How Segregated Is Urban Consumption?

Donald R. Davis

Columbia University and National Bureau of Economic Research

Jonathan I. Dingel

University of Chicago and National Bureau of Economic Research

Joan Monras

Universitat Pompeu Fabra, Barcelona Graduate School of Economics, and Centre for Economic Policy Research

Eduardo Morales

Princeton University and National Bureau of Economic Research

We provide measures of ethnic and racial segregation in urban consumption. Using Yelp reviews, we estimate how spatial and social frictions influence restaurant visits within New York City. Transit time plays a first-order role in consumption choices, so consumption segregation partly reflects residential segregation. Social frictions also affect restaurant choices: individuals are less likely to visit venues in neighborhoods demographically different from their own. While spatial and social frictions jointly produce significant levels of consumption segregation, we find that restaurant consumption is only about half as segregated as residences. Consumption segregation owes more to social than spatial frictions.

We thank Jesse Shapiro, four anonymous referees, Treb Allen, David Atkin, Pierre-Philippe Combes, Victor Couture, Thomas Covert, Alon Eizenberg, Ingrid Gould Ellen, Mogens Fosgerau, Manuel Garcia-Santana, Marçal Garolera, Robin Gomila, Joshua Gottlieb, Jessie Handbury, Art O'Sullivan, Albert Saiz, and many seminar audiences for helpful comments. We thank Bowen Bao, Luis Costa, Amrit K. Daniel, David Henriquez, Yan Hu,

Electronically published June 11, 2019
[*Journal of Political Economy*, 2019, vol. 127, no. 4]

I. Introduction

For half a century, the United States has prohibited ethnic and racial discrimination in housing, jobs, and education. Even so, segregation in each of these domains remains a stubborn feature of modern America (Hellerstein and Neumark 2008; Boustan 2011). Many studies have documented these facts and examined their consequences for socioeconomic outcomes (Massey and Denton 1993; Cutler and Glaeser 1997; Chetty et al. 2014).

Discrimination in consumption venues has also been prohibited for decades, yet racial and ethnic segregation in this domain has been studied much less. A major achievement of the civil rights movement was the Civil Rights Act of 1964, prohibiting discrimination based on race or ethnicity in public accommodations.[1] Jim Crow laws had segregated places where people meet socially in order to maintain segregation of intimate contact (Myrdal 1944, 588). In contemporary America, these shared spaces have the potential to form what Anderson (2011) calls "a cosmopolitan canopy," a place where a diversity of people may interact such that "a cognitive and cultural basis for trust is established that often leads to the emergence of more civil behavior" (xv).[2] Such social capital has potential consequences for many economic outcomes (Guiso, Sapienza, and Zingales 2009; Smith 2010). Gone are the days of whites-only lunch counters. Yet we do not know the degree to which consumption venues are integrated and serve as places where people of different backgrounds encounter each other in everyday life (O'Flaherty 2015, 236–37).

We cannot say a priori whether segregation along demographic lines is greater in consumption or residences. If spatial frictions (i.e., the costs of traversing the city) were arbitrarily high, then consumers would largely patronize the businesses closest to their residences, making residential and consumption segregation very similar. The fact that individuals may

Charlene Lee, Rachel Piontek, Anil Sindhwani, Ludwig Suarez, Shirley Yarin, and, especially, Kevin Dano, Ben Eckersley, Hadi Elzayn, and Benjamin Lee for research assistance. Thanks to the New York Police Department, and especially Gabriel Paez, for sharing geocoded crime data. Dingel thanks the Kathryn and Grant Swick Faculty Research Fund at the University of Chicago Booth School of Business for supporting this work. This work was completed in part with resources provided by the University of Chicago Research Computing Center. Monras thanks the Banque de France Sciences Po partnership. Part of this work is supported by a public grant overseen by the French National Research Agency (ANR) as part of the Investissements d'Avenir program LIEPP (ANR-11-LABX-0091, ANR-11- IDEX-0005-02). Morales thanks the University of Wisconsin–Madison and the Cowles Foundation at Yale University for their hospitality and support. Code and data are provided as supplementary material online.

[1] Examples of public accommodations include hotels, restaurants, and entertainment venues.

[2] There is a large social psychology literature on the mechanisms that may reduce intergroup prejudice. Specifically, the literature on the "contact hypothesis" focuses on the potential for intergroup contact to reduce such prejudice (Allport 1954). Study results in this literature have generally been consistent with the hypothesis that appropriate intergroup contact reduces intergroup prejudice (Pettigrew and Tropp 2006), including some laboratory and field experiments (Paluck and Green 2009).

move around the city for consumption purposes could mitigate the effect of the spatial separation of residences if diverse consumers choose common destinations. Conversely, mobility may allow them to segregate even more if their choices diverge. Divergent choices could result from social frictions, such as aversion to consuming in areas with different population demographics or racial segregation of social networks that influence choices (e.g., family or friends). Consumers from different backgrounds may also choose different venues because of differences in tastes. The interactions of spatial frictions, social frictions, and heterogeneous tastes make the degree of consumption segregation an empirical question.

In this paper, we estimate a measure of consumption segregation along racial and ethnic lines for the residents of New York City and quantify the contributions of spatial and social frictions to it. We do so by estimating a discrete-choice model of restaurant visits using information on more than 18,000 consumption decisions by individuals living in New York City who review on Yelp.com, a website on which users review local businesses. We use our estimated parameters to predict the consumption decisions of all New York City residents. We find that consumption choices are much less segregated than residential locations. Dissimilarity indices contrasting the consumption destinations of different demographic groups are about half the value of dissimilarity indices for residential locations. Both spatial and social frictions have quantitatively large influences on the geography of consumption, with social frictions contributing relatively more to consumption segregation along ethnic and racial lines.

Inferring spatial and social frictions from consumption behavior requires controlling for other determinants of consumers' choices. We exploit several advantages of our data set, described in Section II, to identify these frictions. First, we locate Yelp reviewers' residences and workplaces, allowing us to measure spatial frictions that account for the fact that consumption may originate at home, at work, or on the commute between them and that both automobile and public transit may be used. Second, we combine data from Yelp and the US Census Bureau to characterize reviewer demographics, restaurant characteristics, and neighborhood demographics. This allows us to distinguish demographic differences in tastes from social frictions and to measure the contributions of both individual-level homophily and demographic differences across neighborhoods to social frictions.

However, our data set is not without limitations. First, Yelp reviewers are not representative of the general population. In terms of observable characteristics, reviewers in our estimation sample are more likely to be Asian and female. They live in neighborhoods with higher incomes and more residents between the ages of 21 and 39. We allow our estimates of preference parameters to vary with these observable individual characteristics but cannot do the same for unobservable characteristics. Our results for the whole population of New York City residents thus necessarily predict the level of consumption segregation that prevails if everyone behaves as the observationally equivalent reviewers in our estimation sample. Second, we have limited ability to distinguish between Hispanic and white reviewers in our sample, so we cannot separately identify these two groups' preference parameters. However, we exploit tract-level information on demographics to capture social frictions between Hispanics and whites. Third, we observe every review written by Yelp users, not every restaurant visit. Identifying consumer preferences therefore involves assumptions on review-writing behavior, which we discuss in Section III.

We model consumers' behavior using a conditional-logit specification in which a consumer's valuation of a restaurant may depend on spatial frictions, social frictions, and a large set of observable characteristics of the consumer and the restaurant. All preference parameters are allowed to vary flexibly by race. Our estimation procedure makes use of the McFadden (1978) choice set construction technique to address the computational burden arising from consumers choosing among the thousands of restaurants in New York City.

We present our parameter estimates in Section IV. Our quantification of consumers' aversion to incurring longer travel times reveals a first-order role for spatial frictions in determining the geography of consumption. Depending on the origin of the trip and the mode of transport used, halving the minutes of travel time to a venue implies that a consumer would be two to nearly four times more likely to visit the venue from that origin by that mode. These spatial frictions will cause consumption patterns to partly inherit residential patterns of segregation.

Consumption segregation also reflects the influence of social frictions. These frictions make consumers' decisions depend on the contrast between the residential demographics of the restaurant's location and either the residential demographics of the consumer's home location or the consumer's own racial or ethnic identity. All else equal, a consumer is more likely to visit a venue in a census tract that is more demographically similar to her home tract. Individuals are also more likely to visit restaurants in tracts with a larger share of residents of their own racial group. While consumption may be integrated de jure, these social frictions make consumption less integrated de facto.

Importantly, our estimates of both spatial and social frictions are obtained after controlling for race-specific tastes for observable features of restaurants and areas of the city. For example, we incorporate cuisine category fixed effects. Thus, for example, our finding that Asian consumers are more likely to visit restaurants (of any type) located in neighborhoods with more Asian residents is conditional on the fact that Asian consumers are more likely to visit restaurants serving Asian cuisines. Similarly, we allow consumers' valuations of restaurant prices and ratings to depend on the income level of their home census tract and control for income differences when estimating social frictions associated with racial demographic differences. In robustness checks, we introduce restaurant fixed effects that vary by race and allow for correlation in consumer-specific preferences across restaurants of similar characteristics in nested-logit specifications. These specifications yield similar estimates of spatial and social frictions.

Our estimated model fits the data well. Race-specific preference parameters are key to capturing the level of consumption segregation that we observe in the estimation sample: a specification that assumes that preference parameters are common across all consumers cannot replicate the in-sample isolation of consumers of different races. The specifications that introduce race-specific restaurant fixed effects yield only very modest improvements in fit. This is consistent with the fact that there is little segregation of Yelp reviewers between pairs of restaurants that are observationally equivalent.

Using our estimated model of the restaurant visit decision, we compute measures of consumption segregation for the entire residential population of New York City in Section V. Specifically, we characterize the ethnic and racial segregation of the predicted consumption choices using dissimilarity indices. A dissimilarity index describes the fraction of the population belonging to a group—Asian consumers, for example—that would have to alter their consumption choices in order to match the distribution of predicted restaurant choices made by the remainder of the population. Despite the magnitude of the estimated spatial and social frictions, consumption dissimilarity is notably lower than residential dissimilarity for all ethnic and racial groups.

To quantify the contribution of spatial frictions, social frictions, and demographic differences in tastes to consumption segregation, we recompute the dissimilarity indices using the consumption decisions predicted by our estimated model when the coefficients capturing the corresponding friction are set to zero. Social frictions make a larger contribution to consumption segregation than spatial frictions. Eliminating spatial frictions entirely would reduce consumption dissimilarity indices by 8–20 percent, whereas eliminating social frictions would reduce dissimilarity by 22–41 percent.

We also use our estimated model to examine the impact on consumption segregation of counterfactual changes in transportation policy and in the preference parameters determining the degree of social frictions. Consistent with the modest role of spatial frictions overall, major changes in transportation infrastructure have only small effects on consumption dissimilarity. Reductions in social frictions would integrate consumption.

Finally, we use our estimates to measure the welfare consequences of neighborhood change for incumbent residents in Section VI. Gentrification is associated with changes in both restaurants' and residents' characteristics that affect the value of consumption. We compute the change in welfare that the residents of a census tract in the middle of Harlem would experience if the surrounding census tracts were to exhibit the residential and restaurant characteristics of high-income, majority-white census tracts of the Upper East Side. We find a significant reduction in the value of their consumption opportunities. This is attributable to the increase in social frictions associated with the change in racial demographics. The change in restaurants' characteristics would have very modest effects on their welfare.

Our findings relate to a recent literature on the geography of urban consumption. Studies have documented cross-city variation in the tradable goods available for consumption (Handbury 2013; Handbury and Weinstein 2015), and geographic variation in the supply of nontradables has been posited to shape the relative attractiveness of cities (Glaeser, Kolko, and Saiz 2001; Schiff 2015). Waldfogel (2007) documents that restaurant entry in different cuisine categories is correlated with local demographic composition. This dimension of economic life has grown increasingly important in recent decades.[3] Prior studies of the geography of consumption within the city include Katz (2007), Houde (2012), Couture (2015), and Eizenberg, Lach, and Yiftach (2017). Relative to this prior work, we build a unique data set that combines information on individuals' home and work locations, their demographics, and characteristics of the restaurants they patronize, and we use it to separately identify the effect of spatial and social frictions on consumer decisions.

We study urban consumption using online user-generated content, which is increasingly exploited by social scientists. Among others, Anderson and Magruder (2012) and Luca (2016a) examine Yelp's effects on restaurant outcomes. Harrison et al. (2014) use information disclosed in reviews to detect outbreaks of food poisoning unreported to New York City health authorities. Edelman and Luca (2014) infer racial identities from profile photos to study discrimination on Airbnb.com. Caetano and Maheshri (2019) document consumption segregation by gender using Foursquare data. Luca (2016b) surveys this growing body of research on user-generated content and social media.

We contribute to the large literature on social and economic fragmentation related to demographic differences by measuring consumption segregation in restaurants. The prior literature has largely focused on residential segregation, though there are studies documenting the segregation of workplaces (Hellerstein and Neumark 2008), students' friendship networks

[3] US households' share of food spending devoted to food prepared away from home grew from less than 26 percent in 1970 to more than 43 percent in 2012 (US Department of Agriculture 2014). Analogously, while the number of daily commuting trips has stayed relatively constant for decades, trips for social/recreational purposes have steadily grown (Pisarski 2006).

(Echenique and Fryer 2007), and media consumption (George and Waldfogel 2003; Oberholzer-Gee and Waldfogel 2009). We study racial segregation of consumption in a setting in which consumers travel to consume and come face-to-face with each other. Everyday encounters between people of different backgrounds in shared public spaces may be a basis for building understanding and tolerance (Anderson 2011), though consumption venues are also sometimes sites of racial and ethnic discrimination (Labaton 1994; Lee 2000; Ayres 2001; Antecol and Cobb-Clark 2008; Schreer, Smith, and Thomas 2009). We provide the first quantification of the segregation of these consumption choices.

II. Data

We combine data from Yelp and other sources to estimate our model of the restaurant visit decision and compute measures of consumption segregation. Section II.A describes the Yelp data and Section II.B describes the other sources of data we use. Section II.C presents evidence suggestive of the influence of spatial and social frictions on consumers' restaurant choices.

A. Yelp Data

Yelp.com is a website on which users review local businesses, primarily restaurants and retail stores (Yelp 2013). It describes a venue in terms of its address, average rating, user reviews, and a wide variety of other characteristics. Yelp's coverage of restaurants is close to comprehensive (see app. B.1). In addition to assigning a rating of one to five stars, reviewers describe their personal experience with the business. Crucial for our purposes is that users sometimes disclose information in their reviews about their residential and work locations.

We use data on Yelp users who reviewed a New York City (henceforth, NYC) restaurant venue between 2005 and 2011. As described in detail in appendix B.2, we identify reviewers' residential and work locations from the text of their reviews. We locate reviewers using reviews of all venues, not only restaurants. Specifically, we first search the text of a large number of reviews for 26 key phrases related to location, such as "close to me," "block away," and "my apartment." Then we read the reviews containing these phrases to infer whether the reviewer's home or work is proximate to the reviewed business. Finally, we estimate the residential and work locations of a reviewer as the average of the latitude-longitude coordinates of the sets of venues identified as being close to this reviewer's home and work locations, respectively. Restricting our sample to users whose reviews do not reveal a change in residence or workplace within NYC and whose home and work locations are in census tracts with demographic and in-

TABLE 1
Estimation Sample and NYC Summary Statistics

	Estimation Sample Yelp Rviewers	Manhattan Tracts	NYC Tracts
Reviewer appearance/tract demographics:			
Female	.609	.531	.526
Male	.343	.469	.474
Asian	.243	.112	.126
Black	.075	.129	.227
White or Hispanic	.418	.740	.620
Hispanic		.254	.286
White		.481	.334
Reviewer race indeterminate	.264		
Home tract characteristics:			
Median household income (thousands)	75.6	73.9	55.1
Age 21–39 residents share	.423	.368	.306
Asian isolation index	.197	.273	.326
Black isolation index	.280	.382	.569
White/Hispanic isolation index	.778	.787	.731
Observations	440	279	2,110

Note.—This table summarizes characteristics of the 440 Yelp reviewers in our estimation sample and all census tracts in Manhattan and New York City. Reviewer demographics are inferred from Yelp profile photos. Tract demographics are from the 2010 Census of Population and tract incomes from the 2007–11 American Community Survey. Tracts are weighted by residential population. Isolation indices are as defined in Massey and Denton (1988).

come information, we obtain an estimation sample of 18,015 reviews written by 440 distinct reviewers.[4]

Yelp reviewers typically post a profile photo, which we use to infer their apparent gender and race.[5] Mayer and Puller (2008) compare measures of ethnicity and race inferred from photos with administrative data and find a high degree of accuracy in partitioning subjects into three racial groups: Asian, black, and white or Hispanic. Consequently, when inferring each reviewer's race from her profile picture, we limit ourselves to classifying individuals into these three groups.

Table 1 reports summary statistics for the 440 reviewers included in our estimation sample and the broader NYC residential population. Sixty-one percent of estimation sample reviewers are female, and only 5 per-

[4] We could use a considerably larger sample if we required information only on the reviewer's home location. Table A11 shows that estimates of models that do not require information on individuals' workplaces are similar whether we use the estimation sample or the sample of individuals for whom we have information only on home locations. However, this specification underestimates social frictions compared to our specifications that incorporate workplace information.

