical argument to explain this feature of the restaurant data). In column 2, the skewness of the restaurant distribution enters the regression alone, and its coefficient is positive at 0.06, so the model predicts that trips are longer in areas with a disproportionately large number of restaurant far from home. Unlike the effect of global density, the effect of local density is negative (elasticity of -0.23, in column 5). High density close to home implies relatively lower density far from home, hence the shorter predicted trip times. The predicted effect of time to the closest restaurant is harder to interpret as an elasticity, but its linear effect is close to 1 (about 1.2). So the model predicts that living 1 minute farther from the closest restaurant increases trip time on average by about 1 minute.

Regression on actual trip time I now run regressions on the determinants of actual trip time using the same sample of trips. The estimating equation is:

$$log(t_{nk}) = \alpha + \beta_1 density_k + \beta_4 X_n + \beta_5 Z_n + \epsilon_{nk}. \tag{9}$$

Each observation n is a trip. The dependent variable t_{nk} is trip length in minutes. The independent variables are as in equation (8), except that in some specification I add a vector of individual characteristics X_n for the driver of trip n, such as age, income, speed and fuel costs, and a vector of trip characteristics Z_n, such as the number of individuals on the trip, and the time spent at destination. Table 4 reports the regression results.

Columns 1–4 present regression results with only measures of restaurant density as controls.[38] The effect of global density on trip time is close to zero in each regression, meaning that the number of restaurants available within 45 minutes of travel has little impact on trip time. That increasing global density fails to reduce trip time in the data corroborates the intuition behind the model, that substituting among travel destinations is cheap in dense

[38]The low R^2 of these regressions is consistent with the inherent randomness in discrete-choice decision making. The R^2s in Table 3 are higher because the dependent variable in these regressions is average predicted trip time. In this case, the R^2 reflects the ability of measures of restaurant density to capture the features of the restaurant distribution that determine average predicted trip time.

Table 4: The determinants of trip time

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
log Trip time to restaurant							
log Global density	-0.017^{b}		0.002	0.037^{a}	0.033^{a}	0.036^{a}	0.017
	(0.007)		(0.008)	(0.009)	(0.009)	(0.009)	(0.023)
log Skewness		0.036^{a}	0.009	-0.012^{b}	-0.009	-0.007	0.009
		(0.006)	(0.007)	(0.007)	(0.008)	(0.007)	(0.012)
log Time to closest rest.			0.232^{a}	0.271^{a}	0.252^{a}	0.231^{a}	0.231^{a}
			(0.018)	(0.018)	(0.018)	(0.017)	(0.024)
log Local density				-0.111^{a}	-0.112^{a}	-0.098^{a}	-0.123^{a}
				(0.013)	(0.014)	(0.013)	(0.020)
Controls							
Individual characteristics					X	X	X
Trip characteristics						X	X
MSA fixed-effects							X
Observations	7407	7405	7405	7398	6791	6712	5543
R^2	0.001	0.005	0.035	0.043	0.070	0.168	0.173

Notes: OLS regressions with a constant in all columns. Robust standard errors, clustered at the county level (except in column 9), in parentheses. *a*, *b*, *c*: significant at 1%, 5%, 10%. Dependent variable is log trip time to restaurant in all columns. Individual characteristics include 17 dummies for household income, 4 dummies for education, household size, 6 dummies for age, a dummy for gender, a dummy if black, a dummy for worker's status, the speed of a one mile trip (which depends on individual characteristics and census tract), and gasoline costs per mile at medium speed (which depends on vehicle type and gas prices). Trip characteristics include 5 dummies for each peak hour (7-8am, 8-9am, 15-16pm, 16-17pm, 17-18pm), a dummy for trips on week-end, the number of children on the trip, the number of adults on the trip, the number of non-household members on the trip, and the log time spent at destination.

areas, so individuals often gain from density by visiting preferred location. However, this zero effect does not match the model's prediction of a positive effect. This discrepancy could mean that some assumptions underlying the logit model are too strong.

As predicted, the elasticity of trip length with respect to travel time to the closest restaurant is large and positive, ranging from 0.23 to 0.27 depending on the regression. The elasticity of trip time with respect to skewness is 0.036 when it enters alone in column 2. As predicted, individuals take shorter trips if the mass of restaurant within 45 minutes is disproportionately located close to home. Local density has the predicted negative effect on trip time (elasticity of -0.13), so travelers with a high density of restaurant close to home (passed the closest restaurant) make shorter trips. A general conclusion is that increasing the level of density over a large area has little effect of trip time, but the decision of how far to travel strongly depends on the spatial distribution of restaurants, in ways that can be predicted.

In column 5 and 6, I add controls for individual and trip characteristics to the regression. These additional controls barely affect the coefficients on measures of restaurant density. In column 9, I add MSA fixed-effects and find similar coefficients on measures of restaurant density.

To summarize, the logit model matches the first-order features of the data, but there is a discrepancy between the actual and predicted effect of the number of restaurants within 45 minutes of travel (global density) on trip time. I offer three explanations. First, the IIA may not hold, and remote restaurants may be close substitutes for options available closer from home, so the mass of restaurants far from home in dense areas exert little attraction on a traveler. Two extensions of the model in section 8 relax the IIA property of the logit model. Second, measurement error biases OLS estimates towards zero. Third, omitted variables can bias the coefficient on global density. An instrumental variable strategy alleviates both measurement error and omitted variable biases.

