Gender Homophily and Segregation Within Neighborhoods

Gregorio Caetano and Vikram Maheshri*

October 26, 2016

Abstract

Homophily generates segregation, reducing diversity in peer groups and leading to narrower social interactions. Using novel data from Foursquare, a popular mobile app that documents the activity of millio ns of people, we document robust, highly localized gender homophily: over half of the gender segregation of individuals' recreational and commercial activities in thousands of venues (e.g., shops, restaurants, parks, museums) in eight major US cities occurs within census blocks. We study some of the determinants and consequences of such segregation. A higher variety in the supply of venues on a block attracts more gender-balanced visitors, but, perversely, more intense sorting across those venues ultimately reduces the actual exposure of individuals to gender diversity in venues. Using employment data from the US Census, we find evidence that these homophilic forces widens the gender gap in labor force participation. Our analysis also suggests that localized homophily along other demographic dimensions may be similarly prevalent. JEL: R1, R2, R3, J1, J3. Keywords: Gender Segregation, Homophily, Peer Groups, Urban Sorting, Diversity, Gender Gap.

1 Introduction

Homophily, or the tendency of similar people to associate with each other (McPherson et al. (2001)), is a pervasive, gravitational social force that leads to segregated peer groups. Segregation as a social phenomenon has been been widely studied in a number of important contexts, such as residential neighborhoods, schools and workplaces (Card et al. (2008a); Boustan (2012); Echenique et al. (2006); Fernandez et al. (2000)). While segregation at these levels partially determines peer groups, many further daily choices may expose people to very different social interactions. For

*University of Rochester and University of Houston. We are extremely grateful to Foursquare, Inc. for generous access to their restricted, anonymized data, and to Michael Li and Blake Shaw for helpful discussions about the data. We thank Alexei Alexandrov, Dionissi Aliprantis, Dani Arribas-Bel, Donald Davis, Joan Monras, Richard Murphy, Romans Pancs, Stephen Ross, and various seminar and conference participants for helpful discussions. We also thank Riley Hadden and Hao Teng for excellent research assistance. All errors are our own.

instance, neighbors may shop at different supermarkets, students may select different extracurricular activities, and coworkers may exercise at different gyms. These examples suggest that the effects of residential, educational and occupational segregation on various outcomes may be mitigated by homophilic forces operating at lower levels than are commonly analyzed or observed.

Although these mundane decisions may influence peer group formation, they are difficult to account for due to data availability. This difficulty is compounded by what we term the *paradox of diversity*: as individuals are supplied with a more diverse set of choices, they will tend to be exposed to a less diverse set of peers. The ideal environment for this paradox to prevail, one that is densely populated with diverse individuals and options, is precisely that in which neighborhood effects have been most widely studied: large, metropolitan areas.

In this paper, we exploit a unique data set from a prominent location-based social network, Foursquare, that documents how individuals in eight major US cities[1] sort by gender across tens of thousands of commercial and recreational venues such as shops, restaurants, parks, churches and museums that offer the activities that constitute much of people's social lives. We find evidence of substantial gender homophily in individuals' venue choices, which results in more highly segregated peer groups than would otherwise be measured with residential data alone. We also find strong evidence for the paradox of diversity. Gender segregation at the finer, venue level is facilitated by the urban landscape, as neighborhoods rich in a variety of offerings encourage homophilic forces. Although each of the decisions that we observe are trivial in isolation, cumulatively, they can have measurable effects on important economic and social outcomes. We show that gender segregation in venues tends to widen the gender gap in labor force participation.

Gender segregation is particularly attractive for our analysis since it has been relatively understudied in urban environments, perhaps because of the natural tendency for residential neighborhoods to be well integrated and gender-diverse. This makes the co-existence of an extensive literature on (usually non-gender related) neighborhood effects and another, separate but equally extensive literature on gender peer effects in other environments somewhat incongruous.[2] Our analysis bridges these two literatures by illuminating a mechanism by which seemingly gender integrated neigh-

[1] Our analysis covers New York City, Los Angeles, Chicago, Dallas, Washington DC, San Francisco, Atlanta and Philadelphia.

[2] See, for example, (Chetty and Hendren (2015); Glaeser et al. (1996); Ioannides and Loury (2004)) on neighborhood effects, and (Lavy and Schlosser (2011); Adams and Ferreira (2009); Arcidiacono and Nicholson (2005)) on gender peer effects.

borhoods can have important effects on gender specific outcomes. Residential diversity need not translate into gender diverse peer groups, and the extent to which these two differ is a function of the neighborhood itself. Put plainly, men and women tend to lead an economically important part of their social lives in parallel, despite living amongst one another.

Our first finding is that gender segregation is highly localized: 80-90 percent of such segregation in venues is observed within census tracts, and over half of it is observed within census blocks. Thus, the level of actual gender segregation to which people are exposed is substantially higher than can be measured from residential neighborhood data. Put another way, the level of gender diversity to which people are exposed in day-to-day activities is at least 25 percent (one standard deviation) lower than the reported aggregate levels of residential diversity for roughly half of all neighborhoods in our sample.