[5] While users may choose "male" or "female" for their gender on their Yelp profile, this information is not publicly displayed. Thus, we classify reviewers on the basis of their gender presentation in their profile photo. Reviewers with profile photos for which we could not classify the gender (e.g., cartoon graphics, photos of animals) have both male and female dummy variables equal to zero.

cent of reviewers are of unidentified gender. We could not infer race for 26 percent of reviewers. Asian reviewers constitute 24 percent of the estimation sample, while white or Hispanic reviewers are 42 percent of the sample. Asians are thus overrepresented in our sample, as Asian residents constitute only about 11 percent of the population of NYC. Although only 10 percent of the reviewers with an inferred race were identified as black, these individuals wrote more than 1,000 reviews.

Figure 1 depicts the home and work locations of the reviewers in our estimation sample. Consistent with patterns in the broader population of NYC, these reviewers' workplaces are concentrated in Manhattan below Fifty-Ninth Street, while their residences are more dispersed. The average reviewer in our estimation sample lives in a census tract with median household income near $75,600, which is typical of Manhattan but higher than NYC as a whole ($55,100). The reviewers in our estimation sample tend to live in census tracts with a share of the population between the ages of 21 and 39 (42 percent) that is higher than that of both Manhattan (37 percent) and NYC as a whole (31 percent). These patterns are consistent with statements that Yelp's global user base is younger and higher-income than the population as a whole (Yelp 2013).

Asian and black reviewers in our estimation sample are less residentially segregated than Asian and black residents of NYC as a whole. Table 1 reports "isolation indices" as defined in Massey and Denton (1988), $\Sigma_k(\text{pop}_{gk}/\text{pop}_g)\cdot(\text{pop}_{gk}/\text{pop}_k)$, where pop_g is the population of group g, pop_k is the population of tract k, and pop_{gk} is the population of group

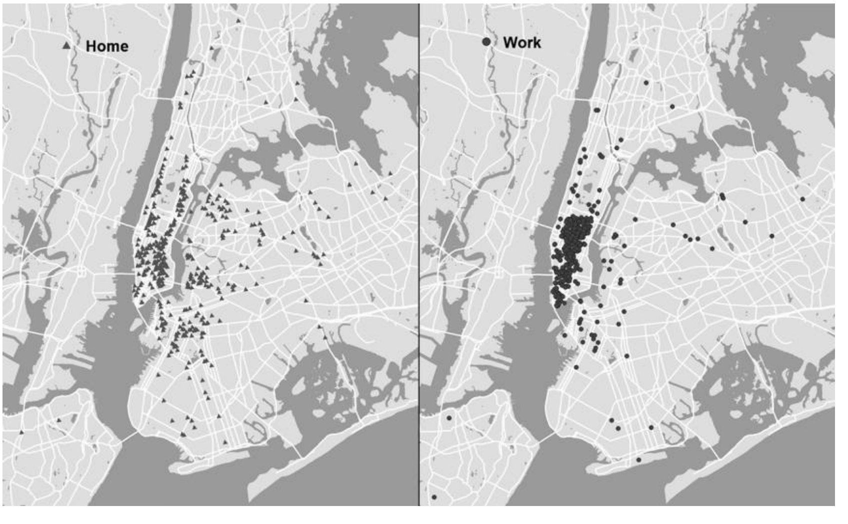

FIG. 1.—Locations of Yelp reviewers in the estimation sample. This figure depicts the distribution of home and work locations of the 440 reviewers in our estimation sample.

g in tract k. These characterize the average group g (e.g., Asian) share of tract residents experienced by members of group g. White/Hispanic reviewers in our estimation sample live in census tracts with a white/Hispanic share of residents that is typical of that experienced by white/Hispanic residents of NYC. By contrast, Asian and black reviewers in our estimation sample live in census tracts that have lower Asian and black shares, respectively, than is typical. The isolation indices for Asian and black reviewers in our estimation sample are about three-quarters of their values for Manhattan residents and half their values for NYC residents as a whole.

Figure A1 displays all the restaurants reviewed by users in our estimation sample. These venues are concentrated in Manhattan below Fifty-Ninth Street, but our estimation sample contains venues in many parts of NYC. Table A1 summarizes the distribution of reviews across NYC restaurants in terms of venues' prices, ratings, cuisine types, and boroughs for both our estimation sample and all Yelp reviewers. The reviewers in our estimation sample exhibit review patterns similar to those of the broader Yelp population reviewing NYC restaurants.

B. *NYC Transit, Demographic, and Crime Data*

To measure spatial frictions, we use car and public transit times between the centroids of census tracts from Google Maps. In addition to direct travel from home or work, we compute the additional transit time an individual would incur by incorporating a visit to a venue as part of her commute. Denote the transit time from location x to location y by $\mathrm{time}(x, y)$. For reviewer i living in h_i and working in w_i, the travel time associated with visiting venue j in tract k_j from her commuting path p_i is computed as

$$\mathrm{time}(p_i, j) = \frac{1}{2}\max\left\{\mathrm{time}\left(h_i, k_j\right) + \mathrm{time}\left(w_i, k_j\right) - \mathrm{time}(h_i, w_i), 0\right\},$$

where the maximum operator imposes the triangle inequality on transit times.

To measure social frictions associated with racial and ethnic demographics, we use data from the 2010 Census of Population that describe each census tract's residential population in terms of five groups: Asian, black, Hispanic, white, and other.[6] These population counts are depicted in figure 2.[7] Using these data, we measure ethnic and racial differences

[6] To be precise, we use the population counts of non-Hispanic Asians, non-Hispanic blacks, all Hispanics, and non-Hispanic whites to respectively define the groups we call "Asian," "black," "Hispanic," and "white." The "other" group includes Native Americans, Hawaiians, other races, and mixed-race categories; it constitutes about 3 percent of the NYC population.

[7] This map was inspired by a *New York Times* 2010 project, "Mapping America: Every City, Every Block (http://projects.nytimes.com/census/2010/explorer).

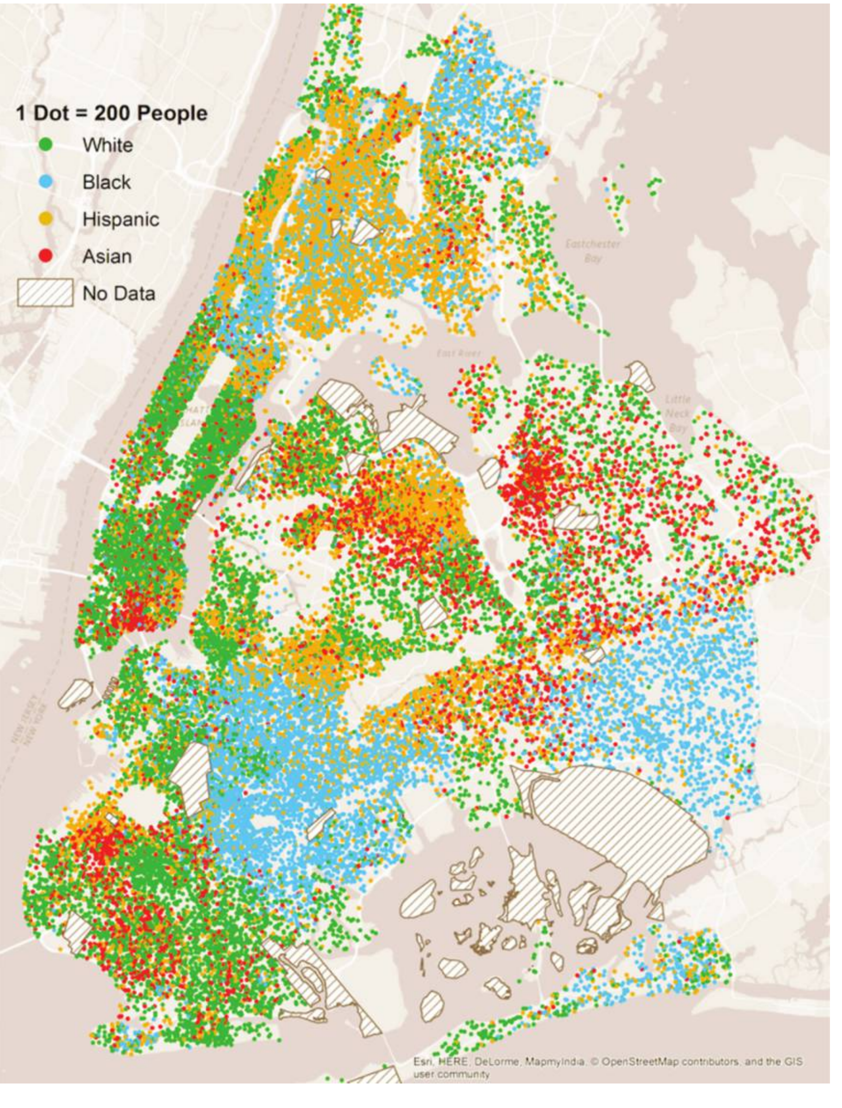

Fig. 2.—New York City population by race or ethnicity, 2010. This figure depicts the residential NYC population in terms of four demographic categories that cover 97 percent of the population. Each dot represents 200 people. Tract-level population data are from the 2010 Census of Population.

between two tracts as the Euclidean distance between the vectors containing the two tracts' five residential population shares. Specifically, defining $\text{shares}_{\text{tract}}$ as the five-element vector containing these population shares, the "Euclidean demographic distance" (henceforth, EDD) between origin and destination tracts is

$$\| \text{shares}_{\text{origin}} - \text{shares}_{\text{destination}} \| / \sqrt{2},$$

where $\|\cdot\|$ indicates the L^2 norm. This measure ranges from zero to one. Figure 3 illustrates the EDD for two origin tracts and many destination tracts. The Morningside Heights origin in the left panel has a diverse population that is similar to that of most NYC tracts, and thus its EDD to most destinations is low. The Manhattan Chinatown origin in the right panel is overwhelmingly Asian and thus quite demographically distant from most tracts, with the exception of the Flushing Chinatown in Queens.

To allow social frictions to depend not only on the demographic composition of the tract in which the restaurant is located but also on the surrounding demographic composition, we calculate the Echenique and Fryer (2007) spectral segregation index (henceforth, SSI) for the modal residential race or ethnicity in each census tract. This index measures the degree to which a census tract borders census tracts of the same residential demographic plurality and the further degree to which those tracts themselves border tracts of the same plurality, ad infinitum. For example, in figure 2, the black census tracts at the center of the cluster of blue dots in Queens, on the right edge of the map, will have higher SSI values than those at the edge of the cluster.

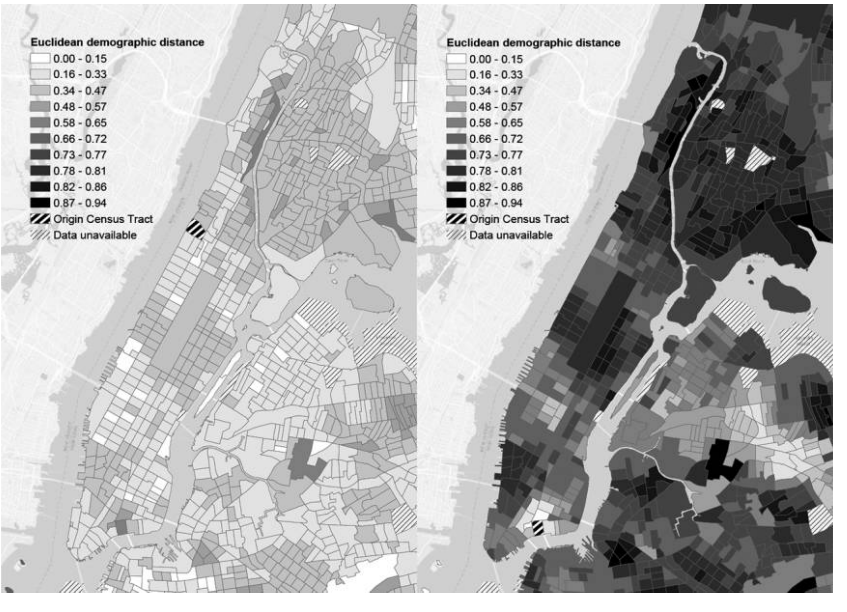

FIG. 3.—Euclidean demographic distances from two census tracts. These maps depict EDD from an origin tract to other NYC tracts. In the left panel, the origin tract is in Morningside Heights; in the right panel, in Manhattan's Chinatown. Demographic data are from the 2010 Census of Population.

We also allow income and crime levels in a restaurant's tract to influence consumer choices. The data on median household incomes come from the 2007–11 American Community Survey 5-Year Estimate. To measure crime rates by location, we compute tract-level robbery statistics for 2007–11 using confidential, geocoded incident-level reports provided to us by the New York Police Department.[8] We use robberies as our crime measure because these are likely the most common and relevant threat to individuals visiting a restaurant. All tract-level characteristics are summarized in tables 1 and A2.

C. *Observed Behavior and Frictions*

Individual users' reviews suggest that both proximity and venue characteristics influence their behavior. Figure 4 maps home, work, and restaurant review locations for two individuals in our sample. The reviewer in the left panel lives and works in midtown Manhattan. The other reviewer works in midtown Manhattan and resides in a southeastern Manhattan development called Stuyvesant Town. Both individuals primarily review venues that are near their home or work locations. At the same time, both reviewers visit more downtown venues than uptown venues, which may reflect differences in the quantity or quality of venues in these areas.

The choices made by all reviewers in our estimation sample suggest the importance of spatial and social frictions. Figure 5 plots, for all reviewers in our estimation sample, the density of three covariates for the set of venues they reviewed and for a random sample of venues that they did not review. The top-left panel depicts transit times from home, the top-right panel transit times from work, and the bottom panel the EDD between the home census tract and the venue's tract. The plots show that, unconditionally, Yelp reviewers are more likely to review restaurants that are closer to their residential and workplace locations and located in tracts with demographics more similar to those of their home tract.[9]

III. Empirical Approach

To measure the relative importance of tastes and spatial and social frictions in determining the restaurant choices of consumers of different races or ethnicities, we introduce a discrete-choice model of restaurant visits. Section III.A describes the assumptions we impose on consumers' preferences. Since we observe restaurant reviews, not all restaurant visits,

[8] Fewer than 3 percent of the Yelp reviews in our estimation sample were posted outside of 2007–11.

[9] The fact that both reviewed and unreviewed venues have shorter travel times from work than from home reflects the fact that most venues and workplaces are in Manhattan.

FIG. 4.—Two reviewers' locations and restaurant reviews. These two maps display two reviewers' home and work locations and the Yelp restaurant venues they reviewed. Dots denote Yelp venues reviewed by this user. The H denotes the average coordinates of those venues identified as home locations in the text of this user's reviews. The W denotes the similarly defined work location.

assumptions on the review-writing behavior of Yelp reviewers are necessary for identification of their preferences. Section III.B introduces these assumptions. Section III.C describes the steps we follow to estimate the model using the data introduced in Section II. In Section III.D, we introduce several extensions that relax the key identifying assumptions in our baseline model.

A. *Demand Specification*

Individuals decide whether to visit any venue and, if they do, which venue to visit. We index individuals by i, venues by j, and by t the occasions on which i needs to decide on whether to visit a venue. In our empirical application, we assume that the set J of potential venues that a consumer may visit is the set of all NYC restaurants listed on Yelp and located in a census tract for which information on residents' median income is available.[10] We denote the outside option of not visiting any venue by $j = 0$.

[10] Specifically, we restrict the set of restaurants to only those with price and rating information listed on Yelp in June 2011, when we collected our data.

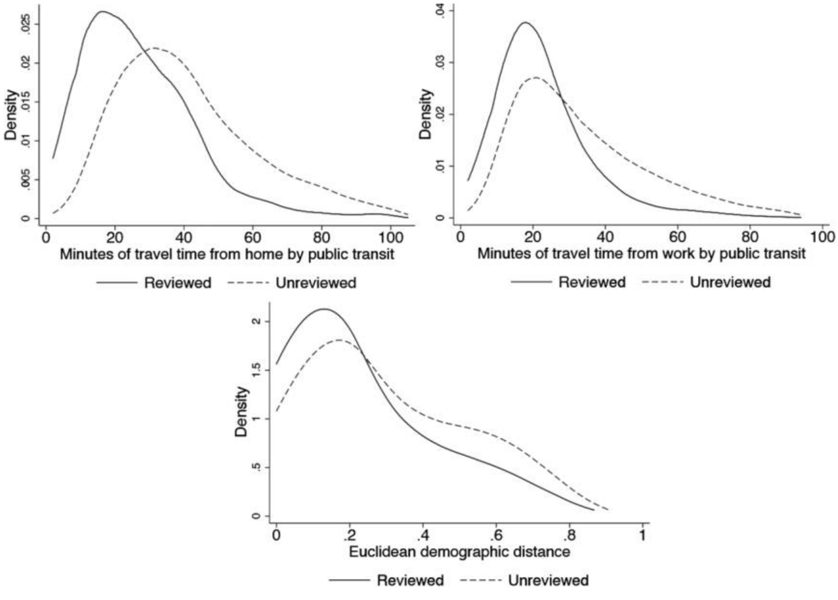

FIG. 5.—Travel times, demographic differences, and consumer choice. These plots are kernel densities for three distributions of reviewer-venue pairs: those venues chosen by reviewers in our estimation sample and a random sample of venues not chosen by these reviewers. The top-left panel plots the densities of travel time from home by public transit; the top-right panel shows travel time from work by public transit; the bottom panel shows EDDs. These plots use Epanechnikov kernels with bandwidths of 3, 3 and 0.1, respectively.