7.3 IV estimation

If individuals sort into areas based on γ or σ, then OLS coefficients are biased because of a correlation between the error term in equation (9) and measures of restaurant density. For instance, the model predicts that individuals with high value of travel time make shorter trips. Therefore, sorting of high γ individuals into dense areas could explain why the model - in which there is no sorting - overestimates trip length in areas with high global density. If this is the case, the IV coefficient on global density will be more positive than its OLS counterpart, and closer to the model's prediction. The reverse happens if individuals with marginal preferences or pronounced taste for variety sort into dense areas; with sorting on σ the coefficient on global density should become even more negative if instrumented.

An instrument z_k for global density in location k must satisfy two criteria. First, it must be relevant, i.e. correlated with global density conditional on other controls:

$corr(global\_density_k, z_k|controls) \neq 0$. Second, the instrument must be exogenous, i.e. uncorrelated with the error term: $corr(\epsilon_{nk}, z_k|controls) = 0$. I instrument global density with the growth in population density from 2000 to 2007, in the county in which an individual lives. Crucially, I am able to select a sample of individuals - old, married, homeowners - who are very unlikely to move out of county in any given year.

The county population data come from the 2000 Census and from the 2005-2009 population averages from the American Community Survey, that I use as a measure of county population in 2007.[39] Population density for each county is population count over area. The growth in population density from 2000 to 2007 is the ratio of the log density in 2007 to that in 2000.

While counties vary in size, they are the census geographic units that most closely match an area accessible through 45 minutes of travel.[40] Current county population density is a strong predictor of restaurant global density. More important, growth in county population density in the 2000s explains variations in the level of restaurant global density in 2011, especially if one controls for initial county population density in 2000.[41] So the instrument is relevant.

If individuals sort into densely populated areas based on unobserved characteristics that affect trip time, then the instrument fails the exogeneity condition. Using growth to instrument a level is an important step towards satisfying the exogeneity condition. I can control for the initial level of population density in 2000, and people probably seldom choose to reside in an area based on an accurate prediction of its density growth prospect a few years hence. Therefore, the main threat to the exogeneity condition comes from individuals who moved between 2000 and 2007 into areas whose population densities were high because of recent growth. In this case, there is sorting on the instrument. This is a particular concern because 15.4% of Americans surveyed by the 2009 ACS had changed residence over the previous year, according to Ihrke, Faber, and Koerber (2011). To remedy this, the identification strategy relies on creating a sample of individuals with a low probability of moving out of county in any

[39]The 2009 ACS is the last to use Census 2000 geography, and using averages allows to experiment with instruments computed at the census-tract level, at which yearly data is not available.

[40]The median county in my sample has a radius of about 38 miles, while 45 minutes of driving usually covers about 25-30 miles.

[41]Without a control for initial population density, the instrument is marginally weak (columns 1 and 5 of Table 5).

given year. The moving rate of individuals aged 45 and older is about 7%, with only a 3% chance each year of moving out of county. Homeowners have a 6.7% moving rate, almost five times smaller than that of renters. Married individuals also have a lower than average moving rate at 9.9%. Keeping in mind that older individuals are also more likely to be married and to own a home, suppose that 2.5% of married homeowner 55 years and older, that I keep in my sample, randomly move out of county every year.[42] In this case, more than 80% of these travelers that I observe in 2008 and early 2009 lived in the same county in 2000. I also verify that OLS regressions on this restricted sample lead to similar results as regressions on the full sample. The summarize, the IV regressions are informative, but one should treat the results with caution.

Table 5 contains the two-stage least squares estimation results. The estimates in columns 1–4 are for the full sample, and those in columns 5–8 are for the sample of individuals with a low probability of moving. In each column, the elasticity of trip time with respect to global density is positive and significantly larger than any of the OLS elasticities. This result is consistent with the sorting of individuals with high value of travel time, who are predicted to make shorter trips, into dense areas. The elasticities range from 0.07-0.12 in columns 1–4 and are twice as large in columns 5–8, suggesting that sample selection is important for identification. For this sample of individuals with a low probability of moving, one cannot reject that the effect of global density on trip time is the same as that predicted by the model. The regression with all measures of density, in column 7, generates results remarkably similar to those obtained from predicted average trip time data in column 5 of table 3. The IV strategy probably mitigates the impact of measurement error, as it moves all coefficients on measures of restaurant density away from 0 and towards the model's prediction.[43]

[42] I keep individuals 55 years and older instead of 45 years and older because a 45-year-old in 2008 was 37 in 2000. I observe both age and homeownership status in the NHTS and I select individuals living in households with two or more members to proxy for marital status.

[43] The variables for local density and travel time to the closest restaurant are also endogenous, but the instrument (even if defined at the census tract instead of at the county-level) is much weaker for these variables, and these IV regressions are sensitive to the set of controls and often lead to unreasonable values.

Table 5: The determinants of trip time, with instrument for global density

Trip time to restaurant	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8))
log Global density	0.085	0.078^c	0.118^b	0.099^b	0.180^b	0.161^b	0.212^a	0.169^b
	(0.054)	(0.047)	(0.050)	(0.049)	(0.077)	(0.065)	(0.070)	(0.082)
log Time to closest rest.			0.249^a	0.262^a			0.340^a	0.262^a
			(0.022)	(0.024)			(0.034)	(0.036)
log Skewness			-0.050^a	-0.032			-0.075^a	-0.055^b
			(0.018)	(0.02)			(0.023)	(0.025)
log Local density			-0.124^a	-0.163^a			-0.188^a	-0.147^a
			(0.019)	(0.023)			(0.033)	(0.033)
Controls								
log County pop. density in '00		X	X	X		X	X	X
Individual characteristics				X				X
Trip characteristics				X				X
Instrument								
Δ log County pop. density '00-07	X	X	X	X	X	X	X	X
Samples								
Full sample	X	X	X	X				
Low probability of moving					X	X	X	X
Observations	7390	7390	6704	6587	3528	3528	3182	3126
First-stage stat.	14	31	30	18	11	27	28	17

Notes: Two-stage least squares regressions with a constant in all columns. Robust standard errors, clustered at the county level in parentheses. a, b, c: significant at 1%, 5%, 10%. Dependent variable is log trip time to restaurant in all columns. The sets of individual and trip characteristics are the same as in Table 4. The sample with a low probability of moving consists of all homeowners 55 years and older living in an household with at least 2 members. The first-stage F statistics are cluster-robust.