Given this robust finding, we ask how gender segregation arises and what policies, if any, can affect local diversity. Segregated peer groups might potentially arise from two homophilic forces: active segregation might occur because individuals prefer the company of similar peers, and passive segregation might occur because similar individuals prefer similar activities, which leads them to visit the same venues. Although active segregation has been more widely discussed in the broad literatures on segregation and discrimination (Schelling (1969); Bruch and Mare (2006); Bobo et al. (2012); Boustan (2012)), passive segregation can also be an important driver of segregation (Banzhaf and Walsh (2013); Caetano and Maheshri (2014)). Empirically, we find that the segregation patterns in our data are most consistent with passive homophilic forces: men and women simply tend to prefer different types of activities. We present a simple model of venue choice in the spirit of Hotelling (1929) to illustrate how variety in the supply of venues in a neighborhood affects diversity at both the neighborhood and the venue levels, and we show that the directions of both of these effects are theoretically ambiguous. We then estimate these effects with three alternative and complementary identification strategies that all support the paradox of diversity: greater venue variety attracts more gender-diverse visitors to a neighborhood, but once there, individuals tend to self-segregate more intensely across venues, thereby *reducing* the amount of gender diversity to which they are exposed. As a result, denser urban areas may actually foster narrower social interactions by providing more opportunities for people to sort into specific venues.

It is tempting to dismiss interactions in shops and recreational venues as trivial and irrelevant, but cumulatively they may be influential. We illustrate one of many potential consequences of gender segregation in venues by connecting our analysis to the literature on gender gaps in the labor market. Large and persistent gender gaps in wages, labor force participation and promotion[3] have been mostly attributed to legal and political institutions, discrimination and human biology.[4] More recently, a growing number of studies have identified social explanations for these gaps.[5] Notably, Bayer et al. (2008) show that job referral networks among residents are highly local and highly homophilic across observable demographic dimensions, including gender, which suggests that gender homophily may lead to self-reinforcing gender gaps in labor market outcomes, as men and women will be more likely to discuss job opportunities among themselves. We conjecture that similar networks may develop between neighborhood residents and venue visitors (who might or might not reside in the same neighborhood) through their day-to-day interactions. We attempt to estimate the effects of these interactions with an identification strategy similar to Bayer et al. (2008) that compares the gender gap in labor force participation among residents of nearby blocks that differ only in the gender diversity of their venue visitors.[6] Among low wage employees, we find that venue diversity reduces the gender gap in labor force participation; among medium and high wage employees, we find an effect in the same direction, albeit smaller and statistically indistinguishable from zero. We also find that block diversity, which attempts to proxy for exposure to diversity outside of venues, has no effect on the gender gaps of any types of employees. These results are consistent with findings that informal social networks are particularly valuable to individuals who are less attached to the labor market (Fernandez et al. (2000); Montgomery (1992); Ioannides and Loury (2004)).

All of our empirical findings are similar for all cities in our sample from dense, older cities such as New York City to sprawling, younger cities such as Dallas. They are also similar for each city across days of the week and across seasons. Moreover, in a parallel analysis, we find qualitatively

[3]Such gender gaps in the US labor market have been documented for over a century (e.g., Goldin and Polachek (1987), Goldin (2014)). Altonji and Blank (1999) offer a survey of this literature.

[4]See for example Becker (1985); Goldin and Polachek (1987); Even and Macpherson (1993).

[5]For example, Fernández et al. (2004) argue that changing social norms due to women's labor force participation during World War II led to lasting, long run effects on the gender gap in employment. Castilla (2005) finds improved job performance among employees that are linked by stronger social networks. Mammen and Paxson (2000) suggest the differences in social status of working women across societies as an explanation for observed international differences in gender gaps. Pan (2015) finds gender homophily in occupation choice.

[6]The key identifying assumption we make is that residents cannot sort *by gender* across census block within block group. We argue that our conclusions are conservative if this assumption is not valid.

similar results for age homophily instead of gender homophily. Finally, our findings are robust to many potential sources of bias and measurement error.[7] We argue that these robust patterns speak to the external validity of our results in several directions.

In a broader sense, our findings highlight some of the difficulties involved in top-down approaches to strengthen the social interactions that form the basis of thriving urban environments, as endogenous homophilic forces at very local levels may undermine well meaning place-based policies. Many policy interventions seek to expose individuals to different peers by manipulating their assignment to neighborhoods, schools, workplaces, and other environments. But in reality, homophily lead individuals to form less diverse peer groups than those encouraged by their assignment, which may substantially attenuate the intended effect of such interventions. For instance, hidden homophily may help explain how interventions such as the Moving To Opportunity program (e.g., Kling et al. (2007)) has been less effective at changing youth outcomes than expected. Moreover, these tendencies might fly "under the radar" if people mostly self-segregate passively. Ultimately, our analysis suggests that homophily in important dimensions such as gender, age, race, income, and political beliefs, much of which may be highly local and thus typically hidden to researchers, shapes society in profound ways outside of the purview of policymakers.