When visiting a venue, individuals must choose whether to visit it from home, from work, or by deviating from their commuting path and whether to travel via public transit or car. We index pairs of origin locations and transportation modes by l and assume that a trip to a venue may be one of six types: from home via car ($l = hc$), from home via public transit ($l = hp$), from work via car ($l = wc$), from work via public transit ($l = wp$), from the commuting path via car ($l = pc$), or from the commuting path via public transit ($l = pp$). We denote the set of these six potential origin-mode pairs as $\mathcal{L} \equiv \{hc, hp, wc, wp, pc, pp\}$.

We allow preferences for restaurants, trip origin locations, and transportation modes to differ across racial or ethnic groups. We index groups by g, which may take three values: white or Hispanic ($g = w$), Asian ($g = a$), and black ($g = b$). We denote the set of these three potential groups as $\mathcal{G} \equiv \{w, a, b\}$.

For an individual i belonging to the racial or ethnic group $g(i)$, we assume that her utility of visiting restaurant j on occasion t from origin-mode l is

$$U_{ijlt} = \gamma^1_{g(i)l} X^1_{ijl} + \gamma^2_{g(i)} X^2_{ij} + \beta^1_{g(i)} Z^1_j + \beta^2_{g(i)} Z^2_{ij} + \nu_{ijlt}, \tag{1}$$

where X^1_{ijl} measures the spatial frictions that i incurs when visiting j from l, the vector X^2_{ij} measures the social frictions that may affect the appeal that restaurant j has to individual i, and Z^1_j and Z^2_{ij} control for other observed venue and individual-venue-specific characteristics, respectively. Specifically, the variable X^1_{ijl} is the log of the number of minutes it takes individual i to travel to restaurant j using the origin-mode pair indexed by l.[11] The vector X^2_{ij} contains the EDD and SSI measures introduced in Section II and the residential population share of each racial and ethnic group in the restaurant's tract. The venue characteristics in Z^1_j are the restaurant's price and Yelp rating, the log median household income of the tract in which the venue is located, 28 area dummies, and nine cuisine dummies.[12] Finally, the vector Z^2_{ij} includes the restaurant's price and Yelp rating interacted with the reviewer's home census tract's median household income, as well as the percent difference and absolute percent difference in median incomes between the home and restaurant tracts.

The variable ν_{ijlt} is a scalar unobserved by the econometrician. We allow all preference parameters $(\gamma^1, \gamma^2, \beta^1, \beta^2)$ to vary across demographic groups g, and since the coefficient γ^1 is l specific, we additionally allow the marginal disutility of a trip to flexibly depend on both its origin and the mode of transit.

Although our data set describes reviewers' home and work locations, it does not indicate the origin-mode l of each trip. We address this data limitation by assuming that consumers jointly optimize the restaurant they patronize and the origin-mode from which they do so, choosing thus the jl combination that maximizes their utility. Accordingly, defining a dummy variable d_{ijlt} that equals one if individual i travels to venue j from origin-mode l at period t, we assume that

$$d_{ijlt} = \mathbf{1}\{U_{ijlt} \geq U_{ij'l't};\forall\ j' \in J,\ l' \in \mathcal{L}\}, \tag{2}$$

where $\mathbf{1}\{A\}$ is an indicator function that equals one if A is true. We also define a variable d_{ijt} that is one if individual i chooses venue j at period t, $d_{ijt} = \Sigma_{l\in\mathcal{L}} d_{ijlt}$, irrespective of the origin-mode of the trip.

In our benchmark specification, we assume that the vector of unobserved utilities for individual i at period t, $\nu_{it} = \{\nu_{ijlt};\forall\ j \in J,\ l \in \mathcal{L}\}$, is independent across individuals and time periods and has a joint type I extreme value distribution: its cumulative distribution function is

[11] The disutility of travel time, $\gamma^1_{g(i)l}$, may vary with l because the direct pecuniary cost of an additional minute of travel time differs across modes of transportation (positive for taxis, zero for the subway). These coefficients may also vary because of heterogeneity across transportation modes or origin locations in nonpecuniary costs (e.g., cleanliness or convenience).

[12] The restaurant's price is captured by dummy variables corresponding to Yelp's four price categories. The area dummies are aggregates of NYC community districts, and the nine cuisine dummies aggregate Yelp's more detailed cuisine categories; see app. B.3.

$F(\nu_{it}) = \exp(-\Sigma_{j\in J}\Sigma_{l\in\mathcal{L}}\exp(-\nu_{ijlt}))$. This distribution yields a conditional-logit discrete-choice model of restaurant visits.

B. Review-Writing Behavior

Let d^*_{ijt} be a dummy variable that equals one if individual i writes a review of restaurant j at time t. The fact that we observe reviews rather than restaurant visits (i.e., we observe d^*_{ijt} but not d_{ijt}) implies that estimating the preference parameters in equation (1) requires making assumptions on the review-writing behavior of Yelp reviewers. We impose three assumptions. First, users do not review restaurants they have not visited. Second, they write at most one review per restaurant (independently of how many times they visit a restaurant). Third, conditional on having visited a restaurant and not having previously reviewed it, they write a review with a probability p^*_{it} that is independent of the restaurant's characteristics and the origin-mode of the trip.

C. Estimation Procedure

We estimate the preference parameters in equation (1) using a maximum likelihood estimator. To derive the relevant likelihood function, we implement the following five steps. Additional details of the mathematical derivations appear in appendix C.1.

Step 1: *Derive restaurant visit probability*. According to the assumptions in Section III.A, the probability that individual i visits venue j from origin-mode l at period t is

$$P(d_{ijlt} = 1|X_i, Z_i, J; (\gamma, \beta)) = \exp(V_{ijl})/\sum_{j'\in J}\left[\sum_{l'\in\mathcal{L}}\exp(V_{ij'l'})\right], \tag{3}$$

where X_i, Z_i, γ, and β are vectors that collect their respective terms and[13]

$$V_{ijl} \equiv \gamma^1_{g(i)l}X^1_{ijl} + \gamma^2_{g(i)}X^2_{ij} + \beta^1_{g(i)}Z^1_j + \beta^2_{g(i)}Z^2_{ij}. \tag{4}$$

The probability that individual i visits venue j at period t is then

$$\begin{aligned} P(d_{ijt} = 1|X_i, Z_i, J; (\gamma, \beta)) &= P\left(\sum_{l\in\mathcal{L}}d_{ijlt} = 1|X_i, Z_i, J; (\gamma, \beta)\right) \\ &= \sum_{l\in\mathcal{L}}P(d_{ijlt} = 1|X_i, Z_i, J; (\gamma, \beta)) \\ &= \sum_{l\in\mathcal{L}}\exp(V_{ijl})/\sum_{j'\in J}\left[\sum_{l'\in\mathcal{L}}\exp(V_{ij'l'})\right]. \end{aligned} \tag{5}$$

[13] Formally, $X_i \equiv \{(X^1_{ijl}, X^2_{ij}); \forall\, j\in J, l\in\mathcal{L}\}$, $Z_i \equiv \{(Z^1_j, Z^2_{ij}); \forall\, j\in J\}$, $\gamma \equiv \{(\gamma^1_{gl}, \gamma^2_g); \forall\, g\in \mathcal{G}, l\in\mathcal{L}\}$, and $\beta \equiv \{(\beta^1_g, \beta^2_g); \forall\, g\in\mathcal{G}\}$.

The first equality applies the definition of d_{ijt}, the second one takes into account that the origin-mode of a trip is unique (i.e., the joint probability that individual i visits restaurant j at period t from two different origins l and l' is zero), and the third equality uses equation (3).

Step 2: *Derive restaurant review probability*. Equation (5) and the review-writing model described in Section III.B imply that the probability of observing a review of venue j written by individual i at period t is

$$\begin{aligned} &P\left(d^*_{ijt} = 1 | X_i, Z_i, J_{it}, J'_{it}; (\gamma, \beta, p^*_{it})\right) \\ &= p^*_{it}\mathbf{1}\{j \neq 0, j \in J'_{it}\}P\left(d_{ijt} = 1 | X_i, Z_i, J; (\gamma, \beta)\right), \end{aligned} \tag{6}$$

where J'_{it} denotes the set of restaurants not previously reviewed by i, that is, $J'_{it} \equiv \{j \in J : d^*_{ijt'} = 0 \text{ for all } t' < t \text{ and } j \neq 0\}$.[14] Combining equations (5) and (6), we can derive the probability that individual i reviews restaurant j at period t conditional on i reviewing any restaurant at that period:

$$\begin{aligned} &P\left(d^*_{ijt} = 1 | d^*_{it} = 1, X_i, Z_i, J'_{it}; (\gamma, \beta)\right) \\ &= \frac{\mathbf{1}\{j \neq 0, j \in J'_{it}\}\sum_{l \in \mathcal{L}} \exp\left(V_{ijl}\right)}{\sum_{j \in J'_{it}}\sum_{l \in \mathcal{L}} \exp\left(V_{ij'l}\right)}, \end{aligned} \tag{7}$$

where $d^*_{it} = \Sigma^J_{j=1} d^*_{ijt}$ is a dummy variable that equals one if i writes a review at t.

Step 3: *Reduce choice set*. The cardinality of the choice set J'_{it} makes it computationally burdensome to construct the denominator of the probability in equation (7). As J'_{it} equals the set of all restaurants in NYC, J, minus those reviewed by individual i prior to period t, the large dimensionality of J implies that the set J'_{it} will also be very large.

To address this dimensionality issue, we adapt the choice set reduction procedure from McFadden (1978) to our empirical setting. For every individual i and period t in which we observe a review written by i, we define a set S_{it} that is a subset of J'_{it}. We construct S_{it} by including the restaurant j for which $d^*_{ijt} = 1$ plus a random subset of all other alternatives in J'_{it}, selecting them from J'_{it} with equal probability. As all elements of S_{it} other than the actual choice of i at t are selected randomly, the set S_{it} itself is random. We denote by $\pi(S_{it} | d^*_{ijt} = 1, J'_{it})$ the probability of assigning the set S_{it} to an individual i who reviewed venue j at t. Our sampling scheme implies that

[14] As reflected in eq. (6), if individuals were to review every restaurant they visit for the first time, $p^*_{it} = 1$, the probability of observing a review would equal either the probability in eq. (5) (for any venue j that user i has not previously visited) or zero (for any previously visited venue).

$$\pi(S_{it}|d^*_{ijt} = 1, J'_{it}) = \begin{cases} \kappa_{it} & \text{if } j \in S_{it} \\ 0 & \text{otherwise,} \end{cases} \tag{8}$$

where $\kappa_{it} \in (0, 1)$ is a constant determined by our choice of the number of venues in S_{it} and the number of venues in J'_{it}.

Given equations (7) and (8), we can write the probability that i reviews restaurant j at period t conditional on a randomly drawn set S_{it} and that i writes a review at t as

$$P(d^*_{ijt} = 1|d^*_{it} = 1, X_i, Z_i, S_{it}; (\gamma, \beta)) = \frac{\mathbf{1}\{j \in S_{it}\}\sum_{l\in\mathcal{L}} \exp(V_{ijl})}{\sum_{j'\in S_{it}}\sum_{l\in\mathcal{L}} \exp(V_{ij'l})}. \tag{9}$$

Importantly, to be able to randomly draw a set S_{it} from the set of nonreviewed restaurants J'_{it}, one needs to observe all reviews previously written by user i.

Step 4: *Derive individual i–specific likelihood function.* Using j_{it} to denote the restaurant reviewed by individual i at period t, the joint probability of observing an individual i writing the T_i reviews $\{j_{i1}, j_{i2}, \ldots, j_{iT_i}\}$ conditional on observing a review written by i in each of the periods $\{1, \ldots, T_i\}$ and on randomly drawing the sets $\{S_{i1}, S_{i2}, \ldots, S_{iT_i}\}$ is

$$\prod_{t=1}^{T_i} \frac{\mathbf{1}\{j \in S_{it}\}\sum_{l\in\mathcal{L}} \exp(V_{ijl})}{\sum_{j'\in S_{it}}\sum_{l\in\mathcal{L}} \exp(V_{ij'l})}. \tag{10}$$

This joint probability is simply the product of the corresponding marginal probabilities in equation (9). Intuitively, all the dynamic effects of a review written by individual i at period t are reflected in the subsequent choice sets $\{J'_{is}, \forall\ s > t\}$, and the effect of each of these choice sets is subsumed in the randomly selected subsets $\{S_{is}, \forall\ s > t\}$.

Step 5: *Derive the log-likelihood function.* Given equation (10) and assuming that we observe a random sample $i = 1, \ldots, N$ of individuals from the population of interest, we can write the log-likelihood function as

$$\sum_{i=1}^{N}\sum_{t=1}^{T_i}\sum_{j\in S_{it}} \mathbf{1}\{d^*_{ijt} = 1\} \ln\left(\frac{\sum_{l\in\mathcal{L}} \exp(V_{ijl})}{\sum_{j'\in S_{it}}\sum_{l\in\mathcal{L}} \exp(V_{ij'l})}\right), \tag{11}$$

where V_{ijl} is the function of the parameter vector of interest (γ, β) defined in equation (4).

All the estimates presented in Section IV are computed as the values of (γ, β) that maximize the function in equation (11). As proved in McFadden (1978), these maximands are consistent estimators of the preference parameters defined in equation (1).[15] Specifically, the fact that we maxi-

[15] Appendix C.2 presents simulation results that illustrate the asymptotic properties of our estimator. Section IV.D illustrates its finite-sample properties.

mize a likelihood function that conditions on the randomly chosen sets $\{S_{it}; \forall i, t\}$ does not affect the consistency of our estimator. However, the variance of this estimator decreases as we increase the cardinality of each set S_{it}. While our benchmark estimates use sets $\{S_{it}; \forall i, t\}$ with 20 restaurants each, tables A4–A6 show that our conclusions are robust to using sets that include 50 or 100 restaurants. In practice, larger choice sets incur greater computation times without appreciably shrinking our standard errors.

D. Discussion

The baseline model described in Sections III.A and III.B embeds several key identifying assumptions. In this section, we discuss why we impose these assumptions, what they imply in our empirical context, and how we relax them in several extensions to this baseline model.

Absence of race-specific preferences for restaurants' unobserved characteristics.—Except for the vector of idiosyncratic errors ν_{it}, the utility function in equation (1) assumes that consumers' preferences exclusively depend on observed restaurant characteristics (see Sec. IV for a complete description of the covariates (X_i, Z_i) entering our demand model). However, individuals from different racial groups may have heterogeneous preferences for restaurants on the basis of characteristics that we do not observe. For example, group g consumers may prefer a specific venue because this venue is frequently patronized by other customers belonging to the same group g. To explore the robustness of our baseline results to the presence of unobserved venue characteristics that determine consumers' preferences, we generalize the utility function in equation (1) to allow for race- and restaurant-specific unobserved effects:

$$\begin{aligned} U_{ijlt} = {} & \gamma^1_{g(i)l} X^1_{ijl} + \gamma^2_{g(i)} X^2_{ij} + \beta^1_{g(i)} Z^1_j + \beta^2_{g(i)} Z^2_{ij} \\ & + \sum_{j' \in J} \alpha_{g(i)j'} \mathbf{1}\{j' = j\} + \nu_{ijlt}, \end{aligned} \tag{12}$$

where α_{gj} captures the group g–specific component of the utility of visiting restaurant j that is determined by unobserved characteristics.[16]

[16] Incorporating the set of fixed effects $\{\alpha_{gj}; \forall\, g \in \mathcal{G}, j \in \mathcal{J}\}$ significantly increases the dimensionality of the parameter vector we estimate. To explore computational costs, we have estimated this specification in two ways. First, we use a procedure similar to that in Sec. III.C but with two adjustments: (*a*) we add the term $\Sigma_{j' \in J} \alpha_{g(i)j'} \mathbf{1}\{j' = j\}$ to the expression for V_{ijl} in eq. (4); (*b*) for each individual i and period t, we form the set S_{it} by drawing only from those restaurants reviewed by at least one sample reviewer that belongs to the same group g as individual i. Importantly, as Train (2009, chap. 3, sec. 7) discusses, while the procedure to sample restaurants described in adjustment b implies that our estimates of the restaurant fixed effects are biased, it does not affect the consistency of our estimates of the spatial and social frictions. Second, the Poisson approximation to the conditional-logit model is described in Taddy (2015) and implemented in Gentzkow, Shapiro, and

Independence of irrelevant alternatives.—Our baseline model assumes that all unobserved determinants of the utility to individual i of patronizing venue j at period t from origin-mode l, as captured in the scalar ν_{ijlt}, are independent across venues, possible origins of the trip, and modes of transport. To illustrate the robustness of our baseline estimates, we relax this independence assumption in multiple directions.