7.4 *Where do welfare gains from density come from?*

Regressions analysis also allows to assess the relative importance of the two sources of gains from density in the model: gains from variety and travel costs savings. For a given decrease in the index, I find the share explained by travel time and fuel cost savings, and I attribute the remainder to gains from variety. I run regressions of restaurant trip time on 9 dummies for the deciles of the variety-adjusted restaurant price index R_k, computed using the estimate of $\hat{\sigma}$ from column 1 in Table 1.[44] I omit the dummy for the first decile, to obtain the average

[44] I include trips longer than 45 minutes in the sample, because in these regressions their inclusion affects the result, as very long trips are significantly more likely in low density areas with high values of the price index.

increase in trip time within each upper decile. From row 1 of Table 6, there is no significant differences in trip length between locations within the first and second deciles of the price index. The index, however, increases from 7.5 to 8.1. So for an individual in the second decile moving to the first decile, I attribute 100% of the \$0.6 decrease in the index (the gains from density) to gains from variety and 0% to gains from shorter trips. Within the last decile, the index is equal to 14.5, so there is a \$7.0 difference in the index between the first and last deciles. Using the travel time difference on a one way trip of about 6.25 minutes from row 10, I derive a \$3.1 difference in travel costs between the first and last deciles.[45] So I attribute about 3.1/7 = 44% of the gains from density to shorter trips. Repeating the same exercise for the remaining deciles leads to shares of travel costs savings in the gains from density that are lower than 44%.[46]

It is instructive to consider additional, non-structural, evidence on the importance of gains from variety in denser areas. The number of restaurants passed by a traveler on her way to her final destination is arguably a measure of whether an individual chooses to visit a 'preferred' destination. The median number of restaurants passed is 14 within the first decile of global density, and it rises steadily up to 89 in the last, denser decile. These numbers strongly suggest that travelers in dense areas often visit destinations closer to their ideal.

These result imply that the main effect of policies promoting denser developments is to allow individuals to visit places that they prefer. Densification can also lead to gains from shorter trips, especially for very low density areas in which even the closest restaurant is far from home.

8. Extensions of the logit model

In this section, I estimate four extensions of the logit model of travel demand. First, I allow for sorting into denser areas by value of travel time, in line with the IV regression results.

[45] Time saving on a round trip is $6.25 * 2 = 12.5$ minutes, valued at $(12.5/60) * 12 = 2.5$ dollars. Fuel costs for a 12.5 minute trip, at average speed, fuel efficiency, and gas prices is about \$0.6, which leads to total travel costs savings of $2.5 + 0.6 = 3.1$ dollars.

[46] The positive effect of density on trip length in IV regressions suggest that the share of travel cost savings in the gains from density is perhaps even smaller. However, I do not have a good instrument for the position of restaurants very close to home, that most affect the index.

Table 6: Gains from shorter trip time

Travel time	(1)	(2)	(3)
Dummy 2^{nd} decile of price index	-0.23	-0.74	-0.90^c
	(0.56)	(0.52)	(0.49)
Dummy 3^{rd} decile of price index	0.64^a	0.37	0.12
	(0.52)	(0.52)	(0.49)
Dummy 4^{th} decile of price index	1.93^a	1.41^b	1.04^b
	(0.55)	(0.57)	(0.52)
Dummy 5^{th} decile of price index	1.43^a	1.01^c	0.49
	(0.52)	(0.54)	(0.51)
Dummy 6^{th} decile of price index	2.93^a	2.70^a	2.05^a
	(0.54)	(0.56)	(0.52)
Dummy 7^{th} decile of price index	2.12^a	2.10^a	1.80^a
	(0.53)	(0.55)	(0.51)
Dummy 8^{th} decile of price index	3.33^a	3.44^a	2.70^a
	(0.55)	(0.57)	(0.54)
Dummy 9^{th} decile of price index	4.78^a	5.18^a	4.36^a
	(0.61)	(0.64)	(0.61)
Dummy 10^{th} decile of price index	6.52^a	6.65^a	5.60^a
	(0.88)	(0.93)	(0.88)
Controls			
Individual characteristics		X	X
Trip characteristics			X
Observations	7510	6896	6896
R^2	0.03	0.05	0.14

Notes: OLS regressions in all columns. Robust standard errors, clustered at the county level, in parentheses. a, b, c: significant at 1%, 5%, 10%. The dependent variables is trip time in minutes. The excluded price index dummy is that for the 1^{st} decile of the index. The sets of individual and trip characteristics are the same as in Table 4.

Second, I let meal prices vary with travel time from home, to investigate a potential source of bias in my estimates. The last two extensions relax the IIA property of the logit model and introduce restaurant diversity into the model. Removing the IIA implies shorter trips in denser areas, because a restaurant far from home in a dense area is often similar to one of the many alternatives closer to home. That is, these extensions can explain why the OLS effect of global density on trip time is smaller than what the logit model predicts.