Finally, our paper attempts to make a methodological contribution to a growing literature that leverages user-generated data to study behaviors that previously proved impossible to observe (Couture (2014); Davis et al. (2014)). Although these datasets offer much promise, they are often plagued by selection issues that make it difficult to extract a meaningful, externally valid signal. Moreover, many broad questions in social science are difficult to approach comprehensively because in practice, one cannot observe all of the choices that jointly determine individuals' social interactions. Furthermore, segregation is an end result of homophily along many potential dimensions, many of which are difficult to observe. Thus, any study like the one conducted in this paper is bound to use data that is both incomplete and unrepresentative. With these obstacles in mind, we develop an empirical approach throughout the paper that reaches only conservative qualitative conclusions, i.e., all of our conclusions would plausibly strengthen with access to more detailed and complete data. Such an approach could be useful for other studies facing similar issues; to that end, we provide a

[7] All of these results, along with a Monte Carlo study of measurement error in our data, are included in the appendix, available at http://bit.ly/1KzNf2X.

detailed sensitivity analysis in the appendix including a comprehensive Monte Carlo study of the implications of potential selection issues in our user-generated data.

The remainder of the paper is organized as follows. In Section 2, we describe our data set, and in Section 3, we show widespread evidence of gender segregation in location choices. In Section 4, we explore the causes of this phenomenon with a simple model of sorting across venues, and we show empirically that the variety of venues that are available in neighborhoods impacts both the levels of diversity in venues and in neighborhoods but in opposite directions. In Section 5, we show suggestive evidence that venue diversity narrows the gender gap in labor force participation. In Section 6, we discuss the external validity of our findings to other environments and along dimensions other than gender. We conclude in Section 7.

2 Data

For a local analysis of individual interactions, we require comprehensive, disaggregated data of their whereabouts across a large number of locations within small neighborhoods; this is difficult to observe directly. We circumvent this issue with novel, proprietary data from Foursquare, Inc., creators of the eponymous mobile app and social network that allows users to document their precise whereabouts electronically. Upon arriving at a venue, Foursquare identifies the venue by GPS on the user's mobile phone, and the user can electronically "check in". We use information on the demographic composition of Foursquare users in each venue to construct a proxy for the actual demographic composition of all individuals (i.e., Foursquare and non-Foursquare users) in the venue. Although this raises important concerns of sample selection, we develop an empirical approach with these concerns specifically in mind. We show that our empirical approach allows us to extract a meaningful signal about the sorting of all individuals across venues from this novel dataset. A comprehensive sensitivity analysis concerning the potential sources of measurement error that might exist in our data is provided in the appendix.[8]

Ours is the first paper to use this large and highly detailed database of venue visitors to study diversity within neighborhoods.[9] Foursquare is particularly suitable for our analysis because it is a

[8]In the appendix, we show that our main results will not change even if the propensity to "check-in", conditional on visiting the venue, differs by gender. We also conduct a Monte Carlo study that shows that our results are, if anything, conservative. That is, in worst case scenarios of measurement error (i.e., when propensity to "check-in" varies by gender as a function of the female share of actual visitors) our main results change only slightly and, if anything, become stronger in the direction of our conclusions.

[9]A small number of studies (e.g., Arribas-Bel and Bakens (2014)) have begun to use Foursquare data obtained

prominent location-based social network that boasts a large number of active users (over 50 million worldwide checking in over 6 billion times as of March 2015), which makes for a highly detailed catalog of activity.

Our data set contains information on all Foursquare activity in venues in eight major US cities: Atlanta, Chicago, Dallas, Los Angeles, New York City, Philadelphia, San Francisco and Washington, DC. Our specific sample regions are defined as the counties in which these cities are primarily located.[10] For each of the 76,377 venues that are tracked in these cities, Foursquare has directly provided to us in fully anonymized form the number of daily check-ins by male and female users from August 1, 2012 to July 31, 2013. This data is aggregated to the venue level, hence we cannot observe any characteristics of individual Foursquare users, nor can we track a particular individual's activity. We restrict our sample to venues that experienced at least 10 check-ins during the sample period to improve our measurements of the gender compositions of venues.[11] In total, these venues experienced 49.6 million check-ins during the sample period with the average venue in our sample experiencing 649 check-ins. Each venue in our data set is also geo-coded by latitude and longitude, which allows us to link venues to unique census tracts, block groups and blocks using neighborhood definitions from the 2010 Decennial Census.