First, in appendix C.3, we introduce an alternative model in which we assume that ν_{ijlt} does not vary with the origin-mode l; that is, $\nu_{ijlt} = \nu_{ijl't}$ for any pair (l, l'). An implication of this alternative model is that, conditional on visiting a restaurant j, every individual i travels to it using the origin and mode of transport that maximizes the term $\gamma^1_{g(i)l}X^1_{ijl}$. As discussed in Section III.A, the covariate X^1_{ijl} equals the (log) number of minutes that it takes individual i to reach venue j from the origin-model l; thus, assuming that $\nu_{ijlt} = \nu_{ijl't}$ and $\gamma^1_{g(i)l} = \gamma^1_{g(i)l'} < 0$ for any pair (l, l') is equivalent to assuming that individuals always choose the origin-mode pair that minimizes travel time to each venue.

Second, in Section IV.E.2, we present estimates of nested-logit models that allow for correlation in the unobserved terms ν_{ijlt} across restaurants j and origin-modes l. Specifically, we allow the terms ν_{ijlt} to be correlated across restaurants that share a number of characteristics. Following Train, McFadden, and Ben-Akiva (1987), appendix C4 describes how we adapt the estimation procedure in Section III.C to these nested-logit demand models.

Absence of within-group unobserved parameter heterogeneity.—One additional limitation of the conditional-logit model described in Section III.A is that it does not allow for within–group g unobserved heterogeneity in the parameters capturing preferences for observable restaurant characteristics. The standard approach to do so is to assume that individual-specific preferences follow a known distribution in the population of interest. In our setting, estimating such a model is infeasible: unobserved heterogeneity in the vector (γ, β) makes the choice set construction procedure in McFadden (1978) inapplicable and, therefore, requires estimating a likelihood function that, for each individual i and period t in our sample, depends on the actual choice set J'_{it}. This is computationally infeasible in a city with thousands of restaurants like NYC.[17]

While we do not allow for within–group g unobserved heterogeneity in preferences, the utility function in equation (1) does allow preferences

Taddy (2019). In our setting, this approximation is exact if all individuals in the sample have equal expected utility from each restaurant trip. The results from these two estimation procedures are described in Sec. IV.E.1 and app. D.3, respectively.

[17] Katz (2007) and Pakes (2010) show that there is a moment inequality approach that allows one to handle both large choice sets and unobserved heterogeneity in preferences for observed choice characteristics. We discuss in app. C.5 the relative advantages and disadvantages of this moment inequality approach for our particular application.

to vary within groups with the observed characteristics of the home and work census tracts of each individual. Namely, X_{ij}^2 and Z_{ij}^2 contain interactions of individual i and restaurant j characteristics. For example, we allow individuals living in tracts of different income levels to value restaurants' prices and ratings differently.

Exogeneity of home and work locations.—Section III.A implicitly assumes that individuals' home and work locations are exogenously given. In practice, individuals choose where to live and work, and these locations may be determined as a function of restaurant characteristics. However, the endogenous location of home and work will not bias our estimates of the preference parameters (γ, β) if the distribution of the vector of unobserved characteristics affecting individuals' restaurant choices, ν_{it}, is independent of the characteristics determining the optimal selection of home and work location. Note that this is compatible with the vector of observed characteristics (X_i, Z_i) affecting individuals' endogenous home and work locations.[18]

Restaurant and origin independence of review-writing probabilities.—As described in Section III.B, our baseline model assumes that the probability that an individual writes a review about a visited restaurant does not depend on the restaurant itself nor on the origin of the trip. This allows the review-writing decision to depend on the consumption experience in a number of ways. First, our assumption allows arbitrary variation in the propensity to write a review across reviewers and time. It can therefore account for the fact that reviewers are more likely to contribute to an online platform when they are nearing a reputational reward, such as Yelp's "elite" status (Luca 2016b). Second, our model is consistent with individuals being more likely to write reviews about dining experiences that surprised them, either negatively or positively. Surprises are, by definition, independent of the variables that are in the information set of consumers when deciding which restaurant venue to patronize and, therefore, independent of the consumers' restaurant choice.

In robustness checks, we address three possible violations of our assumption that the review-writing probability is independent of the patronage choice. First, one could claim that users are more likely to review restaurants with few prior reviews or that are not chain establishments well known by most consumers. As we show in appendix C7, one may control

[18] In fact, if individuals' home and work locations are determined as a function of the expected utility of restaurant consumption, then our estimates are less likely to be biased the larger the set of characteristics that we explicitly control for through the vector (X_i, Z_i). Intuitively, the fewer the variables that are accounted for by the unobserved term ν_{it}, the more likely it is that this composite is independent of the characteristics determining each individual's choice of home and work location. A detailed discussion of this point is contained in app. C.6.

for characteristics that affect the review-writing probabilities of Yelp reviewers by introducing them explicitly as covariates in our conditional-logit model. In Section IV.E.3, we report estimates in which we control for a restaurant's total number of reviews and whether it belongs to a chain with more than eight NYC locations.

Second, it is possible that users are more likely to review restaurants that they want to signal they have patronized. The specifications in which we introduce race-specific restaurant fixed effects allow us to control for the possibility that, for example, individuals want to signal that they have visited a restaurant that is idiosyncratically popular with members of group g. We discuss these estimates in Section IV.E.1.

Finally, users may be more or less likely to review restaurants that they visited from a particular origin, such as a business lunch near their workplace. To address this possibility, we introduce race-origin-mode-specific fixed effects in specifications reported in Section IV.E.3.

Lack of serial correlation in unobserved preferences.—Our baseline model assumes that individuals' unobserved restaurant preferences (captured in the vector ν_{it}) are independent over time. As discussed in detail in appendix C8, if, contrary to our assumption, the preference shocks ν_{it} are serially correlated, the fact that we identify users' preferences from their reviews and that users do not review a restaurant twice will generate a selection bias in our estimates of consumers' preference parameters. Specifically, positive serial correlation would cause attenuation bias: upward bias in the estimates of coefficients on characteristics that consumers dislike (e.g., spatial and social frictions) and downward bias in the estimates of coefficients on characteristics that appeal to consumers (e.g., restaurants' rating). We illustrate the possible size of this bias through a simulation in appendix C8. To reduce this selection bias, we report estimates that use only the first half and first fifth of each user's reviews in Section IV.E.3.

IV. Estimation Results

This section reports the results of estimating discrete-choice models of the form described in Section III using the data introduced in Section II. The models differ in the set of spatial and social frictions we incorporate. In Section IV.A, we introduce spatial frictions while omitting social frictions. In Section IV.B, we incorporate these measures. In all cases, we include a set of venue and reviewer-venue characteristics that may influence consumer demand.

All the specifications presented in the main text are estimated using a fixed set of randomly generated choice sets, $\{S_{it}, \forall i, t\}$, so that variation in the estimates across columns and tables is exclusively due to variation in the included covariates.

A. *Spatial Frictions*

Columns 1–3 of table 2 estimate spatial frictions while omitting social frictions. All the coefficients on spatial frictions are negative. There is only modest heterogeneity by race, with somewhat lower disutilities of transit times for Asian reviewers.[19] They are less precisely estimated for black reviewers, because of the smaller number of observations. The coefficients on transit times from work are larger than those on travel times from home. This may reflect a higher cost of time spent away from work. The magnitude of the coefficients is modestly but nearly always higher for travel by car than by public transit. This could reflect the fact that NYC public transit fares are invariant to distance while taxi fares are not.

Columns 1–3 of table 2 deliver a clear finding: spatial frictions play a first-order role in individuals' consumption choices within the city. Consumers are less likely to visit venues that, in terms of mass transit and automobile travel time, are more distant from their home and work locations, as well as the commuting path between these. Consider two hypothetical restaurants, identical in their characteristics except for the number of minutes away from the individual's optimal origin of the trip. The first restaurant is 15 minutes from the individual's workplace by car; the second restaurant is 30 minutes away. The estimated coefficients in columns 1–3 of table 2 imply that the individual would be about four times as likely to visit the more proximate venue from work by car (e.g., for a black reviewer, $2^{2.02} \approx 4.06$). Similarly, if the two restaurants were 15 and 30 minutes from the commuting path by public transit, the individual would be about twice as likely to visit the more proximate venue.

Finally, note that reviewers' choices also depend on restaurants' characteristics in predictable ways. Restaurants with higher ratings and lower prices are generally more attractive. However, restaurants in the $$ price category are attractive relative to $ restaurants, indicating that prices also reflect quality.[20] Reviewers residing in census tracts with higher incomes exhibit more willingness to pay higher prices. Asian reviewers' most preferred cuisine category is Asian cuisine, while white/Hispanic reviewers' most preferred categories are Latin American, American, and vegetarian/vegan.

B. *Social Frictions*

Columns 4–6 in table 2 introduce social frictions.

Reviewers exhibit homophily at the individual level. Reviewers are more likely to visit venues located in tracts with a larger residential population of

[19] Table A10 shows that there is little heterogeneity in the coefficients on spatial frictions along age, gender, and income dimensions.

[20] Restaurants in the $ category are often fast-food venues.

the same race. Asian reviewers are more likely to visit a restaurant in a tract with more Asian residents relative to all other residential racial and ethnic categories (white residents are the omitted category). Similarly, black reviewers are more likely to visit a restaurant in a tract with more black residents. In the case of white or Hispanic reviewers, homophily is less evident, perhaps because we do not distinguish between whites and Hispanics when classifying reviewers' profile photos. In column 6, there is a positive coefficient on Hispanic population share, but there is also a positive coefficient on Asian residents relative to white residents.

The negative coefficients on EDD reveal a role for environmental similarity. Reviewers are less likely to visit venues located in census tracts with demographics different from those of their home census tract. Since we control for transit times between tracts, this result cannot be attributed to the joint impact of residential segregation and disutility of travel. Similarly, our controls include income differences between tracts, so this result cannot be attributed to spatial differences in incomes predicting consumers' choices. Thus, the coefficient on EDD likely captures mechanisms linked to racial and ethnic differences. Individuals may have preferences regarding the residential demographics of the neighborhoods in which they reside and those in which they consume. Alternatively, consumers may be more likely to visit restaurants located near their friends' residences, with these social ties being predicted by neighborhoods' demographic similarity.

The estimated coefficient on EDD implies an economically significant role for this social friction. Consider a user who contemplates visiting two venues that are identical except for their EDDs, which differ by one standard deviation. Our estimates imply that an Asian user would be 25 percent more likely to visit the venue in the more demographically similar census tract.[21] A black user would be 51 percent more likely to visit the more similar tract. We can also express the economic significance of demographic differences as a trade-off between demographic distance and transit time. To hold constant an Asian consumer's utility visiting a venue from home via public transit, a venue one standard deviation more demographically distant would have to be about 21 percent closer in terms of travel time.[22] For a black consumer, it would have to be 44 percent closer.

[21] Comparing two venues j and j' that are identical in every covariate except X_{ij}^2,

$$P(d_{ij} = 1|V_i)/P(d_{ij'} = 1|V_i) = \exp\left(\gamma_{g(i)}^2\left(X_{ij}^2 - X_{ij'}^2\right)\right).$$

Table A2 shows that the standard deviation of EDD across all pairs of census tracts in NYC is 0.226, so the coefficient of -1 in col. 4 of table 2 implies that a venue that has an EDD 0.226 lower than an otherwise identical venue will be visited with 25 percent higher probability: $\exp(-1.00 \times (-0.226)) \approx 1.25$.

[22] To hold U_{ijlt} constant, a change of ΔX_{ij}^2 would be offset by the change $\Delta X_{ijl}^1 = -\gamma_{g(i)}^2 X_{ij}^2/\gamma_{g(i)l}^1$. Since the coefficient $\gamma_{g(i)hp}^1$ estimated in col. 4 of table 2 is -1.06, the change required to offset a one standard deviation increase in EDD is $-1.00\times 0.226/1.06 \approx -0.21$.

TABLE 2
SPATIAL AND SOCIAL FRICTIONS ESTIMATES

	Spatial Frictions			Social Frictions		
	Asian (1)	Black (2)	White/Hispanic (3)	Asian (4)	Black (5)	White/Hispanic (6)
Log travel time from home by public transit	−1.07***	−.996***	−1.15***	−1.06***	−.938***	−1.13***
	(.101)	(.119)	(.058)	(.107)	(.127)	(.059)
Log travel time from home by car	−1.19***	−1.24***	−1.38***	−1.17***	−1.19***	−1.36***
	(.086)	(.141)	(.059)	(.091)	(.158)	(.060)
Log travel time from work by public transit	−1.27***	−2.16	−1.92***	−1.24***	−1.85*	−1.87***
	(.145)	(2.43)	(.298)	(.149)	(1.11)	(.287)
Log travel time from work by car	−1.69***	−2.02***	−2.01***	−1.60***	−1.79***	−1.95***
	(.188)	(.584)	(.181)	(.176)	(.459)	(.171)
Log travel time from commute by public transit	−.955***	−.997***	−1.11***	−.943***	−.930***	−1.10***
	(.063)	(.098)	(.042)	(.067)	(.105)	(.044)
Log travel time from commute by car	−1.08***	−1.43***	−1.46***	−1.04***	−1.32***	−1.43***
	(.060)	(.171)	(.056)	(.061)	(.177)	(.058)
Euclidean demographic distance (EDD)				−1.00***	−1.84***	−1.19***
				(.121)	(.280)	(.130)
Spectral segregation index (SSI) of k_j				.150***	.075	.045*
				(.051)	(.093)	(.027)
EDD × SSI				−.149	−.171	−.068
				(.117)	(.239)	(.083)
Share of tract population that is Asian				1.03***	.011	.363***
				(.120)	(.345)	(.138)
Share of tract population that is black				.220	1.08***	.140
				(.319)	(.399)	(.265)
Share of tract population that is Hispanic				−.251	.467	.415**
				(.235)	(.381)	(.188)
Share of tract population that is other				.059	3.56	.484
				(2.07)	(3.43)	(1.99)
Dummy for $$ bin	.309***	.645***	.317***	.375***	.771***	.355***
	(.087)	(.194)	(.082)	(.087)	(.197)	(.083)

TABLE 2 (*Continued*)

	Spatial Frictions			Social Frictions		
	Asian (1)	Black (2)	White/Hispanic (3)	Asian (4)	Black (5)	White/Hispanic (6)
Dummy for $$$ bin	.175	−.283	−.075	.287**	−.090	−.026
	(.115)	(.334)	(.120)	(.116)	(.341)	(.120)
Dummy for $$$$ bin	.086	−.313	−.398*	.220	−.074	−.347
	(.185)	(1.18)	(.219)	(.188)	(1.22)	(.221)
Yelp rating of restaurant	.583***	.036	.335***	.579***	.053	.344***
	(.064)	(.137)	(.059)	(.064)	(.138)	(.059)
African cuisine category	.271	−.046	.319	.280	−.198	.298
	(.297)	(.548)	(.259)	(.299)	(.553)	(.261)
American cuisine category	.421***	.542***	.596***	.432***	.523***	.591***
	(.054)	(.118)	(.050)	(.054)	(.119)	(.050)
Asian cuisine category	.931***	.201	.308***	.886***	.255*	.307***
	(.054)	(.132)	(.054)	(.054)	(.134)	(.054)
European cuisine category	.204***	−.339**	.247***	.195***	−.326**	.235***
	(.059)	(.153)	(.056)	(.059)	(.154)	(.056)
Indian cuisine category	.374***	−.422	−.018	.370***	−.451	−.039
	(.091)	(.299)	(.097)	(.091)	(.301)	(.097)
Latin American cuisine category	.491***	1.03***	.699***	.517***	1.01***	.690***
	(.070)	(.134)	(.061)	(.070)	(.136)	(.062)
Middle Eastern cuisine category	.264***	.066	.204**	.280***	.104	.203**
	(.100)	(.250)	(.094)	(.101)	(.251)	(.094)

Vegetarian/vegan cuisine category	.365***	−.041	.596***	.392***	.001	.587***
	(.138)	(.408)	(.116)	(.138)	(.409)	(.116)
$$ bin × home tract median income	.041***	−.002	.049***	.034***	−.022	.042***
	(.011)	(.032)	(.009)	(.011)	(.032)	(.009)
$$$ bin × home tract median income	.086***	.109**	.089***	.075***	.077	.081***
	(.014)	(.052)	(.013)	(.014)	(.053)	(.013)
$$$$ bin × home tract median income	.088***	−.119	.105***	.074***	−.167	.095***
	(.022)	(.224)	(.022)	(.022)	(.234)	(.023)
Yelp rating × home tract median income	.010	.007	.017***	.011	.008	.016**
	(.008)	(.023)	(.007)	(.008)	(.023)	(.007)
Percent absolute difference in median incomes	−.218***	.469***	−.350***	−.062	.850***	−.100*
	(.045)	(.114)	(.046)	(.050)	(.126)	(.053)
Percent difference in median incomes ($k_j - h_i$)	−.233	1.04	.791***	.114	.619	.719**
	(.292)	(.826)	(.293)	(.305)	(.853)	(.300)
Log median household income in k_j	.119	−.869	−.694***	−.109	−.360	−.625**
	(.258)	(.733)	(.259)	(.267)	(.744)	(.262)
Average annual robberies per resident in k_j				−3.41***	2.43**	−3.74***
				(.676)	(1.20)	(.771)
Number of trips	6,447	1,079	6,936	6,447	1,079	6,936

Note.—Each column reports an estimated conditional-logit model of the decision to visit a Yelp venue. Standard errors are in parentheses. Unreported controls are 28 area dummies.