8.1 *Sorting by value of time*

IV regressions indicate that individuals who choose to live in high restaurant density areas make shorter trips, and the model suggests that sorting by value of travel time can explain this result. Hence, I estimate an extension of the basic logit model in which an individual's value of travel time γ is a function of global density in location k. I estimate σ jointly with a new parameter β capturing the strength of the relationship between γ and density. The estimation results are in Table 1 and the details are in Section F of the online appendix. The elasticity of substitution $\hat{\sigma} = 8.4$ is slightly lower than in the model without sorting, which implies marginally higher gains from density. As expected, the model with sorting predicts the OLS regression results of a near zero effect of global density on trip time. Recall that the model without sorting predicts the IV regression results of a positive effect of global density on trip time.

8.2 *Variation in meal prices*

Information from Yelp.com allows to assess the estimates' sensitivity to price variation across areas. A particular concern is that higher density implies higher restaurant prices or better quality. I only have Yelp data for 23% of restaurants, all in the 20 largest MSAs in the sample. There is only a small positive correlation between average meal price and restaurant density within 45 minutes of travel, which is too small to significantly impact welfare estimates. Price differences across MSAs are also too small to affect the welfare results. Average quality ratings have almost no correlation with density levels, and display little geographic variation.[47] One should interpret these results with caution, because the Yelp price data is coarse and available only for a selected sample of restaurants, and Yelp ratings often originate from local residents and are not necessarily comparable across areas. The lack of data on the exact restaurant choice of an individual further limits my ability to investigate the impact of price and quality variation across areas, by integrating it into the model. So while restaurants are one of the best example of horizontal differentiation, future datasets, perhaps drawn

[47] Berry and Waldfogel (2010) also do not find evidence that average restaurant quality increases with market size.

from constantly evolving online applications, will hopefully help illuminate the importance of vertical differentiation.

An alternative way of assessing the importance of price variation across density levels is to focus on a standardized restaurant item like McDonald's Big Mac. The Economist's Big Mac index is already a popular way of assessing price differences across countries. Landry (2013) collected data on Big Mac prices within New York City, and identified sizable spatial variation in prices. Perhaps surprisingly, the average price in locations within 0 to 4 miles from Manhattan's Penn Station is $3.95, and this average price actually increases to $4.02 in locations farther from the high density core of New York City (4 to 17 miles from Penn Station). So while it is suggestive that New York City serves some of the most expensive Big Macs in the U.S., these numbers are inconsistent with a significant drop in prices as density declines within a city, and do not undermine the validity of the paper's key finding of a large drop in the restaurant price index as one travels to the high density center of a city.

Assuming a constant meal price can also bias my estimates if the characteristics of restaurants close to home differ systematically from those of restaurants far from home. For instance, a majority of travelers live in suburbs, relatively far from downtown restaurants which tend to be slightly more expensive and upscale and to feature rarer restaurant types (e.g. French). In the nested-logit extension, I account for the exact location restaurants of each type, and for local tastes. Here, I use data from Yelp to compute average characteristics of restaurants that vary with travel time from home. Again, restaurants characteristics in general and quality ratings in particular are surprisingly constant across areas and travel time, but I do find that restaurants far from home are on average pricier, have more reviews, a more upscale ambience and attire, are more likely to require a reservation and less likely to have easy parking. The average price of a restaurant within 0–5 minutes of travel from home is 7.5% smaller than that for restaurants 40–45 minutes away. If my assumption that quality differentials exactly compensate price differentials is incorrect, then the positive correlation between trip time and restaurant price biases my estimates upward. So I re-estimate the logit model, but with meal price varying in each time bin. I normalize meal price of restaurant between 20–25 minutes from home at exactly 13 to be consistent with other estimations. I find $\hat{\sigma} = 8.4$ (column 4 of table 1). This elasticity of substitution is slightly lower than that from the model with constant meal prices, and the welfare gains from density are correspondingly slightly larger.

8.3 Redundant chain restaurants

In 2009, there were 267,499 chain restaurants in the US, representing 46% of the total number of commercial restaurants.[48] Restaurants within a given chain are never exactly similar, but clearly two McDonald's are highly substitutable with one another, and a model in which restaurants in the same chain are perfectly substitutable may be a better representation of reality. This is perhaps the simplest way to relax the logit assumption that all restaurants be equally substitutable (IIA property), and to introduce restaurant diversity into the model. Intuitively, areas consisting mostly of repeated chain restaurants have low diversity, and the model is now flexible enough to take this into account.

To estimate this model, I code the 50 largest restaurant chains in my data, which represent 23% of all restaurants in the sample, and are likely to occur more than once within 45 minutes of travel. Because travel is costly, a restaurant that is perfectly substitutable with another restaurant closer from home is never visited, so I eliminate all repeat chain restaurants from each traveler's choice set. Estimation is then exactly as in the logit model, with the estimator given by equation (7). I find $\hat{\sigma} = 8.4$ (column 5 of table 1). This extension generates predictions on the effect of global density on trip time that are only marginally closer to the data. The gains from density are slightly larger in the model with substitutable chains due to a lower σ than in the basic model. Note that repeat chain restaurants account for a smaller proportion of restaurants in denser areas, contrary to an intuition that would be correct if restaurants were randomly distributed.

8.4 Nested-logit model

In a nested-logit model, individuals first choose a category of restaurants (e.g. pizza, Chinese, burger or vegan) and then decide which restaurant to visit within that category. The IIA

[48] See https://www.npd.com/wps/portal/npd/us/news/press-releases/pr_110124/, retrieved 10 September 2010.

property of the logit model does not hold, because restaurants within the same category are more substitutable. I use the utility specification in Sheu (2011), which generates the same choice probabilities as a nested-CES model.