In Table 1, we summarize our sample by city and by venue classification. Not surprisingly, larger cities such as New York and Los Angeles have more venues and check-ins. Males tend to check in slightly more than females on average, but there is substantial and robust variation in the gender composition of venues in all cities. It is immediate that there is more variation in the average gender composition of venues across categories than across cities and more variation in gender composition within categories than within cities.[12] The 9 categories of venues are further classified into 225 narrow subcategories; detailed summary statistics disaggregated by subcategory can be found in the appendix.

indirectly via the Foursquare API (application programming interface). Foursquare data obtained via the API unfortunately does not disaggregate check-ins along any demographic dimension.

[10]The counties are Fulton (Atlanta), Cook (Chicago), Dallas, Los Angeles, New York, Philadelphia, and San Francisco respectively. We treat the entire District of Columbia as the "county" for Washington. Most of the cities in our sample are entirely contained in their corresponding county with the notable exception that New York County only contains the borough of Manhattan.

[11]This restriction does not bias our results (see appendix).

[12]For each city in our sample, check-ins across venues are approximately distributed log-normally.

Table 1: Summary Statistics

City	Venues	Check-ins	μ	σ	$p_{75} - p_{25}$	Tracts	B. Groups	Blocks
Atlanta	4,115	2.84	0.46	0.17	0.19	180	361	1,307
Chicago	13,665	8.11	0.49	0.16	0.19	1,100	2,235	6,237
Dallas	5,065	2.40	0.45	0.16	0.19	421	774	1,986
Los Angeles	23,108	10.2	0.46	0.15	0.18	1,902	3,584	9,182
New York City	16,203	16.2	0.49	0.17	0.19	282	945	2,501
Philadelphia	3,933	2.10	0.47	0.16	0.19	301	568	1,757
San Francisco	6,601	4.78	0.42	0.15	0.16	182	440	1,898
Washington, DC	3,687	2.98	0.43	0.16	0.17	152	272	1,069
Category							**Unique Subcategories**	
Food	31,398	16.6	0.45	0.13	0.17		65	
Shops/Services	20,903	9.97	0.52	0.21	0.28		66	
Bars	6,441	6.52	0.44	0.12	0.13		20	
Outdoors	4,795	4.62	0.44	0.16	0.21		22	
Cafes	4,483	3.88	0.47	0.14	0.18		3	
Entertainment	4,189	4.08	0.46	0.13	0.15		29	
Hotels	1,798	2.24	0.40	0.11	0.13		5	
Gyms	1,625	1.41	0.49	0.23	0.34		12	
Spiritual	745	0.29	0.48	0.17	0.23		3	

Notes: Check-ins reported in millions. μ and σ refers to the mean and standard deviation of the proportion of females in venues, and p_{25} and p_{75} refer to the 25^{th} and 75^{th} percentiles of the proportion of females in venues.

Because we observe daily check-ins at each venue, we can assess whether there are any dynamic trends in our data over the sample period. As shown in the first panel of Figure 1, there is substantial day-of-week variation in check-ins since venues are more highly frequented on weekends, but the gender composition of check-ins is nearly constant. This suggests that we can aggregate the data at least to the weekly level to analyze gender diversity. We do so and check for aggregate weekly trends in our data in the second panel of Figure 1. There is no systematic weekly variation in check-in frequency and no discernible seasonality or aggregate trend. More importantly, the gender composition of check-ins is roughly constant throughout the sample period. This suggests that we can aggregate the data set to the annual level to analyze gender diversity without loss of generality.

Figure 1: Check-ins and Gender Composition Over Time

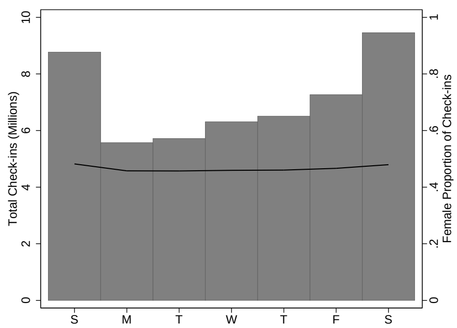

(a) Day-of-Week Variation

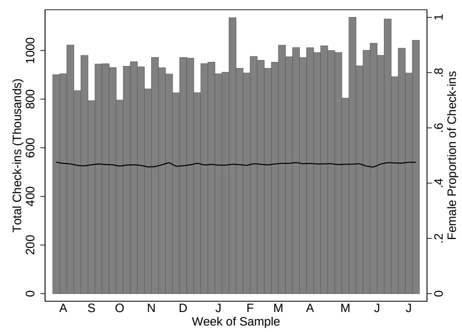

(b) Week-of-Year Variation

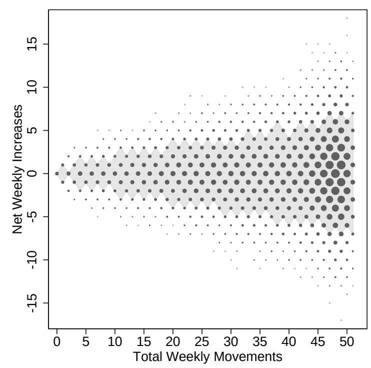

(c) Within Venue Variation

Notes: (a), (b): Bars represent total check-ins, lines represent gender composition of aggregate check-ins. The 53rd week of the sample is omitted because it only contains a single day. (c): In this scatter plot of venues in our data, larger dots correspond to a greater numbers of venues. A venue experiences a weekly increase (decrease) in gender composition if the proportion of female check-ins rises (falls) by at least one percentage point.