* Statistically significant at 10 percent.
** Statistically significant at 5 percent.
*** Statistically significant at 1 percent.

The consequences of social frictions vary across individuals of different races because of differences in population sizes. Reviewers of all races are more likely to visit restaurants that have lower values of the EDD covariate, and black individuals have the most negative coefficient on EDD. Yet the mean value of EDD for venues visited by black reviewers (0.34) is in fact greater than the mean value of EDD for venues not visited by white reviewers (0.29). This finding is consistent with the idea in Anderson (2015, 10) that "white people typically avoid black space, but black people are required to navigate the white space as a condition of their existence."

Do demographic differences between census tracts matter more when the venue is located deep within a segregated area? To assess this, we use a spectral segregation index (SSI) that describes a tract's demographic isolation in terms of its racial or ethnic plurality. In table 2, the coefficients on both SSI and the interaction of EDD and SSI are modest in magnitude. When EDD is close to one, their sum is close to zero. Thus, a restaurant in a tract near the edge of a racially or ethnically distinct area is about as likely to be visited as a tract with the same demographic differences located deep inside that area. Individuals' choices are therefore mostly predicted by the demographic composition of the area immediately surrounding the restaurant.

Relative to columns 1–3, the coefficients on spatial frictions and cuisine categories in columns 4–6 of table 2 are slightly attenuated toward zero. Residential segregation means that spatial frictions and social frictions are positively correlated, so spatial frictions will be overestimated if social frictions are omitted. Similarly, if restaurants in tracts with more Asian residents are more likely to serve Asian cuisine that appeals to Asian reviewers, then the cuisine coefficients will be overestimated if social frictions are omitted. Comparing columns 1–3 and 4–6 of table 2 suggests that this occurs, but in the vast majority of cases the estimated coefficients differ by less than a standard error.

Since our coefficients are normalized by the standard deviation of the logit error ν_{ijlt}, comparisons of the levels of coefficients across columns implicitly assume that this standard deviation is constant. The change in the coefficients on spatial frictions and cuisine categories across models should be compared to the change in a coefficient that is plausibly not biased by the omission of social frictions, such as the Yelp rating of the restaurant. Comparing columns 1–3 and 4–6 of table 2, the coefficient on Yelp rating is slightly attenuated for Asian reviewers and actually larger for black and white/Hispanic reviewers. Thus, the attenuation of coefficients on spatial frictions and cuisine categories in columns 4–6 is not solely attributable to a change in the standard deviation of ν_{ijlt} across models. Comparing these results to table A11 shows that including the workplace origin is key to our estimates of social frictions.

This section has documented patterns in consumer behavior that tend to segregate consumption. We find roles for both environmental similarity and individual homophily. Regardless of the particular mechanisms underlying how demographic differences shape consumption in the city, our quantification indicates that these social frictions play an important role in shaping consumer behavior. These elements will contribute to our estimates of urban consumption segregation in Section V.

C. *Model Fit*

In this section, we discuss how well our estimated model fits the data.

1. In-Sample Isolation

We first compare our model's prediction of consumption segregation to that observed in the estimation sample. Following Gentzkow and Shapiro (2011), we compute isolation indices using "leave-out" means to address finite-sample bias. Denote the number of reviews of venue j by members of racial group g by $v_{gj} = \Sigma_{i:g(i)=g}\Sigma_t d^*_{ijt}$, the total number of reviews by those members by $v_g = \Sigma_j v_{gj}$, the number of reviews of venue j by individuals who are not members of group g by $v_{\neg g,j} = \Sigma_{i:g(i)\neq g}\Sigma_t d^*_{ijt}$, and the total number of reviews of venue j by $v_j = \Sigma_g v_{gj}$. The Gentzkow and Shapiro "leave-out isolation index" measures the extent to which members of group g disproportionately review venues whose other reviewers are also members of group g:

$$\hat{S}_g = \sum_j \frac{v_{gj}}{v_g} \cdot \left(\frac{v_{gj} - 1}{v_j - 1}\right) - \sum_j \frac{v_{\neg gj}}{v_{\neg g}} \cdot \left(\frac{v_{gj}}{v_j - 1}\right).$$

To generate a model-predicted value of $\hat{S}_g$ that is comparable to that in the data, we simulate model-predicted visits to restaurants for the observations in the estimation sample. Our estimated model predicts that each user will visit a venue with a probability given by equation (7). One draw for each observation from this probability distribution generates one simulated value of $\hat{S}_g$. We simulate the model 500 times to obtain a distribution of $\hat{S}_g$ values. The value observed in the estimation sample and the 90 percent confidence interval for simulated values are presented in table 3.

The actual data exhibit values of $\hat{S}_g$ within the 90 percent simulated confidence intervals. Appendix D shows that allowing preference parameters to vary across races is key to matching the observed consumption segregation: more restrictive specifications that pool preference parameters across races underpredict the isolation indices observed in the data.

TABLE 3
Model Fit: Isolation Indices

	Estimation Sample	Model Predictions
Asian isolation index	.087	[.054, .088]
Black isolation index	.087	[.041, .092]
White/Hispanic isolation index	.045	[.023, .055]

Note.—The reported leave-out isolation indices are the value for the estimation sample and the 90 percent confidence interval for model-predicted outcomes from 500 generated samples of the same size. Isolation indices are as defined in Gentzkow and Shapiro (2011).

2. Schelling-Style Segregation

There is an additional concern to address regarding our measure of consumption segregation. We cannot observe the complete racial and ethnic composition of the patrons of every restaurant in NYC. Accordingly, our baseline specification assumes that consumer preferences do not depend on this restaurant characteristic. Thus, our baseline model predicts that two restaurants with the same observable characteristics (tract, cuisine, price, and rating) will exhibit the same racial composition of patrons. The work of Schelling (1969, 1971) gives reasons why this may not be the case. He develops models in which neighborhoods may tip to extreme segregation even if the typical city resident prefers a much less segregated neighborhood. Card, Mas, and Rothstein (2008) provide evidence of such tipping in US residential patterns, and Zhang (2011) emphasizes the theoretical robustness of such predictions. The correlate concern in the present context is that there may be high degrees of segregation among restaurants with the same observables (tract, cuisine, price, and rating) if there is endogenous racial sorting due to preferences for same-race copatrons.

To examine the plausibility of our assumption that restaurants with the same observables will exhibit the same racial composition of patrons, we collect information on the racial composition of all Yelp reviewers for 125 pairs of restaurants that are identical in terms of their cuisine category, price category, Yelp rating, and census tract.[23] Define the "race gap" within a pair of restaurants p to be the Euclidean distance

$$\text{gap}_p \equiv \| \text{share}_j - \text{share}_{j'} \| / \sqrt{2},$$

with share_j being a three-element vector of the fraction of users reviewing restaurant j who are Asian, black, and Hispanic/white.[24] We compare the observed distribution of gap_p for the 125 pairs of observationally

[23] See app. D.2 for details. For reasons of feasibility, we restrict our attention to restaurants that belonged to a tract-cuisine-price-rating pair and had between 10 and 40 reviews.

[24] We drop reviewers whose race is not determined on the basis of their photos from these computations.

equivalent restaurants to the one that would arise if, consistent with our model, individuals were randomly assigned to one of the two restaurants within each pair.

Figure 6 depicts the distribution of gap_p for both the data and the random draws. The mean of the race gap for the observed data is 0.19. The mean for the random distribution is 0.175. The p-value for the one-sided test of equal means is .074. Appendix D.2 reports a similar result for differences in pairs of restaurants' contributions to the Gentzkow and Shapiro (2011) isolation index.

If consumption were segregated within sets of restaurants that our model treats as observationally equivalent, one might worry that our model would underpredict the true degree of consumption segregation. Our examination suggests that this is not the case. Conditional on the observable covariates that we employ to predict consumption segregation, Yelp reviews do not exhibit much further racial segregation.

D. *Parametric Bootstrap*

To examine the finite-sample behavior of our estimator, we perform a parametric bootstrap. Using the estimated model reported in columns 4–6 of table 2, we simulate 500 samples of the same size as our estimation sample. We then estimate our model on each of these generated sam-

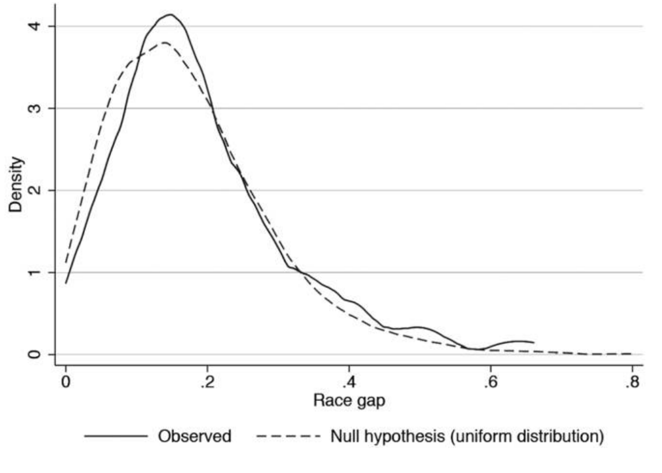

FIG. 6.—Racial gap between pairs of observationally equivalent restaurants. These kernel densities depict the distribution of the Euclidean distance between two restaurants' shares of patrons belonging to three racial categories for 125 pairs of restaurants that are identical in terms of their cuisine category, price category, Yelp rating, and census tract. The null hypothesis, in line with our model, is that individuals are randomly assigned to one of the two restaurants within each pair.

ples, obtaining a distribution of our estimator under the estimated data-generating process. Appendix D.5 reports the results in detail.

Generally, the parametric bootstrap shows that the estimator performs well. For the covariates describing social frictions and restaurant characteristics, the bootstrapped distributions are close to normal, their means are close to the point estimates we computed on the original sample, and their standard deviations are very similar to our estimated (asymptotically valid) standard errors. For the spatial-friction covariates, the bootstrapped samples occasionally produce estimates that are extreme outliers. The reason seems to be that we identify these six spatial-friction parameters exclusively from restaurant-reviewing outcomes $d^*_{ij} = \Sigma_l d^*_{ijl}$, without actually observing the origin-mode-level outcomes d^*_{ijl}, and that transit times from the same origin are highly collinear.[25] When we assume that the error term ν_{ijlt} and the disutility of travel do not vary across origin-modes l—implying that there is a single spatial-friction parameter to estimate and that consumers visit each restaurant via the origin-mode pair with the minimum travel time (see app. C.3)—the standard error we compute for our estimator of the disutility caused by this spatial friction is very similar to the bootstrapped one.

When we apply this bootstrap procedure to the isolation indices predicted by our estimated parameters, we find that they fit those associated with the estimated data-generating process. Table D4 shows that the confidence intervals for isolation indices predicted by the average of the bootstrapped parameters are very similar to those reported in table 3. Figure D6 shows that the distributions of the endpoints of these 90 percent confidence intervals are nearly centered around the data-generating process's values.

E. *Robustness Checks*

1. Restaurant Fixed Effects

It is feasible to introduce race-specific restaurant fixed effects into our model. However, as we cannot identify these fixed effects for restaurants that are not visited by reviewers in the estimation sample, this generalized model cannot be used to compute citywide measures of consumption segregation and counterfactuals.[26] We therefore employ the specification with race-specific restaurant fixed effects only to examine the

[25] For example, the correlation between travel time from work by car and from work by public transit exceeds .9 for all three racial groups.

[26] Estimating models with restaurant fixed effects is computationally costly: estimation takes days rather than minutes. This is true in our setting whether we estimate our fixed effects directly or assume the approximation in Taddy (2015). We report the results of Taddy's estimation procedure in app. D.3.

robustness of the coefficients on observable reviewer-restaurant covariates reported in table 2 and to assess the relative fit of these specifications.

The results suggest that our baseline specification in table 2 is sufficient to capture the relevant variation in consumers' choices. Table 4 reports the result of estimating the specification with restaurant fixed effects. The estimated coefficients on our measures of spatial and social frictions are similar to those reported in table 2. Table 5 reports the result of a likelihood ratio test comparing the fit of the restaurant-fixed-effects specification to the specifications in table 2 that use only observable characteristics. For

TABLE 4
Restaurant Fixed Effects

	Asian (1)	Black (2)	White/Hispanic (3)
Log travel time from home by public transit	−1.12***	−1.07***	−1.24***
	(.111)	(.141)	(.061)
Log travel time from home by car	−1.23***	−1.28***	−1.49***
	(.092)	(.153)	(.062)
Log travel time from work by public transit	−1.29***	−1.87**	−1.82***
	(.141)	(.803)	(.210)
Log travel time from work by car	−1.78***	−1.92***	−2.10***
	(.201)	(.420)	(.176)
Log travel time from commute by public transit	−1.04***	−1.03***	−1.17***
	(.073)	(.106)	(.041)
Log travel time from commute by car	−1.15***	−1.40***	−1.58***
	(.066)	(.158)	(.061)
EDD between h_i and k_j	−.796***	−2.40***	−1.15***
	(.133)	(.324)	(.146)
EDD × SSI	−.527***	−.533	−.003
	(.182)	(.433)	(.087)
$$ bin × home tract median income	.040***	−.003	.050***
	(.011)	(.034)	(.010)
$$$ bin × home tract median income	.080***	.052	.076***
	(.014)	(.056)	(.013)
$$$$ bin × home tract median income	.062***	−.200	.082***
	(.022)	(.259)	(.023)
Yelp rating × home tract median income	.020**	.012	.025***
	(.010)	(.027)	(.009)
Percent absolute difference in median incomes ($h_i - k_j$)	−.200***	1.54***	−.189***
	(.057)	(.181)	(.061)
Percent difference in median incomes ($k_j - h_i$)	.053	.505	.182
	(.362)	(.995)	(.383)
Number of trips	6,447	1,079	6,936
Number of fixed effects	2,867	892	3,497

Note.—Each column reports an estimated conditional-logit model of individuals' decisions to visit a Yelp venue. Standard errors are in parentheses. The unreported covariates are restaurant fixed effects.

* Statistically significant at 10 percent.

** Statistically significant at 5 percent.

*** Statistically significant at 1 percent.

TABLE 5
Observables versus Restaurant Fixed Effects

Sample	Log Likelihood Values: Observables	Log Likelihood Values: Restaurant Fixed Effects	χ^2 Test Statistic	p-Value
Asian reviewers	−55,257.58	−47,449.99	15,615.19***	.00
Black reviewers	−9,185.82	−6,759.39	4,852.87	1.00
White/Hispanic reviewers	−59,121.74	−51,591.93	15,059.63***	.00

Note.—The observables specifications include 60 covariates, while the fixed-effect specifications include 14 covariates plus the 10,945 restaurant fixed effects. Thus, the χ^2 test has 10,899 degrees of freedom.
* Statistically significant at 10 percent.
** Statistically significant at 5 percent.
*** Statistically significant at 1 percent.

black reviewers, we do not reject the hypothesis that the observables specification fits the data as well as the fixed-effects specification. For Asian and white/Hispanic reviewers, the fixed-effects specification is superior in terms of fit, but the estimated coefficients are largely consistent with those in table 2. The fixed-effects specification exhibits worse fit in terms of isolation indices, predicting values greater than the estimation-sample values reported in table 3.[27]

2. Nested-Logit Specification

We relax the independence of irrelevant alternatives property of the conditional-logit model of Section III.A by specifying a nested-logit structure. Appendix C.4 derives this estimator in detail. We define nests by two schemes: (*a*) restaurants of the same cuisine category, Yelp rating, and area and (*b*) restaurants of the same cuisine category, price category, and census tract. For this exercise, we define 39 cuisine categories (more disaggregated than the nine categories shown in table 2) and employ the 28 areas, four price categories, and nine Yelp ratings (from one to five stars) described in Section II. These two schemes group the 10,945 restaurants into 3,064 and 7,622 nests, respectively. Table D3 reports the estimates. The estimated coefficients on spatial frictions, social frictions, and restaurant characteristics are all similar to the values reported in columns 4–6 of table 2. The within-nest-correlation parameter λ is generally near one, consistent with the conditional-logit assumption we imposed in Section III.A.

[27] In finite samples, our estimates of the restaurant fixed effects partly reflect the frequency with which a restaurant is randomly sampled in choice sets. Even if all restaurants are sampled with the probability given by eq. (8), uneven sampling in finite realizations will affect the estimates of the fixed effects regardless of consumers' preferences. This may worsen the fit of the isolation index. As noted in n. 16, this does not threaten our estimates of the coefficients on observable covariates.