There are 85 categories of restaurants, indexed by c, and representing different types of cuisine.[49] I first assume that taste for categories is constant across categories and locations, and then I relax this assumption. As before, each restaurant receives a type I extreme value idiosyncratic shock ϵ_{kci} with scale parameter $1/(\sigma-1)$, but now each restaurant also receives a category-specific type I extreme value idiosyncratic shock ς_{kc} with scale parameter $1/(1-\mu)$. The utility from choosing restaurant i from category c in location k is:

$$u_{kci} = ln\left(r_{kci}\right) + \epsilon_{kci} + \varsigma_{kc}.$$

I provide the full derivation of the probability $prob(t_{kci}|\sigma,\mu,T_k,F_k)$ of traveling to each restaurant in Section F of the online appendix.

Let n index each trip t_{nk} in the sample. To estimate the model without data on the category of restaurant visited on each trip, I define $R_{nk}(t_{nk})$ as a set of restaurants at travel time 'close' to actual trip time t_{nk} in the choice set T_k. With a slight abuse of notation, let i index all restaurants in R_{nk}, so the log-likelihood function becomes:

$$\ell(\sigma,\mu,T_N,\mathbb{T}_K,\mathbb{F}_K) = \sum_{n=1}^{N} log\left(\sum_{i\in R_{nk}(t_{nk})} prob(t_{kci}|\sigma,\mu,T_k,F_k)\right),$$

and the maximum likelihood estimator is:

$$(\hat{\sigma},\hat{\mu}) = \underset{\sigma,\mu}{\text{argmax}}\ \ell(\sigma,\mu,T_N,\mathbb{T}_K,\mathbb{F}_K). \tag{10}$$

Defining R_{nk} as the set of all restaurants within 5 minutes of actual trip time t_{nk}, I find $\hat{\sigma}=8.4$ and $\hat{\mu}=9.9$, with μ imprecisely estimated. Given that $\sigma=\mu$ corresponds to the logit model, this result may suggest that the nests for types of cuisine are unnecessary.

These estimates, however, are problematic because they assume constant tastes for categories, so that pizza and vegan restaurants are equally desirable. This assumption overstates

[49] The categories, and their percentage share, are listed in Section F.4 of the online appendix. The category is 'undefined' for 17% of restaurants, usually smaller independent places serving standard fares. The next largest categories are Pizza (9.3%), Mexican (9.3%), American (9.1%), Burger (7.5%) and Chinese (6.3%). There is no category for casual dining 'family' restaurants, and such establishments are generally included in the 'Undefined' and 'American' category.

the attraction of small and arbitrary categories. Moreover, the restaurant cuisines available in a particular area probably reflect the tastes of individuals who live there. That is, individuals may live on average closer to their preferred restaurant type, and one should not infer a low preference for variety from these short trips. I therefore add to the model a location-specific distribution of tastes for restaurant categories, with a parameter b_{kc} capturing the taste for category c in location k.[50] The utility function becomes:

$$u_{kci} = ln\left(b_{kc}^{1/(\sigma-1)} r_{kci}\right) + \epsilon_{kci} + \varsigma_{kc}.$$

In Section F of the online appendix, I explain how introducing restaurant supply in the model allows to derive the taste parameters b_{kc} from observed restaurant density in each category.[51] Estimating the model with location specific tastes for restaurant categories, I find $\hat{\sigma}$ = 9.2 and $\hat{\mu}$ = 3.6, with $\hat{\mu}$ having a large standard error of 0.25 (column 7 of 1). That σ is larger than μ suggests that travelers have strong preferences for restaurant cuisine, but care less about the particular restaurant serving that cuisine. The nested-logit model predicts trip length in areas with high global density better than the logit model (30% closer to the data). If travelers care mostly about restaurant categories, of which there are relatively few, then in dense areas the large number of additional restaurants rapidly become redundant. Long trips in dense areas become unnecessary, because some of the thousands of restaurants far away from home belong to categories available closer to home.

I also compute a variety-adjusted restaurant price index from the nested-logit model with location-specific tastes for cuisines. Assuming that individuals always move to areas with exactly the same share of each cuisine as that in their original location, the 90^{th} to 10^{th} percentile differential in the index is about 42%, or four percentage point higher than that computed from the logit model. Hence, accounting for restaurant diversity and local tastes

[50]Schiff (2012) shows that even if preferences are identical everywhere, densely populated areas in a free-entry model feature more categories of restaurants, because they contain enough people with marginal tastes to make restaurants in the least popular categories profitable. This argument is intuitive, but it alone cannot account for the wide range of restaurant diversity that I measure in areas at the same density level. For instance, some areas in Texas contain a large share of Mexican restaurants, which likely reflects local tastes for this type of cuisine.

[51]Note that I obtain analytical results from a continuous version of the model with uniform density. The full derivation is available in an online appendix.

increases the gains from density. Clearly, data on restaurant choice - and not just on trip length - would help refine this finding.

9. Conclusion

This paper shows how to estimate the consumption value of density by combining travel data with microgeographic data on local businesses. Individuals' substitution patterns among travel destinations reveal gains from urban density that are large but localized. These gains originate in part from shorter trip times, but mostly arise because increased choice in denser areas allows individuals to visit destinations that they prefer. This result explains why empirical studies fail to uncover large reductions in travel as density increases, and shows that popular policies designed to reduce vehicle travel will instead have larger impacts on increasing gains from variety.