To further support this choice of aggregation, we check whether the gender compositions of individual venues follow a trend over time. For each venue, we compute the net number of week-on-week increases (increases minus decreases) in the proportion of female check-ins over the sample period, and we plot them against the total number of changes in the proportion of female check-ins in Figure 1.[13] Larger dots represent more venues in the sample, and the shaded region is defined to include 95% of all venues. It is immediate that most venues experience roughly as many relative increases in female popularity as relative decreases in female popularity. Because the gender

[13] A venue is defined to experience a week-on-week increase (decrease) in the female share if its female share increases (decreases) by a threshold of at least one percentage point over consecutive weeks. The total number of changes in the proportion of female check-ins is equal to the sum of increases and decreases. We replicated panel (c) of Figure 1 with alternative thresholds of 5, 10 and 15 percentage points and obtained qualitatively similar results.

composition of a venue tends to vary around a fixed value, it is appropriate to interpret longitudinal variation in check-ins as measurement error, which we minimize by aggregating our data to the annual level in order to focus on the more relevant cross-sectional variation in our data.[14]

3 Measuring Gender Homophily and Segregation In Neighborhoods

Homophily will lead the gender compositions of the venues to diverge from one another as individuals sort across them. In the extreme case, if some venues are only visited by females and others are only visited by males, then the venues are fully segregated and exhibit no gender diversity. One important and widely used measure of segregation is the Theil (1967) index.[15] Formally, if s_{jk} is the share of females at venue j located in neighborhood k, then the Theil index of neighborhood k is given by

$$T_k = \frac{1}{n_k} \sum_{j \in k} \left(\frac{s_{jk}}{\bar{s}_k} \cdot \log \frac{s_{jk}}{\bar{s}_k} \right) \tag{1}$$

where n_k is the number of venues in the neighborhood and $\bar{s}_k$ is the simple average of s_{jk} across all venues in the neighborhood.[16] If the neighborhood is fully integrated (i.e., no observable homophily and hence maximal diversity), then all of its venues will have the same gender composition as the neighborhood overall, and $T_k = 0$. Neighborhoods with less diverse venues have larger values of T_k.[17] In practice, k can correspond to the entirety of a city (c), a census tract (t), a census block group (g) or a census block (b), so T_k represents the extent to which venues in k are segregated by gender.

[14] As a robustness check, we replicated all main results of the paper by month-of-year and by day-of-week and found similar results (see appendix).

[15] Weitzman (1992) proposes a general, recursively defined measure of diversity that satisfies numerous attractive mathematical, economic and conceptual properties. In certain contexts, he shows it to be equivalent to the widely used Shannon index, which measures the amount of "true diversity" or the effective number of different types of "objects". In our application, objects correspond to venues by demographic composition, and the Shannon index reduces to the Theil index up to an additive constant.

[16] Our results are virtually unchanged if we denote s_{jk} as the share of men in venue j in neighborhood k.

[17] The maximum value that the Theil index can take is $\log n_k$, which varies with the density of venues in a neighborhood. Where applicable, our results using the Atkinson (1970) index (the Theil index divided by $\log n_k$, thus normalized to values between 0 and 1) are all qualitatively equivalent. As we explain below, we use the Theil index instead of the Atkinson index in our analysis because of its decomposability properties.

Figure 2: Densities of Theil Indices for Various Neighborhood Definitions

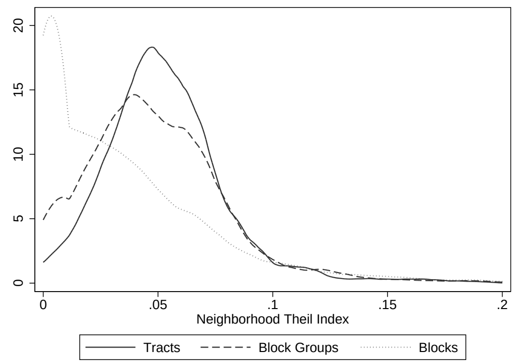

Notes: All densities are estimated using a bandwidth of 0.005 and an Epanechnikov kernel. For clarity, we present the density only for values of the domain less than 0.2; fewer than 1% of neighborhoods of any type have a Theil Index in excess of 0.2. Theil Indices are pooled across neighborhoods in all cities.

We compute the Theil index for each tract, block group and block in the cities in our sample and present the densities of these indices in Figure 2. The bulk of the density of T_t lies away from zero, which reveals that individuals sort within tracts. Similarly, the bulk of the density of T_g lies away from zero, which reveals that individuals also sort within block groups. The density of T_b is close to zero for approximately 10% of the sample, so roughly 90% of blocks in these cities are further sorted by gender in venues. Mathematically, $T_b \leq T_g \leq T_t$ for all $b \in g \in t$. Because these three densities roughly coincide for higher values of the Theil index, all of the sorting in highly homophilous tracts and block groups occurs within their constituent blocks as opposed to across them.