Table D3 reports the p-values for likelihood ratio tests of the null hypothesis that the true model is a conditional logit versus the alternative hypothesis that the true model is the corresponding nested-logit model. While the tests formally reject the hypothesis that $\lambda = 1$ for three of the six specifications, these nested-logit specifications yield very similar coefficients on observable covariates and deliver very similar predictions of the in-sample isolation measures (see table D1). Since estimating these nested-logit models comes at considerably greater computational cost, we employ the conditional-logit specification when computing consumption segregation and counterfactual outcomes.

3. Additional Observable Characteristics and Sample Restrictions

In appendix A, we implement several robustness checks that restrict the set of observations or introduce additional covariates to address concerns discussed in Section III.D. These are reported in tables A4–A6. We also report these robustness checks for the specification in which consumers visit each restaurant via the origin-mode pair with the minimum travel time in tables A7–A9. We briefly summarize the robustness checks here.

To address concerns that the error term ν_{ijlt} may exhibit serial correlation, we restrict the sample to either the first half or the first fifth of reviews written by each user. These restrictions increase the standard errors and cause one coefficient to be unidentified in the resulting small sample of black reviewers. However, they do not systematically increase the absolute values of coefficients, suggesting that there is not substantial attenuation bias caused by serial correlation in the unobserved preference shocks ν_{ijlt}. For a more detailed discussion, see appendix C.8.

To address concerns related to how we located reviewers, we show that both dropping the 5 percent of restaurant reviews used to locate reviewers and splitting the sample on the basis of the number of Yelp reviews revealing a reviewer's home location alter the coefficients of interest little. To address concerns that early adopters of Yelp may be less sensitive to spatial and social frictions, we restrict the sample to reviewers who joined Yelp later. The estimated coefficients are broadly similar to those in table 2, without any systematic increase or decrease. Controlling for 39 cuisine categories (rather than nine) slightly attenuates our point estimates of homophily and makes the environmental similarity coefficients slightly more negative. We find that our results change little when controlling for tract-level differences in private vehicle ownership rates.

To address concerns that the decision to write a review may depend on certain restaurant characteristics, we control for two additional covariates: the restaurant's number of Yelp reviews and a dummy indicating whether it belongs to a restaurant chain. Introducing these covariates substantially alters the coefficients on the restaurant's rating and price, but the coefficients on spatial and social frictions change little.

To address concerns that users may be more or less likely to review restaurants visited from a particular origin, table A12 introduces origin-mode-specific intercepts. This yields broadly similar coefficients on log travel times, though the accompanying standard errors are considerably larger. The coefficients on social frictions are modestly attenuated.

In summary, the results in table 2 are broadly unchanged by restricting the estimation sample or introducing additional covariates.

V. Consumption Segregation

In this section, we use data on the demographic composition of all census tracts in NYC and the estimates presented in columns 4–6 of table 2 to compute NYC-wide measures of segregation in consumption for Asian, black, Hispanic, and white consumers. In Section V.A, we define the measure of segregation we use. Section V.B presents our estimates of consumption segregation and examines the contributions of spatial and social frictions to them. In Section V.C, we illustrate the mechanisms underlying the citywide results by focusing on the consumption patterns observed in particular neighborhoods within the city. Finally, in Section V.D we examine a number of counterfactual experiments.

A. *Dissimilarity Indices*

To measure consumption segregation, we use the "dissimilarity index" commonly employed in the literature on residential segregation. For each group $\mathfrak{g}$, we compute

$$\text{Dissimilarity}(\mathfrak{g}) = \frac{1}{2}\sum_{j\in J}\left|\Pr\left(d_{ij} = 1|\mathfrak{g}(i) = \mathfrak{g}\right) - \Pr\left(d_{ij} = 1|\mathfrak{g}(i) \neq \mathfrak{g}\right)\right|, \tag{13}$$

where $\mathfrak{g}$ is Asian, black, Hispanic, white, or other.[28] This index sums, across all restaurants, the absolute difference between the probability that a randomly selected individual belonging to group $\mathfrak{g}$ visits a restaurant and the probability that a randomly selected individual who does not belong to group $\mathfrak{g}$ visits the same restaurant. The higher the value of this index, the larger the differences in consumption choices. Think of the probability $\Pr(d_{ij} = 1|\mathfrak{g}(i) = \mathfrak{g})$ as the fraction of individuals of

[28] As discussed in Sec. III, we estimate parameters for three racial groups indexed by g. However, we compute dissimilarity indices for five racial groups indexed by $\mathfrak{g}$. We compute $\Pr(d_{ij} = 1|\mathfrak{g}(i) = \mathfrak{g})$ for Hispanic, white, and other using $\mathfrak{g}$-specific data on residential locations and the estimated white/Hispanic ($g = w$) preference parameters.

group $\mathfrak{g}$ that would visit restaurant j if they were all going out to dine; the dissimilarity index in equation (13) indicates the share of individuals in group $\mathfrak{g}$ that would need to alter their consumption choices in order to match the distribution of predicted restaurant choices made by the remainder of the population. A pairwise dissimilarity measure can also be computed to compare the consumption choices of any two groups, $\mathfrak{g}_1$ and $\mathfrak{g}_2$:

$$\text{Dissimilarity}(\mathfrak{g}_1, \mathfrak{g}_2) = \frac{1}{2}\sum_{j\in J}\left|P\left(d_{ij} = 1|\mathfrak{g}(i) = \mathfrak{g}_1\right) - P(d_{ij} = 1|\mathfrak{g}(i) = \mathfrak{g}_2)\right|.$$

These dissimilarity indices are invariant to the size of the groups being compared.

If we were to observe all visits to restaurants for a sufficiently large representative sample of NYC residents, we could estimate $\Pr(d_{ij} = 1|\mathfrak{g}(i) = \mathfrak{g})$ directly by its sample analogue. To our knowledge, such a large and representative data set does not exist.[29] We therefore apply the parameters estimated in Section IV to the broader population of NYC in order to consistently estimate the probability $\Pr(d_{ij} = 1|\mathfrak{g}(i) = \mathfrak{g})$ for all restaurants j and races $\mathfrak{g}$. Appendix E details how we construct consistent estimates of these probabilities. A further advantage of relying on the demand model described in Section III as the basis for computing our measures of consumption segregation is that it can also be used to quantify the contributions of spatial and social frictions to our overall measure of consumption dissimilarity.

B. Citywide Consumption Segregation

Table 6 reports residential and consumption dissimilarity indices for each demographic group. Panel A reports overall dissimilarity indices for each group; panel B reports pairwise dissimilarity indices. We compute residential dissimilarity at the level of census tracts and consumption dissimilarity at the level of restaurant venues.[30] We employ the parametric bootstrap

[29] As explained by Gentzkow et al. (2019), computing consistent measures of segregation can be difficult in a sample that is small relative to the dimensionality of the choice set faced by individuals. In our context, this sample would have to be large relative to the number of restaurants in NYC. The advantage of the behavioral model introduced in Sec. IV is that it expresses the probabilities $\Pr(d_{ij} = 1|\mathfrak{g}(i) = \mathfrak{g})$ as a function of a relatively small number of estimated parameters.

[30] The choice of spatial unit is nontrivial when computing dissimilarity indices, as Echenique and Fryer (2007) stress. Restaurants are a natural spatial unit within which interactions may occur, while tract-level results may be sensitive to how the Census Bureau chose to partition the city. To facilitate comparison of measures of residential and consumption segregation, table A14 reports consumption dissimilarity indices computed at the level of census tracts. The resulting indices are broadly similar to those computed at the level of restaurant venues.

TABLE 6
Residential and Consumption Segregation

	Residential Dissimilarity (1)	Consumption Dissimilarity: Estimated (2)	No Spatial (3)	No Social (4)	Neither Friction (5)
		A. Dissimilarity Index			
Asian	.521	.315	.290	.245	.232
		[.305, .335]	[.280, .314]	[.233, .268]	[.222, .259]
Black	.653	.352	.322	.273	.260
		[.337, .397]	[.307, .372]	[.258, .320]	[.248, .309]
Hispanic	.486	.142	.114	.106	.088
		[.134, .162]	[.108, .137]	[.099, .125]	[.083, .109]
White	.636	.190	.153	.112	.093
		[.180, .209]	[.143, .174]	[.106, .130]	[.090, .112]
White or Hispanic	.470	.205	.189	.150	.156
		[.197, .236]	[.182, .224]	[.143, .182]	[.149, .191]
		B. Pairwise Dissimilarity			
Asian-black	.796	.495	.448	.388	.357
		[.480, .534]	[.429, .491]	[.370, .429]	[.340, .402]
Asian-Hispanic	.584	.288	.273	.220	.217
		[.277, .310]	[.262, .299]	[.208, .246]	[.206, .247]
Asian-white	.519	.278	.255	.212	.203
		[.268, .298]	[.245, .279]	[.200, .236]	[.193, .233]
Black-Hispanic	.558	.328	.297	.261	.250
		[.312, .372]	[.284, .348]	[.246, .308]	[.238, .300]
Black-white	.822	.354	.324	.263	.255
		[.337, .401]	[.309, .375]	[.249, .310]	[.243, .306]
Hispanic-white	.658	.159	.115	.095	.037
		[.144, .176]	[.098, .135]	[.085, .104]	[.028, .047]

Note.—This table reports dissimilarity indices. Panel A reports the index for each demographic group's residential/consumption locations compared to members of all other demographic groups. Panel B reports the index for residential/consumption locations between each pair of demographic groups. The demographic group "other" is included in computations but not reported. Col. 1 reports indices based on tracts' residential populations. The remaining columns report venue-level dissimilarity indices based on the coefficient estimates in cols. 4–6 of table 2. Col. 2 uses the estimated coefficients. Col. 3 sets the coefficients on travel time covariates to zero. Col. 4 sets the coefficients on demographic-difference covariates to zero. Col. 5 sets the coefficients on travel time and demographic difference covariates to zero. Bootstrapped 95 percent confidence intervals from 496 draws are reported in brackets.

distributions of our estimator from Section IV.D to produce 95 percent confidence intervals for the consumption dissimilarity estimates.

Column 1 provides the dissimilarity index for residential segregation while column 2 provides the analogous index for consumption segregation. Comparing these columns, all groups are significantly more integrated in their consumption than in their residences. The ratio of residential dissimilarity to consumption dissimilarity is 3.4 for Hispanics, 3.4 for whites, 1.9 for blacks, and 1.7 for Asians. NYC's levels of residential dissimilarity are similar to the nationwide average level of dissimilarity for black residents in 2010, while the levels of dissimilarity in consumption implied by our estimates are comparable to the levels of black residential dissimilarity observed in America's most integrated metropolitan areas (Glaeser and Vigdor 2012). The 20–45 percentage point difference between consumption segregation and residential segregation of NYC residents is one to two times the largest declines in black residential dissimilarity observed from 1970 to 2010 across US metropolitan areas (Glaeser and Vigdor 2012). At the median historical rate of decline, residential segregation would have to continue its decline for nearly a century to reach levels comparable to our estimated levels of consumption segregation.

The fact that all groups are significantly more integrated in their consumption than in their residences is not a necessary consequence of the model assumptions imposed in Section III. While the willingness of residents to travel outside of their home census tract to consume may tend to reduce consumption segregation relative to residential segregation, demographic differences in cuisine tastes or demographically linked social frictions could cause consumption segregation to exceed residential segregation. Our numbers show that social frictions and heterogeneity in tastes do not overturn the integrating effect of consumers' mobility.

Across demographic groups, black and Asian individuals exhibit the highest values of consumption dissimilarity, but in part because we assign white and Hispanic consumers the same preference parameters as a result of our inability to differentiate between white and Hispanic Yelp reviewers in their photos. To the extent that white and Hispanic consumers differ in their preferences, we will underestimate consumption segregation for these two groups. In panel B, the largest pairwise dissimilarities are found between Asian and black consumers. Black and white consumers' choices are also dissimilar, while Hispanic-white and black-Hispanic consumption choices are more integrated.[31]

In order to measure the contributions of spatial and social frictions to consumption segregation, we again use the estimates in columns 4–6 of table 2 and recompute the dissimilarity indices that arise from setting some of the estimated coefficients to zero. This calculation holds fixed both the set of restaurants and their characteristics and, thus, should not be interpreted as capturing the total effect of eliminating spatial or social frictions (which would likely generate supply responses). In computing the dissimilarity indices reported in column 3 of table 6, the coefficients on travel time covariates are set to zero, eliminating the role

[31] Despite their common preference parameters, white and Hispanic consumers still exhibit pairwise dissimilarity due to observable differences in residential locations that generate differences in the social friction covariates employing tract-level racial and ethnic demographics (e.g., EDD) and other covariates employing tract-level characteristics (e.g., the interaction of price dummies and median household income in the reviewer's home tract).

of spatial frictions.[32] In column 4, the coefficients on the demographic-differences covariates are set to zero, eliminating the role of social frictions.[33] In column 5, both these sets of coefficients are set to zero, so that consumers in different groups exhibit different predicted consumption behavior due only to differences in their residential income levels and income-linked valuations of venues' prices and ratings; race-specific valuations of restaurants' cuisines, prices, and ratings; race-specific responses to robberies per resident; and race-specific area fixed effects.[34]

Comparing columns 2 and 3, the elimination of spatial frictions has a mild integrating effect, causing consumption dissimilarity to fall by an average of 3 percentage points. That is, individuals from different demographic groups value restaurant destinations sufficiently similarly that eliminating spatial frictions would, all else equal, make their choices more integrated. Because of spatial frictions, consumption segregation at least partly inherits the pattern of residential segregation. Comparing columns 2 and 4, eliminating the roles of environmental dissimilarity and homophily (i.e., social frictions) reduces consumption dissimilarity by an average of 6.6 percentage points, or more than twice the effect of eliminating spatial frictions. If consumer behavior did not respond to differences between the residential demographics of the restaurant tract and both the consumer's individual identity and home tract demographics, predicted consumption behavior would be much more integrated.[35] Eliminating both spatial and social frictions would reduce consumption dissimilarity by about one-third on average.

The relative contributions of spatial and social frictions are consistent across demographic groups. For each group, social frictions make a notably greater contribution to the observed level of consumption dissimilarity than spatial frictions. Moreover, the relative overall levels of estimated consumption dissimilarity appear to reflect dissimilarity attributable to demographic differences in tastes. In the absence of both spatial and social frictions, black consumers would exhibit the greatest consumption dissimilarity, and Hispanic consumers the least, just as they do in the estimated levels in column 2.

The pairwise dissimilarity indices reported in panel B of table 6 reflect the rich set of covariates incorporated in our behavioral model. For ex-

[32] This is a ceteris paribus exercise. In reality, spatial frictions and social frictions may not be entirely independent. For example, if social frictions reflect segregated friendship networks and users visiting restaurants near their friends' residences, then the elimination of spatial frictions would eliminate this component of social frictions.

[33] Specifically, we set to zero the coefficients on all the EDD, SSI, and share covariates reported in table 2.

[34] For example, Asian reviewers are, all else equal, more likely to visit Asian and Indian restaurants, and reviewers residing in higher-income tracts are more likely to visit restaurants with higher prices.

[35] This decomposition holds residential segregation fixed. However, one could expect that residential segregation would be different in the absence of social frictions.

ample, while Asian-black and black-white residential dissimilarity indices are similar, consumption dissimilarity is notably greater between Asian and black consumers than between black and white consumers. This partly reflects the magnitude of the Asian-black interactions in table 2. Column 5 of table 6 reveals that it also reflects divergent choices due to differences in income levels and tastes. The Asian-black pairwise dissimilarity results, which are not affected by the fact that we pool preference parameters for white and Hispanic reviewers, match the findings in panel A of table 6: Asian-black residential dissimilarity is much greater than Asian-black consumption dissimilarity, and this consumption dissimilarity reflects social frictions more than spatial frictions.

These results are robust to a number of the alternative estimating assumptions discussed in Section IV.E.3. Table A15 shows that we obtain similar results when using parameter estimates obtained from our nested-logit specifications, limiting the sample to reviewers' early reviews, splitting the sample on the basis of the amount of locational information used to identify a reviewer's residence, restricting the sample to late adopters, controlling for vehicle ownership rates, controlling for chain establishments and the restaurant's number of Yelp reviews, using more disaggregated cuisine categories, and constraining all trips to start at home. We also find very similar results when computing dissimilarity using a minimum–transit time specification (see app. C.3), as reported in table A16.

C. *Illustrative Examples*

We illustrate the mechanisms behind the results in table 6 in two parts of the city. First, we examine three neighborhoods in Manhattan: the Upper East Side, Central Harlem, and East (Spanish) Harlem (respectively, Manhattan community districts 8, 10, and 11). The fact that each of these neighborhoods is residentially segregated, with a distinct demographic majority, makes the general process at work easy to visualize in figure 7.[36]

Panel A in figure 7 captures residential segregation. Each dot represents 5 percent of the tract population. The Upper East Side of Manhattan stretches from Fifty-Ninth Street to Ninety-Sixth Street, Central Park to the East River. While NYC is only 33 percent white, the Upper East Side is 81 percent white. If we restrict attention to the 14 census tracts between Third Avenue and Central Park, the median tract is 92 percent white. Among nearly 61,000 residents of these tracts, only 726 were black. In short, this is a highly segregated area of the city.