The consumption benefits of density that I estimate in the restaurant industry demonstrate that cities, and downtown cores in particular, enjoy a large advantage in non-tradable service provision. Ultimately, the canonical model of spatial equilibrium (Rosen, 1979, Roback, 1982) implies that positive amenities translate into higher land rents. Investigating the capitalization of greater access to goods and services into real estate prices, at fine spatial scales, could be a productive area for future research.

References

Ackerberg, Daniel A. and Marc Rysman. 2005. Unobserved product differentiation in discrete choice models: Estimating price elasticities and welfare effects. *RAND Journal of Economics* 36(4):771–778.

Albouy, David. 2008. Are big cities bad places to live? Estimating quality of life across metropolitan areas. Technical report, NBER, Working Paper 14472.

Albouy, David and Bert Lue. 2011. Driving to opportunity: Local wages, commuting, and sub-metropolitan quality of life. University of Michigan.

Anderson, Simon P., Andre de Palma, and Jacques-Francois Thisse. 1992. *Discrete Choice Theory of Product Differentiation*. Cambridge MA: MIT Press.

Ben-Akiva, Moshe and Steven Lerman. 1985. *Discrete Choice Analysis: Theory and Application to Travel Demand*. Cambridge MA: The MIT Press.

Berry, Steven and Joel Waldfogel. 2010. Product quality and market size. *Journal of Industrial Economics* 58(1):1–31.

Bhat, Chandra, Susan Handy, Kara Kockelman, Hani Mahmassani, Qinglin Chen, and Lisa Weston. 2000. Development of an urban accessibility index: Literature review. Technical Report Report TX-01/7-4938-1, Center for Transportation Research, The University of Texas at Austin, prepared for the Texas Department of Transportation.

Bockstael, Nancy E., Kenneth E. McConnell, and Ivar E. Strand Jr. 1989. A random utility model for sportfishing: Some preliminary results for florida. *Marine Resource Economics* 6(3):245–260.

Broda, Christian and David E. Weinstein. 2006. Globalization and the gains from variety. *Quarterly Journal of Economics* 121(2):541–585.

Broda, Christian and David E. Weinstein. 2010. Product creation and destruction: Evidence and price implications. *American Economic Review* 100(3):691–723.

Carlino, Gerald A. and Albert Saiz. 2008. Beautiful city: Leisure amenities and urban growth. Technical report, FRB of Philadelphia Working Paper No. 0822.

Clark, Terry Nichols. 2003. Lakes, opera and juice bars: Do they drive development? In Terry Nichols Clark (ed.) *The City as an Entertainment Machine*, volume 9 of *Research in Urban Policy*. New York NY: JAI Press/Elsevier, 103–140.

Clawson, Marion. 1959. Methods of measuring the demand for and value of outdoor recreation. Technical Report Reprint No. 10, Washington D.C.: Resources of the Future, Inc.

Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon. 2011. The identification of agglomeration economies. *Journal of Economic Geography* 11(2):253–266.

Congress of the New Urbanism. 2013. *Charter of the New Urbanism*. McGraw-Hill.

Couture, Victor, Gilles Duranton, and Matthew A. Turner. 2012. Speed. University of Toronto.

Davis, Peter. 2006. Spatial competition in retail markets: Movie theaters. *The RAND Journal of Economics* 37(4):964–982.

Diamond, Rebecca. 2013. The determinants and welfare implications of US workers' diverging location choices by skill: 1980–2000. Stanford University.

Dubin, Jeffrey A. and Daniel L. McFadden. 1984. An econometric analysis of residential electric appliance holdings and consumption. *Econometrica* 52(2):345–362.

Ewing, Reid and Robert Cervero. 2010. Travel and the built environment. *Journal of The American Planning Association* 76(3):265–294.

Feenstra, Robert. 1994. New product varieties and the measurement of international prices. *American Economic Review* 84(1):157–177.

Glaeser, Edward L. and Matthew E. Kahn. 2004. Sprawl and urban growth. In Vernon Henderson and Jacques-François Thisse (eds.) *Handbook of Regional and Urban Economics*, volume 4. Amsterdam: North-Holland, 2481–2527.

Glaeser, Edward L. and Janet E. Kohlhas. 2004. Cities, regions and the decline of transport costs. *Papers in Regional Science* 83(1):197–228.

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

Goldberg, Pinelopi Koujianou. 1998. The effects of the corporate average fuel efficiency standards in the us. *The Journal of Industrial Economics* 46(1):1–33.

Handbury, Jessie. 2012. Are poor cities cheap for everyone? Non-homotheticity and the cost of living across US cities. University of Pennsylvania.

Handbury, Jessie and David E. Weinstein. 2012. Is new economic geography right? Evidence from price data. Technical report, NBER Working Paper No. 17067.

Helpman, Elhanan. 1998. *The size of regions*, chapter Topics in Public Economics: Theoretical and Applied Analysis. Cambridge: Cambridge University Press, 33–54.

Hotelling, Harold. 1947. Letter of June 18, 1947, to Newton B. Drury. Included in the report 'The Economics of Public Recreation: An Economic Study of the Monetary Evaluation of Recreation in the National Parks' by Prewitt, 1949. Mimeographed. Washington D.C.: Land and Recreational Planning Division, National Park Service.

Houde, Jean-François. 2012. Spatial differentiation and vertical mergers in retail markets for gasoline. *American Economic Review* 102(5):2147–82.