The Theil index possesses the attractive property of being additively decomposable, which allows for segregation in an entire city to be split into one term that captures segregation within neighborhoods and another term that captures segregation across neighborhoods.[18] Formally, we

[18] Although the Theil index is not the only such measure that is additively decomposable, it is the only one that is homogeneous of degree zero (Bourguignon (1979)), which makes it invariant to rescaling. This is important in our application because males may be more or less likely to check in on Foursquare than females; hence in order to maintain the external validity of our estimates we should make only relative comparisons of homophily. In addition, as Shorrocks (1980) points out, other commonly used measures of segregation, diversity, exposure or inequality with other attractive properties which are based on the Herfindahl index (e.g., the index of segregation introduced in Ellison and Glaeser (1997)) or the Gini coefficient are not additively decomposable, so they are less useful and appropriate in our context.

can decompose the total Theil index of city c into within- and across- tract components as

$$T_c = \underbrace{\sum_{t \in c} \alpha_t \cdot T_t}_{\text{within-tracts}} + \underbrace{\sum_{t \in c} \alpha_t \cdot \log \frac{\bar{s}_t}{\bar{s}_c}}_{\text{across-tracts}} \tag{2}$$

where the weights $\alpha_t = \frac{n_t s_t}{n_c s_c}$ correspond to the contribution of each tract to overall venue diversity in c (s_k represents the share of females across all venues in neighborhood k). T_c can be similarly decomposed to the block group or block levels. The key benefit of this simple decomposition is that we can analyze neighborhood segregation (and hence homophily in venues) independently of how individuals sort across neighborhoods. In Table 2, we present the proportion of city-wide gender segregation in venues that is attributable to homophily within neighborhoods, i.e., the contribution of the first term of equation (2).[19] Intuitively, this captures how much of the variation in the gender composition of city venues is "local." It is immediate that the majority of homophily and resulting gender segregation in city venues is highly localized.

Table 2: Venue Sorting Within Neighborhoods

	Proportion of city-wide segregation attributable to homophily within:		
	Tracts	Block Groups	Blocks
Atlanta	0.89	0.83	0.59
Chicago	0.82	0.74	0.47
Dallas	0.79	0.71	0.48
Los Angeles	0.83	0.74	0.50
New York City	0.92	0.88	0.78
Philadelphia	0.85	0.78	0.50
San Francisco	0.83	0.78	0.57
Washington, DC	0.88	0.84	0.61

Note: Bootstrapped standard errors for all entries in all cities are less than 0.005 and are omitted for clarity.

To better interpret the measures in Table 2, we can benchmark the observed gender compositions of venues against the gender compositions of venues that would be hypothetically observed if there

[19] We calculated bootstrapped standard errors with 500 repetitions for the means of T_t, T_g and T_b for each city separately. All means are statistically significantly different from zero at the 99% level.

was no homophily within neighborhoods.[20] This exercise reveals how much additional segregation we can measure because we can observe sorting across venues within neighborhoods as opposed to only sorting across neighborhoods (as in the vast majority of studies). By observing sorting at the more disaggregated venue level, we are able to detect 2-4 times more homophily than in data aggregated to the block level, and 4-12 times more homophily than in data aggregated to the tract level.[21] For Manhattan, these numbers are on the higher end: we are able to detect 4 (12) times more homophily than we would have with data aggregated to the block (tract) level.

Finally, because we find gender homophily in all of the choices that we are able to observe – men and women systematically visit different tracts within a city, different block groups within a tract, different blocks within a block group, and different venues within a block – it is likely that we are underestimating the extent to which homophily actually mitigates exposure to diversity in peer groups. For example, individuals may sort to the same restaurant at different times of the day, to different tables in the same restaurant, or even to different conversations at the same table.

Remark 1. Endogenous sorting at highly local levels may dramatically affect the transmission of peer effects and the intensity of neighborhood effects. For instance, many studies have found that interventions aimed at exposing poor individuals to more affluent neighbors have not been as successful as expected (e.g., Oreopoulos (2003), Kling et al. (2007)). Even if exposure to different peers truly mattered, an intervention that assigns individuals to different groups but still allows for further endogenous sorting may be ineffective (Weinberg (2007)). For example, assignment in the "Moving to Opportunity" (MTO) program occurred at the residential census tract level, which we find here to be a poor proxy for the actual exposure of individuals to diversity. Indeed, Kling et al. (2007) find some positive effects of MTO for women who are less likely to interact with individuals from their original neighborhood than men.[22] A more recent analysis of MTO (Chetty and Hendren (2015)) has found positive effects among individuals who were assigned to a new neighborhood at a very young age. This could partly be due to the fact that there is less scope for endogenous sorting along demographic dimensions for the very young. Psychologists have identified non-cognitive

[20]We also conduct a falsification exercise in which individuals are not allowed to sort within blocks to validate this benchmarking exercise and ensure that our results are not simply artifacts of sampling error. The details and results of this exercise are provided in the appendix.