Central Harlem comprises Fifth Avenue to Eighth Avenue, Central Park North (110th Street) to the Harlem River. This is the storied center

[36] Table A17 reports shares of residents and predicted consumers by race for the three community districts depicted in these maps.

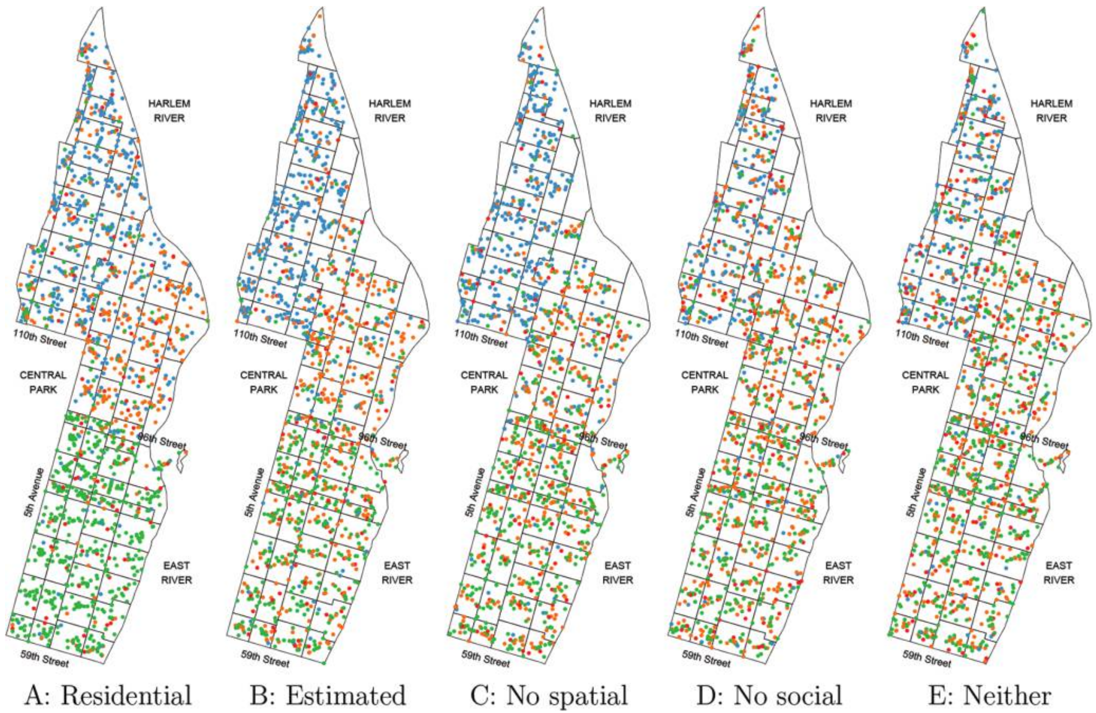

Fig. 7.—Residential and consumption segregation in three Manhattan communities. These maps depict community districts 8 (Upper East Side), 10 (Central Harlem), and 11 (East Harlem) in Manhattan. The five maps correspond to the five scenarios reported in table 6. Each dot represents 5 percent of the tract's residential population or predicted restaurant visitors. Asian residents or consumers are represented by red dots, black by blue dots, Hispanic by orange dots, and white by green dots. In different tracts, dots represent a different number of people, so the maps depict variation in shares, not levels.

of urban black America. While its black population share has fallen from its 1990 level of 88 percent, it remains 63 percent black. The next-largest group is Hispanics at 22 percent, with modest levels of whites (10 percent) and Asians (2 percent). While Central Harlem's residential population is becoming more diverse, panel A of figure 7 makes it evident that it remains a highly segregated area.

East (Spanish) Harlem stretches from Ninety-Sixth Street to the Harlem River, and Fifth Avenue to the East River. The Hispanic fraction of the population has remained roughly constant in the last 20 years at about 50 percent. There is a large, even if declining, black population in East Harlem, at roughly 30 percent, located most densely where East Harlem abuts Central Harlem to the west and in the more northerly areas of the district. Asians and whites are present in small but growing numbers.

Panel B of figure 7 shows the degree of consumption segregation within these areas. Each dot represents 5 percent of the visits to that tract. Two features jump out from this panel. The first is that predicted consumption in panel B is strikingly less segregated than residences in panel A. This is consistent with a comparison of columns 1 and 2 of table 6. The boundaries between black and Hispanic consumers are porous, consistent with the black-Hispanic interactions in table 2 and the pairwise dissimilarity index in table 6. Asian consumers are more prevalent in the Upper East Side, for example, than Central Harlem, reflecting both a shorter distance to Asian residential population centers and smaller social frictions between Asians and whites than between Asians and the two other groups. The second feature is that, nonetheless, there remains a very high level of segregation. As summarized in table A17, black consumers dominate consumption in Central Harlem, Hispanics in East Harlem, and whites in the Upper East Side. Segregation of consumption is much less than residences but still strong.

The following three panels in figure 7 follow columns 3–5 of table 6 by illustrating the degree of consumption segregation for the respective cases in which spatial, social, or both types of frictions are set to zero when constructing predicted consumption patterns. Panel C is based on our estimates in which we wholly eliminate spatial frictions, effectively making every restaurant in the city instantly available to any resident of the city. Comparing panel C to panel B, there is a diminution of the degree of consumption segregation. Yet the change seems modest, consistent with a comparison of columns 2 and 3 of table 6.

When we move to panel D of figure 7, we allow spatial frictions to again be at their estimated level but now set social frictions to zero. Visually, comparing panels B and D in figure 7, we see a large decline in the degree of consumption segregation in each of these neighborhoods. It is important to recognize that this did not need to be true. Residential segregation plus spatial frictions to consumption could have been enough to maintain very high levels of consumption segregation; as we have observed, they just happen not to do so. Social frictions matter a great deal for consumption integration, consistent with column 4 of table 6.

Finally, we consider the elimination of both spatial and social frictions by setting both sets of coefficients to zero. Column 5 of table 6 tells us that there is some further modest decline in the degree of consumption segregation relative to the no-social-frictions case. However, it is sufficiently modest that it does not stand out clearly to the eye in panel E of figure 7 in any of the three neighborhoods. This again suggests that spatial frictions explain a small share of consumption segregation.

Next, we examine another part of New York City with considerable, but less extreme, segregation. Figure 8 depicts neighborhoods near the portion of the East River separating the Lower East Side of Manhattan from Brooklyn.[37] Panel A shows strong patterns of residential segregation. Manhattan community district 3 includes Chinatown (with predominantly Asian residents), the East Village (with predominantly white residents in the northern and western portions), and the Lower East Side (with predominantly Hispanic residents). Across the river, Brooklyn community district 1 is home to concentrations of white residents in Greenpoint and Williamsburg and concentrations of Hispanic residents, especially in the areas leading out to (mostly Hispanic) Bushwick. Brooklyn community district 2 includes the mostly white Brooklyn Heights as well as Fort Greene, with a mixture of black, Hispanic, and white residents. Finally, Bedford-Stuyvesant (Brooklyn community district 3) is a traditionally black area that now has white residents in the area near Fort Greene and Hispanic residents in areas proximate to Williamsburg and Bushwick.

Panel B of figure 8 shows that estimated consumption is strikingly more integrated. Chinatown shows considerable inflows of white and Hispanic consumers, who are proximate residents, but more modest inflows of black consumers, who are more remote. The East Village and Lower East Side are also notably more integrated, again with only modest numbers of black consumers. Greenpoint and Williamsburg host predominantly white consumers who are augmented by Hispanic consumers residing, presumably, in Williamsburg and Bushwick. Brooklyn Heights and Fort Greene are similarly host to mostly white consumers augmented by black and Hispanic consumers. Finally, consumption in Bedford-Stuyvesant is notably more racially integrated than its residences.

Panels C, D, and E of figure 8 reaffirm our finding that consumption segregation is driven more by social frictions than spatial frictions. To the eye, panel C, which sets spatial frictions to zero, is nearly identical to panel B. For example, the areas of Chinatown and Bedford-Stuyvesant continue to be dominated by Asian and black consumers, respectively. By contrast,

[37] Table A18 reports shares of residents and predicted consumers by race for the four community districts depicted in these maps.

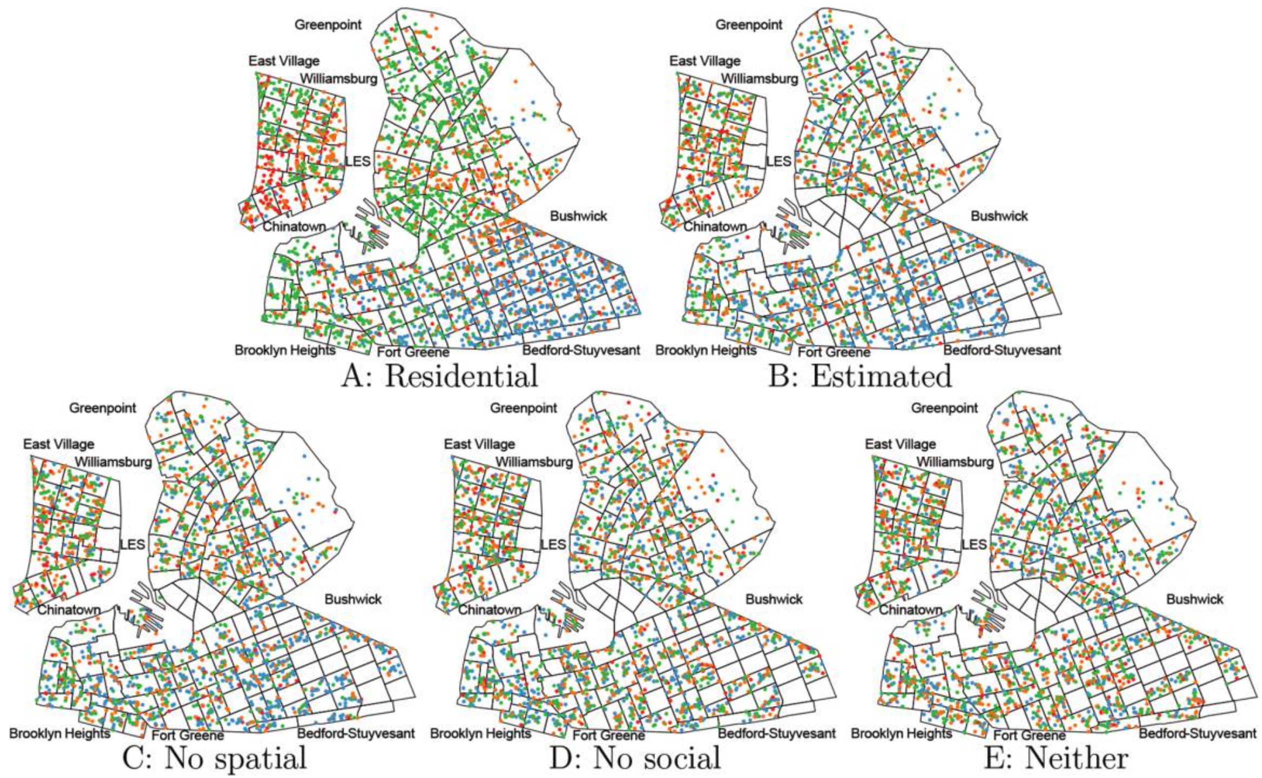

Fig. 8.—Residential and consumption segregation in Lower Manhattan and West Brooklyn. These maps depict community district 3 in Manhattan and 1, 2, 3, and 4 in Brooklyn. See the notes to fig. 7.

panel D, which sets social frictions to zero, noticeably differs from panel B. The consumption enclaves of Chinatown and Bedford-Stuyvesant do not entirely disappear, but they are much more integrated. Indeed, setting both frictions to zero, as in panel E, does little to integrate consumption more visibly than the removal of social frictions alone.

D. Counterfactuals

We next study how changes in transportation infrastructure and in the level of social frictions may affect the level of consumption segregation we observe.

1. Transportation Policy and Technology

Table 6 shows that the complete elimination of spatial frictions would only modestly reduce consumption dissimilarity indices. This bounds the conceivable impact of driverless cars or other "frictionless" technologies on consumption segregation. How might more immediately feasible transportation projects affect consumption segregation? We consider two counterfactuals relevant to NYC policy makers. First, we forecast the effect of the new Second Avenue subway on consumption segregation. Second, we look at the effects of a general slowdown in NYC transit.

The Second Avenue subway is an ongoing multi–billion dollar expansion of the NYC subway system. When completed, the line will stretch from 125th Street in East Harlem all the way down the East Side to Hanover Square in the Financial District. We compute counterfactual travel times for this transportation infrastructure improvement and forecast the effect on consumption segregation.[38] The results are reported in table 7. The Second Avenue subway line has almost no effect on consumption segregation. This is the joint effect of the fact that spatial frictions play a relatively small role in determining consumption segregation and that the Second Avenue line is inferred to have relatively small effects on travel times for the majority of the city's residents.

Second, we study consumption segregation when automobiles and public transit are 20 percent slower. In 2014, NYC lowered the speed limit within the city from 30 miles per hour to 25 miles per hour (Bankoff 2014). Its subway speeds are also down about 20 percent in the last few years (Rosenthal, Fitzsimmons, and LaForgia 2017). The effects on consumption segregation are all small in magnitude. Our conclusion is that

[38] Appendix E.2 details how we construct the counterfactual transit times. Our computation captures only the value of new subway connections in the network graph and does not assign any value to benefits such as alleviating overcrowding, which is a major motivation for the Second Avenue project.

TABLE 7
Counterfactual Consumption Dissimilarity

	Estimated (1)	2nd Ave. (2)	Slowdown (3)	Social Change (4)
	A. Dissimilarity Index			
Asian	.315	.315	.318	.300
Black	.352	.352	.354	.330
Hispanic	.142	.142	.145	.133
White	.190	.190	.194	.170
White or Hispanic	.205	.206	.208	.189
	B. Pairwise Dissimilarity			
Asian-black	.495	.494	.498	.469
Asian-Hispanic	.288	.288	.291	.274
Asian-white	.278	.278	.281	.264
Black-Hispanic	.328	.328	.330	.309
Black-white	.354	.354	.357	.328
Hispanic-white	.159	.159	.164	.144

Note.—This table reports venue-level dissimilarity indices based on the coefficient estimates in cols. 4–6 of table 2. Col. 1 uses the estimated coefficients. Col. 2 introduces the decrease in public transit times due to the Second Avenue subway. Col. 3 increases all travel times by 20 percent. Col. 4 reduces the (absolute value of the) coefficients on all social friction covariates by 20 percent.

even quite substantial interventions in the transport sphere are unlikely to have a major impact on the integration of consumption.

2. Social Frictions

Next we examine the effects of a decline in demographic-linked social frictions on consumption segregation. At first glance, this may appear to be an odd policy exercise, since social frictions likely reflect a variety of factors, such as tastes and social networks, that are not immediately under policy makers' control. Yet there are government policy initiatives at the federal, state, and municipal levels that aim to encourage understanding and prevent tensions between different demographic groups.[39]

The counterfactual that we examine is the reduction of the magnitude of social frictions by 20 percent.[40] The results are reported in column 4 of

[39] At the federal level, the Community Relations Service within the Department of Justice (https://www.justice.gov/crs) "is the Department's 'Peacemaker' for community conflicts and tensions arising from differences of race, color, national origin, gender, gender identity, sexual orientation, religion and disability." It "facilitates the development of viable, mutual understandings and solutions to the community's challenges." In New York City, the Commission on Human Rights (https://www1.nyc.gov/site/cchr/community/community.page) has a dual mandate, to educate the public about legal protections for a variety of groups and "encouraging understanding and respect among New York City's many communities."

[40] Our choice of this magnitude is admittedly arbitrary. Reducing social frictions by 20 percent might be an implausibly large change. On one side, 42 percent of Americans were

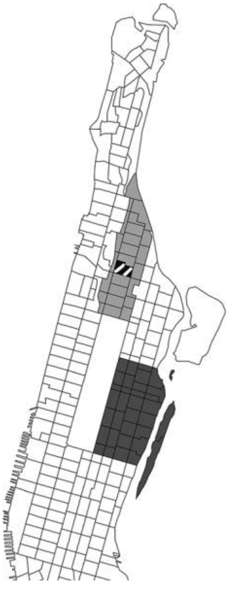

FIG. 9.—Harlem gentrification scenario. We compute the change in black residents' expected utility in the striped tract if the surrounding light gray tracts were to exhibit the characteristics of the dark gray tracts. Table 8 reports the changes in these characteristics.

table 7. The direct effect of reducing social frictions on consumption segregation is modest relative to the overall level of consumption segregation. The residual consumption segregation reflects taste differences that may well attenuate over time with reductions in other dimensions of segregation but are unlikely to disappear over any short or even medium horizon.