Ihrke, David K., Carol S. Faber, and William K. Koerber. 2011. Geographical mobility: 2008 to 2009. Technical report, U.S. Department of Commerce, Economics and Statistics Administration, U.S. Census Bureau.

Krugman, Paul R. 1991. Increasing returns and economic geography. *Journal of Political Economy* 99(3):483–499.

Landry, Anthony. 2013. Borders and big macs. *Economics Letters* 120:318–322.

Lee, Sanghoon. 2010. Ability sorting and consumer city. *Journal of Urban Economics* 68(1):20–33.

Li, Nicholas. 2012. An Engel curve of variety. University of Toronto.

Melo, Patricia C., Daniel J. Graham, and Robert B. Noland. 2009. A meta-analysis of estimates of urban agglomeration economies. *Regional Science And Urban Economics* 39(3):332–342.

Moretti, Enrico. 2012. *The New Geography of Jobs*. New York: Houghton Mifflin Harcourt.

Murphy, Daniel. 2013. Urban density and the substitution of market purchases for home production. University of Michigan.

Niemeier, Debbie A. 1997. Accessibility: an evaluation using consumer welfare. *Transportation* 24(4):377–396.

Okrent, Abigail M. and Julian M. Alston. 2010. Demand for food in the United States. Technical report, RMI-CWE Working Paper number 1002.

Popkin, Susan J., Bruce Katz, Mary K. Cunningham, Karen D. Brown, Jeremy Gustafson, and Margery A. Turner. 2004. A decade of HOPE VI: Research findings and policy challenges. Technical report, Urban Institute, Washington, DC.

Roback, Jennifer. 1982. Wages, rents and the quality of life. *Journal of Political Economy* 6:1257–1278.

Rosen, Sherwin. 1979. Wages-based indexes of urban quality of life. In P. Mieszkowski & M. Straszheim (ed.) *Current Issues in Urban Economics*. Baltimore: John Hopkins Univ. Press.

Sato, Kazuo. 1976. The ideal log-change index number. *Review of Economics and Statistics* 58(2):223–228.

Schiff, Nathan. 2012. Cities and product variety. University of British Columbia.

Sheu, Gloria. 2011. Price, quality, and variety: measuring the gains from trade in differentiated products. US Department of Justice.

Shoup, Donald. 2005. *The high cost of free parking*. Chicago: Planners Press.

Train, Kenneth. 2009. *Discrete Choice Methods with Simulation*. Second edition. Cambridge (MA): Cambridge University Press.

Vartia, Yrjo. 1976. Ideal log-change index numbers. *Scandinavian Journal of Statistics* 3(3):121–126.

Waldfogel, Joel. 2008. The median voter and the median consumer: Local private goods and population composition. *Journal of Urban Economics* 63(2):567–582.

Wang, Hui. 2010. Consumer valuation of retail networks: Evidence from the banking industry. Peking University.

West, Brian H., Ralph N. McGill, and Scott Sluder. 1999. Development and validation of light-duty vehicle modal emissions and fuel consumption values for traffic models. Technical report, Federal Highway Administration.

Appendix A. Derivation of aggregate relative price index

The linear utility specification in a nested-logit model with one nest for restaurants and one nest for all other goods is similar to that in the nested-logit model of Section 8, and generates choice probabilities which can easily be interpreted in terms of expenditure shares. An individual first solves the maximization problem within the restaurant nest (exactly as in Section 3) and within the nest for all other goods. Then she solves the aggregate utility maximization problem by choosing expenditure shares on restaurants and on all other goods. Denote the quantity of the restaurant good purchased in area k by r_k, the quantity of all other goods purchased by g_k, the price index for restaurants (from equation 5) by R_k and the price index for all other goods by G_k. Solving the aggregate utility maximization problem leads to $r_k/g_k = (R_k/G_k)^{-\nu}$, where ν is the elasticity of substitution between restaurants and all other goods. Denote the price elasticity of demand for restaurant by ϵ_r, such that $(\partial r/\partial R)(R/r) = \epsilon_r$. It is straightforward to show that $\epsilon_r = -\nu$, so that $\epsilon_r = -1$ corresponds to the limiting case $\nu = 1$.[52] The aggregate relative price index between area k and k' is:

$$P_{k,k'} = \frac{(R_{k'}^{1-\nu} + G_{k'}^{1-\nu})^{1/(1-\nu)}}{(R_k^{1-\nu} + G_k^{1-\nu})^{1/(1-\nu)}}. \tag{A1}$$

As shown in Sato (1976) and Vartia (1976), one can express the relative price index above in terms of expenditure shares. For instance, if $s_{Rk'}$ is the expenditure share on restaurants in area k', then:

$$P_{k,k'} = \left(\frac{G_{k'}}{G_k}\right)^{w_{Gk'}} \left(\frac{R_{k'}}{R_k}\right)^{w_{Rk'}},$$

where:

$$w_{Rk'} = \frac{(s_{Rk'} - s_{Rk})/(ln(s_{Rk'}) - ln(s_{Rk}))}{(s_{Rk'} - s_{Rk})/(ln(s_{Rk'}) - ln(s_{Rk})) + (s_{Gk'} - s_{Gk})/(ln(s_{Gk'}) - ln(s_{Gk}))}.$$

I assume that the price index for all other goods is constant across areas, so that $G_k = G_{k'}$. As $\epsilon_r = -1$, the expenditure share on restaurants is constant, so that $s_{Rk} = s_{Rk'}$ for any areas k and k'. If $s_{Rk'}$ is arbitrarily close to s_{Rk}, then we can find $w_{Rk'}$ as: $\lim_{s_{Rk} \to s_{Rk'}} w_{Rk'} = s_{Rk'}$.