[21]To obtain these figures, we take the reciprocal of the proportion of observed venue sorting due to homophily within neighborhoods (e.g., $(1-0.89)^{-1} = 9.09$ for Tracts in Atlanta).

[22]See Table G-5 in their supplementary appendix.

factors such as aggression and anxiety as the most powerful predictors of early childhood friendship formation (see Ladd (1999) for a survey) unlike adolescent friendship formation, which depends more on demographics (e.g., Currarini et al. (2009)). These results highlight the need to understand the extent to which homophily operates within assignment groups.[23]

Neighborhood Residents vs. Visitors

Typically researchers can observe only the broad location choices that individuals make such as the neighborhoods where they reside. Because we also observe the choices of which neighborhoods people visit, we can compare the relative amounts of homophily among residents and visitors.

Figure 3: Residential Homophily vs. Visitor Homophily

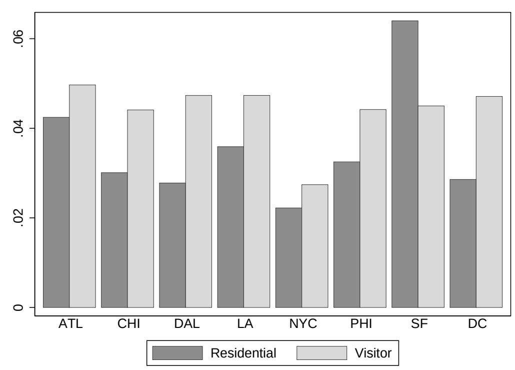

Note: Residential homophily is calculated as the Theil index of the gender composition of block residents from the 2010 Census. For comparability, visitor homophily is calculated as the Theil index of the gender composition of check-ins in blocks. Bootstrapped standard errors for all estimates are below 0.005 and are omitted for clarity.

In Figure 3, we compare how residents sort across blocks with how visitors sort across blocks for each city in our sample. Residential homophily is calculated as the Theil index of the gender composition of block residents for each city from the 2010 Census. Visitor homophily is calculated as the Theil index of the gender composition of block visitors for each city from our data, which is equivalent to the second term in a block level decomposition of T_c according to equation (2). We find that for all cities except one, there is less residential homophily than visitor homophily.[24]

[23] There have been several notable findings of peer effects in non-neighborhood settings where the assignment to peer groups leaves little or no room for further endogenous sorting. For instance, Sacerdote (2001) finds that exposure to a roommate with high academic and social ability has a variety of positive effects on students, and Mas and Moretti (2009) finds that workplace productivity increases with peer effects in two person teams.

[24] The exception is San Francisco. We speculate that this is due to San Francisco's sizable gay population, which concentrates residentially in certain neighborhoods whose venues are visited by a very gender diverse population. Indeed, San Francisco, like all other cities in the sample, exhibits less residential homophily than visitor homophily by age (see appendix).

Our findings suggest that by understating homophily, studies that rely on residential data alone will substantially overstate individuals' exposure to diversity. For example, if we were to observe residential data only, then we would estimate that the average woman in neighborhood k would randomly encounter another woman with probability equal to the female share of residents in that neighborhood. However, if we were to observe data disaggregated to the venue level, then we could better estimate that the average woman in k would randomly encounter another woman with probability $\sum_{j \in k} \frac{f_{jk}}{f_k} \cdot s_{jk}$ where f_{jk} is the number of women visiting venue j, and f_k is the total number of women visiting neighborhood k.

To provide some empirical context for this thought experiment, we calculate how much more likely we would estimate that a woman would encounter a woman in venue level data than in residential data, and present the empirical cumulative distribution of this difference (which corresponds to the extent to which residential data fails to capture homophily that is observable in venue data) in Figure 4. To be conservative, we define neighborhoods as census blocks, which are the smallest residential geographic units that are used by researchers. In roughly 40 percent of the blocks, we would underestimate exposure to peers of the same gender by at least 25 percent (or about one standard deviation of the probability that a woman would encounter another woman in a venue) if we used residential data instead of venue data. In roughly 20 percent of the blocks, we would underestimate exposure to peers of the same gender by at least 50 percent.

Figure 4: Excess Predicted Homophily in Venue Data

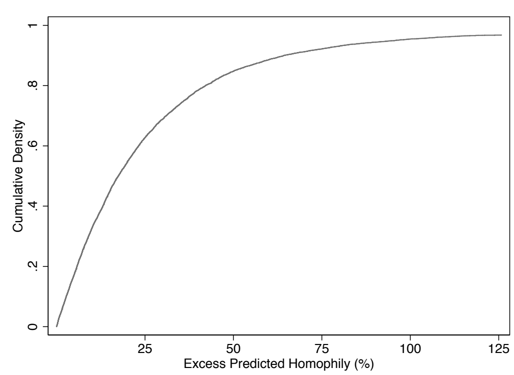

Note: In this figure, we present the empirical cumulative distribution of how much more likely we would predict that a woman would encounter another woman in a census block using venue level data than using residential data.