VI. Gentrification

Public concerns about gentrification are primarily about displacement of one group by another. Nominally this is an issue of social class, but in practice it also has a racial component: neighborhoods shifting composition from less affluent minorities toward more affluent whites. Since tastes for consumption venues differ by class and race, there is a second concern that the commercial mix may also shift away from the set of venues well

worried "a great deal" about race relations in 2017, a peak since Gallup began asking the question in 2001. On the other, race relations improved considerably over recent decades in a number of dimensions. To mention a single measure, approval of black-white marriage rose from 4 percent in 1958 to 87 percent in 2011. The realm of conceivable changes in social frictions is quite broad, and we are not aware of evidence disciplining the conceivable magnitude of policy-induced changes in social frictions.

TABLE 8
Harlem Gentrification Scenario

Change in	Mean	Standard Deviation
Asian residential share	.076	.057
Black residential share	−.565	.095
Hispanic residential share	−.112	.083
White residential share	.611	.173
Robberies per resident	−.005	.003
SSI	−1.177	.273
Yelp rating	.137	1.037
Price ($ to $$$$)	.667	.988
Median household income (thousands)	76.543	58.998
EDD	.541	.121
Number of restaurants	102	

Note.—The table reports the changes in characteristics for the gentrification scenario depicted in fig. 9.

suited to poor and minorities to those favored by affluent whites. This amounts to a concern about the potential for gentrification to erode the benefits of living in a neighborhood for those who stay behind. In New York City, the presence of rent-stabilized apartments and New York City Housing Authority buildings, with low turnover rates, makes this question salient.

We employ our estimates to quantify the welfare impact of gentrification for incumbent residents by examining the consequences of changes in resident and restaurant composition. The gentrification scenario we study is depicted in figure 9. We select one low-income, majority-black census tract in Harlem (the striped polygon) and compute the change in black residents' expected utility if the surrounding census tracts in Central Harlem (in light gray) were to exhibit the residential and restaurant characteristics of high-income, majority-white census tracts of the Upper East Side (in dark gray). In doing so, we hold the number of restaurants fixed while allowing their characteristics to change to those typical of the Upper East Side. The changes in restaurant and residential characteristics are summarized in table 8.

As a result of this gentrification, black residents of the unchanged Harlem census tract experience a decrease in the expected utility of patronizing restaurants. Stated in terms of spatial frictions, this welfare loss is equivalent to each gentrifying restaurant in Central Harlem becoming nearly four times as far away from the residents' homes by both car and public transit. This welfare loss can be decomposed using a simple approximation of the utility change that we derive in appendix E.3:

$$U_i' - U_i \approx \left(\sum_{j\in\mathcal{J}^G} P_{ij}\right) \times \left[\exp\left(\gamma_g^2\Delta\bar{X}_{ij}^2 + \beta_g^1\Delta\bar{Z}_j^1 + \beta_g^2\Delta\bar{Z}_{ij}^2\right) - 1\right],$$

TABLE 9
Welfare Losses Due to Gentrification of Surrounding Harlem Neighborhoods

Transit Time Increase Equal to Welfare Loss	Initial Visit Share	Change in Value of Characteristics ($\gamma\Delta\overline{X}_i$, $\Delta\overline{Z}_i$) Social Frictions	Restaurant Traits	Other Traits
291%	.101	−1.76	−.069	.692

Note.—Welfare loss is expressed as the percentage increase in transit times from home that would be equivalent to the welfare loss associated with the covariate changes due to gentrification. See app. E.3 for details. Initial visit share is $\Sigma_{j\in\mathcal{J}^G} P_{ij}$. Social frictions are EDD, SSI, EDD × SSI, and racial and ethnic population shares of k_j. Restaurant traits are price, rating, cuisine category, and price and rating interacted with median household income. Other traits are destination income, differences in incomes, and robberies per resident.

where $\mathcal{J}^G$ is set of restaurants that change as a result of gentrification, P_{ij} is the probability that an incumbent resident would visit restaurant j prior to gentrification, and $(\Delta\bar{X}_{ij}^2, \Delta\bar{Z}_j^1, \Delta\bar{Z}_{ij}^2)$ are the average changes in restaurant and residential characteristics due to gentrification. Table 9 reports the results. The 102 restaurants in the gentrifying area account for one-tenth of predicted visits by incumbent residents prior to gentrification, so changes in the characteristics of these restaurants and residents could have large welfare effects.

Our decomposition of the welfare loss differs from the emphases of popular discussions of gentrification. The changes in restaurants' prices and cuisines have very small effects on welfare. The rise in neighborhood income levels makes consuming in these tracts more appealing to incumbents. Instead the source of welfare losses for incumbent black consumers is primarily due to the social frictions that arise with the shift of the surrounding tracts from mostly black to mostly white residents.

This exercise illustrates potential welfare costs of gentrification to incumbent residents beyond increases in housing rents. In this Harlem example, as well as a Brooklyn example reported in appendix E.3, the consumption value of the location for incumbent residents falls by a meaningful amount. This decline is not due to changes in restaurants' characteristics but increased social frictions associated with changes in surrounding neighborhoods' racial demographics.

VII. Conclusions

We use a novel data source to describe restaurant consumption in NYC and exploit properties of the conditional-logit discrete-choice model to identify how consumers value venues' and locations' characteristics. Our data set allows us to characterize how consumption in the city depends on travel times, demographic differences, crime rates, restaurant characteristics, and reviewer characteristics. We then use our estimates to compute measures of consumption segregation for NYC residents. While NYC is distinctive in terms of its population density and diversity of restaurants, we believe that our paper's results for the largest city in the United States both are interesting in their own right and establish a basis for studying consumption segregation in other settings.

Both spatial and social frictions influence consumption choices. Consistent with theories of spatial competition, spatial frictions play a large role in determining the spatial distribution of consumption within the city. Our estimates show that measures of travel time, from both home and work by both public transit and car, are relevant for predicting the restaurants patronized by NYC Yelp reviewers. Across origin-mode pairs, halving the minutes of travel time to a venue would imply that the reviewer would be two to nearly four times more likely to visit the venue from that origin by that mode.

Social frictions are suggested by the finding that consumers are less likely to visit restaurants in neighborhoods with different residential demographics. A venue in a location one standard deviation more demographically distant from a user's home location is 25–50 percent less likely to be visited. Reviewers are also more likely to visit restaurants in neighborhoods where a larger fraction of the residents share the user's race. These social frictions are asymmetric, in the sense that the negative effects of tract-level demographic differences are larger for black consumers, yet black consumers experience larger tract-level demographic differences during their average restaurant visit.

While our estimation approach exploits data on the decision to eat at restaurants across NYC, the consequences of spatial and social frictions are likely to apply to a much broader scope of life in the city. These would include both the broader scope of consumption of nontradable services, from bars to retailers, and the vast array of nonmarket activities that cause residents to traverse the city. For example, in our gentrification exercises, the reduction in the value of restaurant consumption for incumbent residents is much more attributable to increases in social frictions than changes in restaurant characteristics. These demographic changes presumably have consequences for many nonrestaurant dimensions of urban life.

We use our estimates to characterize predicted consumption segregation for the city's population. While spatial frictions, social frictions, and demographic differences in tastes cause dissimilarity in consumption choices across racial and ethnic groups, dissimilarity indices for consumption are considerably lower than the dissimilarity indices for residential locations. Life in NYC is less segregated than one might infer from looking at residential segregation alone. Our analysis of these patterns reveals that social frictions contribute more to consumption segregation than spatial frictions. A consequence of this finding is that improved transportation linkages within the city would only modestly integrate consumption further, given existing residential patterns. Eliminating social frictions would result in substantially more integrated consumption.

References

Allport, Gordon W. 1954. *The Nature of Prejudice*. Cambridge, MA: Addison-Wesley.

Anderson, Elijah. 2011. *The Cosmopolitan Canopy: Race and Civility in Everyday Life*. New York: Norton.

———. 2015. "The White Space." *Sociology of Race and Ethnicity* 1 (1): 10–21.

Anderson, Michael L., and Jeremy R. Magruder. 2012. "Learning from the Crowd: Regression Discontinuity Estimates of the Effects of an Online Review Database." *Econ. J.* 122 (563): 957–89.

Antecol, Heather, and Deborah A. Cobb-Clark. 2008. "Racial and Ethnic Discrimination in Local Consumer Markets: Exploiting the Army's Procedures for Matching Personnel to Duty Locations." *J. Urban Econ.* 64 (2): 496–509.

Ayres, Ian. 2001. *Pervasive Prejudice? Unconventional Evidence of Race and Gender Discrimination*. Chicago: Univ. Chicago Press.

Bankoff, Caroline. 2014. "Things to Know about NYC's New 25-Miles-per-Hour Speed Limit." *New York Mag.*, November 7. http://nymag.com/daily/intelligencer/2014/11/things-to-know-about-nycs-new-speed-limit.html.

Boustan, Leah. 2011. "Racial Residential Segregation in American Cities." In *The Oxford Handbook of Urban Economics and Planning*, edited by Nancy Brooks, Kieran Donaghy, and Gerrit-Jan Knaap. Oxford: Oxford Univ. Press.

Caetano, Gregorio, and Vikram Maheshri. 2019. "Gender Segregation within Neighborhoods." *Regional Sci. and Urban Econ.*, forthcoming.

Card, David, Alexandre Mas, and Jesse Rothstein. 2008. "Tipping and the Dynamics of Segregation." *Q.J.E.* 123 (1): 177–218.

Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez. 2014. "Where Is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States." *Q.J.E.* 129 (4): 1553–1623.

Couture, Victor. 2015. "Valuing the Consumption Benefits of Urban Density." Manuscript, Univ. California, Berkeley.

Cutler, David M., and Edward L. Glaeser. 1997. "Are Ghettos Good or Bad?" *Q.J.E.* 112 (3): 827–72.

Echenique, Federico, and Roland G. Fryer. 2007. "A Measure of Segregation Based on Social Interactions." *Q.J.E.* 122 (2): 441–85.

Edelman, Benjamin, and Michael Luca. 2014. "Digital Discrimination: The Case of Airbnb.com." Working Paper no. 14-054, Harvard Bus. School.

Eizenberg, Alon, Saul Lach, and Merav Yiftach. 2017. "Retail Prices in a City: An Empirical Analysis." Discussion Paper no. DP11590, Centre Econ. Policy Res., London.

Gentzkow, Matthew, and Jesse M. Shapiro. 2011. "Ideological Segregation Online and Offline." *Q.J.E.* 126 (4): 1799–1839.

Gentzkow, Matthew, Jesse M. Shapiro, and Matt Taddy. 2019. "Measuring Polarization in High-Dimensional Data: Method and Application to Congressional Speech." Working Paper no. 22423, NBER, Cambridge, MA.

George, Lisa, and Joel Waldfogel. 2003. "Who Affects Whom in Daily Newspaper Markets?" *J.P.E.* 111 (4): 765–84.

Glaeser, Edward L., Jed Kolko, and Albert Saiz. 2001. "Consumer City." *J. Econ. Geography* 1 (1): 27–50.

Glaeser, Edward L., and Jacob L. Vigdor. 2012. "The End of the Segregated Century: Racial Separation in America's Neighborhoods, 1890–2010." Civic Report no. 66, Manhattan Inst., New York.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales. 2009. "Cultural Biases in Economic Exchange?" *Q.J.E.* 124 (3): 1095–1131.

Handbury, Jessie. 2013. "Are Poor Cities Cheap for Everyone? Non-homotheticity and the Cost of Living across U.S. Cities." Manuscript, Wharton School, Univ. Pennsylvania.

Handbury, Jessie, and David E. Weinstein. 2015. "Goods Prices and Availability in Cities." *Rev. Econ. Studies* 82 (1): 258–96.

Harrison, Cassandra, Mohip Jorder, Henri Stern, et al. 2014. "Using Online Reviews by Restaurant Patrons to Identify Unreported Cases of Foodborne Illness—New York City, 2012–2013." *Morbidity and Mortality Weekly Report* 63 (20): 441–45.

Hellerstein, Judith K., and David Neumark. 2008. "Workplace Segregation in the United States: Race, Ethnicity, and Skill." *Rev. Econ. and Statis.* 90 (3): 459–77.

Houde, Jean-François. 2012. "Spatial Differentiation and Vertical Mergers in Retail Markets for Gasoline." *A.E.R.* 102 (5): 2147–82.

Katz, Michael. 2007. "Supermarkets and Zoning Laws." PhD diss., Harvard Univ.

Labaton, Stephen. 1994. "Denny's Restaurants to Pay $54 Million in Race Bias Suits." *New York Times*, May 25.

Lee, Jennifer. 2000. "The Salience of Race in Everyday Life: Black Customers' Shopping Experiences in Black and White Neighborhoods." *Work and Occupations* 27:353–76.

Luca, Michael. 2016a. "Reviews, Reputation, and Revenue: The Case of Yelp.com." Working Paper no. 12-016, Harvard Bus. School.

———. 2016b. "User-Generated Content and Social Media." In *Handbook of Media Economics*, vol. 1B, edited by David Stromberg, Simon P. Anderson, and Joel Waldfogel. Amsterdam: North-Holland.

Massey, Douglas S., and Nancy A. Denton. 1988. "The Dimensions of Residential Segregation." *Soc. Forces* 67 (2): 281–315.

———. 1993. *American Apartheid: Segregation and the Making of the Underclass.* Cambridge, MA: Harvard Univ. Press.

Mayer, Adalbert, and Steven L. Puller. 2008. "The Old Boy (and Girl) Network: Social Network Formation on University Campuses." *J. Public Econ.* 92 (1–2): 329–47.

McFadden, Daniel. 1978. "Modelling the Choice of Residential Location." In *Spatial Interaction Theory and Planning Models*, edited by A. Karlqvist, L. Lundqvist, F. Snickars, and J. Weibull, 75–96. Amsterdam: North-Holland.

Myrdal, Gunnar. 1944. *An American Dilemma: The Negro Problem and Modern Democracy.* New York: Harper.

Oberholzer-Gee, Felix, and Joel Waldfogel. 2009. "Media Markets and Localism: Does Local News *en Español* Boost Hispanic Voter Turnout?" *A.E.R.* 99 (5): 2120–28.

O'Flaherty, Brendan. 2015. *The Economics of Race in the United States.* Cambridge, MA: Harvard Univ. Press.

Pakes, Ariel. 2010. "Alternative Models for Moment Inequalities." *Econometrica* 78 (6): 1783–1822.

Paluck, Elizabeth Levy, and Donald P. Green. 2009. "Prejudice Reduction: What Works? A Review and Assessment of Research and Practice." *Ann. Rev. Psychology* 60:339–67.

Pettigrew, Thomas F., and Linda R. Tropp. 2006. "A Meta-Analytic Test of Intergroup Contact Theory." *J. Personality and Soc. Psychology* 90 (5): 751–83.

Pisarski, Alan E. 2006. *Commuting in America III: The Third National Report on Commuting Patterns and Trends*. Washington, DC: Transportation Res. Board. http://onlinepubs.trb.org/onlinepubs/nchrp/ciaiii.pdf.

Rosenthal, Brian M., Emma G. Fitzsimmons, and Michael LaForgia. 2017. "The Making of a Meltdown." *New York Times*, November 18.

Schelling, Thomas C. 1969. "Models of Segregation." *A.E.R.* 59 (2): 488–93.

———. 1971. "Dynamic Models of Segregation." *J. Math. Sociology* 1 (2): 143–86.

Schiff, Nathan. 2015. "Cities and Product Variety: Evidence from Restaurants." *J. Econ. Geography* 15 (6): 1085–1123.

Schreer, George E., Saundra Smith, and Kirsten Thomas. 2009. "'Shopping While Black': Examining Racial Discrimination in a Retail Setting." *J. Appl. Soc. Psychology* 39 (6): 1432–44.

Smith, Sandra Susan. 2010. "Race and Trust." *Ann. Rev. Sociology* 36:453–75.

Taddy, Matt. 2015. "Distributed Multinomial Regression." *Ann. Appl. Statis.* 9 (3): 1394–1414.

Train, Kenneth E. 2009. *Discrete Choice Methods with Simulation.* New York: Cambridge Univ. Press.

Train, Kenneth E., Daniel L. McFadden, and Moshe Ben-Akiva. 1987. "The Demand for Local Telephone Service: A Fully Discrete Model of Residential Calling Patterns and Service Choices." *Rand J. Econ.* 18 (1): 109–23.

US Department of Agriculture. 2014. "Food-Away-from-Home." Report, US Dept. Agriculture, Washington, DC. http://www.ers.usda.gov/topics/food-choices-health/food-consumption-demand/food-away-from-home.aspx.

Waldfogel, Joel. 2007. *The Tyranny of the Market: Why You Can't Always Get What You Want*. Cambridge, MA: Harvard Univ. Press.

Yelp. 2013. "Yelp.com Welcomes 100 Million Unique Visitors in January 2013." http://officialblog.yelp.com/2013/02/yelpcom-welcomes-100-million-unique-visitors-in-january-2013.html.

Zhang, Junfu. 2011. "Tipping and Residential Segregation: A Unified Schelling Model." *J. Regional Sci.* 51 (1): 167–93.