[52]It is standard to assume that $\nu > 1$, but the welfare estimates are not sensitive to using $\epsilon_r = -1.02$, which is the price elasticity suggested by Okrent and Alston (2010) in their meta-analysis, instead of $\epsilon_r = -1$.

With data on expenditure shares, it is now possible to compute the aggregate relative price index and to measure the average household's willingness to pay to enjoy a 20% decrease in the restaurant price index. The 2009 CEX suggests that food away from home, that I take as a proxy for restaurants, accounts for 5.3%, of total expenditures, so I set $s_{Rk'}$ = 0.053. The aggregate relative price index becomes $P_{k,k'} = \left(\frac{G_{k'}}{G_k}\right)^{w_{Gk'}} \left(\frac{R_{k'}}{R_k}\right)^{w_{Rk'}} = (1)\left(\frac{R_{k'}}{R_k}\right)^{s_{Rk'}} = 0.8^{0.053}$ = 0.98824. The average total household expenditures in the CEX 2009 is about $49,067, so the average household's willingness to pay for a 20% decrease the restaurant price index is 49,067$(1 - 0.98824)$, which equals $576.[53]

Appendix B. Data

A. Travel speed

To obtain estimates of car travel speed for a trip of a given distance in a given location, I regress the log of trip speed on the log of trip distance for the entire NHTS sample of car trips, with a fixed effect at the census tract level to capture speed variation across areas. Measuring speed as a function of trip distance allows longer trips to be faster, because they are taken mostly on faster roads like highways. In an alternative specification, I add individual characteristics to the regression.[54] Let n index each trip and its driver, and k index the census tract in which an individual lives. Let $speed_{nk}$ and $distance_{nk}$ denote the speed and distance of trip n in census tract k. The estimating equation is:

$$ln(speed_{nk}) = \alpha + \beta_1 ln(distance_{nk}) + \beta_2 X_n + \theta_k + \epsilon_{nk}, \qquad \text{(B1)}$$

where X_n is a vector of individual characteristics. I find that longer trips are on average faster, with an elasticity of 0.42, precisely estimated. Richer, younger and more educated individuals drive faster.

[53]To compute welfare gains when the elasticity of demand is higher, say at $\epsilon_R = -2$, I use the fact that $\frac{s_{Rk}}{s_{Rk'}} = \frac{R_k}{R_{k'}}^{1+\epsilon} P_{k,k'} \approx \frac{R_k}{R_{k'}}^{1+\epsilon}$. I find that for a household with average expenditures, a 20% decrease in the restaurant price index increases the expenditure share on restaurants from $s_{Rk} = 0.053$ to $s_{Rk'} = 0.06625$ and generates welfare gains valued at $617, a 7% increase relative to the case with $\epsilon = -1$.

[54]See Couture, Duranton, and Turner (2012) for additional details on these regressions. Note that for census tracts in which I have less than 10 trips (6.6% of the sample), I aggregate census tract fixed-effects at the county level before computing price indices.

B. *Fuel cost*

The NHTS contains data on daily gasoline prices in the region each trip originates from, as well as information on vehicle type in four categories (car, van, SUV, pickup truck and other trucks). To each type of vehicle, I assign the fuel efficiency of the best selling vehicle of this type in 2000.[55] For instance, the best selling pickup in 2000 was the F-150. The other best selling vehicles in each category are the Toyota Camry (car), Ford Explorer (SUV) and Dodge Caravan (van). The relationship between fuel consumption and trip distance is nonlinear, because vehicles are less efficient at very low or very high speeds. To adjust for this, I draw from results by West, McGill, and Sluder (1999), who analyze fuel consumption at different speeds for 9 models of cars and light trucks, built in the year 1988 to 1997. Based on their numbers, I assume that each vehicle reaches maximum fuel economy at speeds between 25 and 60 miles per hour, and consumes 20% more at speeds exceeding 60 mph or below 25 mph. I take maximum fuel economy to be the value reported for 'highway' consumption by the Department of Energy, which is 19 miles per gallon (mpg) for the Ford F-150, 28 mpg for the Toyota Camry, 19 mpg for the Ford Explorer and 24 mpg for the Dodge Caravan.[56] To compute speed at any point on a trip, I use fitted values from equation (B1), which provides average speed as a function of distance. As an example, if an individual takes a 10 miles trip in a pick-up, and reaches 25 mph after 5 miles but never reaches a speed beyond 60 mph, then his total fuel consumption in gallons is: $5 \text{ miles}/(19 * (1 - 0.2) \text{ mpg}) + 5 \text{ miles}/19 \text{ mpg}$. Then fuel cost for a trip is the product of total fuel consumption and fuel price. Fuel price is a variable in the NHTS, which is equal to the retail price of gasoline as provided by the US Energy Information Administration on the Monday closest to an individual's travel day. The prices vary across 5 Petroleum Administration for Defense Districts: East Coast, Gulf Coast, Midwest, Rockies and West Coast (plus Alaska and Hawai).

[55] 2009 models are not representative of the actual stock of vehicles on the road in 2009, which is on average 10 years old. The data on the best selling vehicles of 2000 is from Edmunds.com, an online source of information on the American vehicle market. Available online at http://www.edmunds.com/car-reviews/top-10/top-10-best-selling-vehicles-in-2000.html, retrieved 10 September 2013.

[56] The data is available online. For instance information on the F-150 is at http://www.fueleconomy.gov/feg/bymodel/2000_Ford_F150_Pickup.shtml, retrieved 10 September 2013.