4 Local Determinants of Gender Diversity

4.1 A Simple Model of Venue Choice

The substantial gender homophily in venue choices creates a wedge between the levels of gender diversity that are observed in venues and in neighborhoods. In order to explore how variety in the supply of venues might be a determinant of this wedge, we consider a simple, stylized model of how individuals choose between venues within and across different neighborhoods in the spirit of Hotelling (1929).

On the supply side, we model a neighborhood k as a collection of venues indexed by j, each of which possess a single particular characteristic x_j that lies on the unit interval and differentiates the venues. This characteristic can be thought of as the venue's particular type of offering. The spatial distribution of venues on the unit interval corresponds to the variety of venues in the neighborhood. For example, a Mexican restaurant and a Chinese restaurant would lie closer to each other on the interval than a Mexican restaurant and a shoe store. More generally, when venues are more spread out, they collectively represent a greater variety of venue offerings, which we denote as V_k. To simplify our analysis as much as possible, we consider the simplest setting that could feature sorting across venues: a single neighborhood with two fixed venues (i.e., $j \in \{1, 2\}$). In such a neighborhood, $V_k = |x_1 - x_2|$.[25]

On the demand side, we assume that there is a mass of individuals, each of whom are indexed by i and possess a utility function over venues $U_i(x; \delta_i)$. U_i is assumed to be a single peaked function around the point δ_i, which represents individual i's ideal point (e.g., $U_i(x) = u - (\delta_i - x)^2$). Once again, we consider the simplest specification of demand that could feature sorting across venues: individuals belong to one of two equally sized groups of potential venue visitors: males and females. The δ_i are drawn from different distributions depending on i's gender. Each individual is assumed to choose at most one venue that maximizes their utility provided $U_i > 0$.[26] If more than one venue offers an individual maximal positive utility, then ties are broken randomly.

[25] In general, x_j could also refer in part to the physical locations of venues. In such a formulation, connected subsets of the unit interval could correspond to physical neighborhoods, and we could study sorting across neighborhoods as well. For simplicity, we abstract away from this formulation because our empirical analysis exploits only very local variation in venue variety.

[26] This condition accommodates an outside option within the model. Individuals for whom $U_i \leq 0$ for all available x_j should be understood to choose the outside option, which reflects visiting another neighborhood or staying at home.

Figure 5: Venue Variety and Diversity

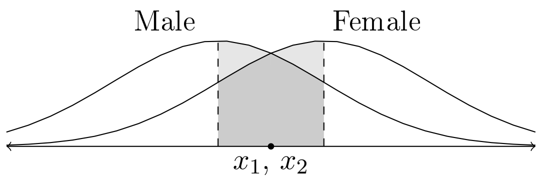

(a): Low V, Max D^V, Max D^N

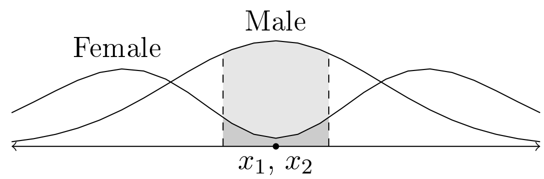

(a′): Low V, Max D^V, Low D^N

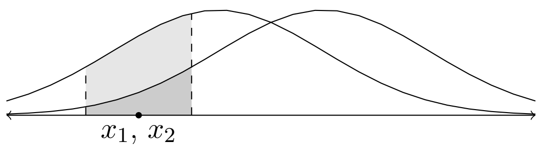

(b): Low V, Max D^V, Low D^N

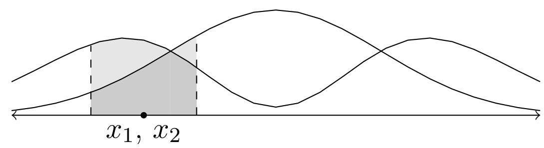

(b′): Low V, Max D^V, High D^N

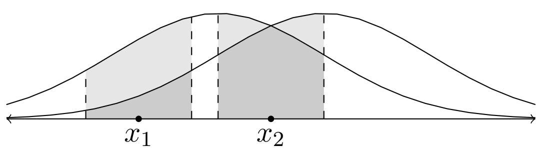

(c): Med. V, Med. D^V, Med. D^N

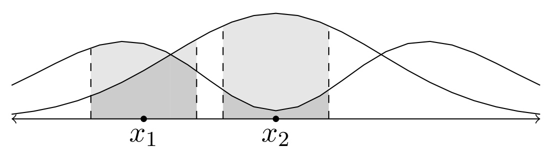

(c′): Med. V, Low D^V, Med. D^N

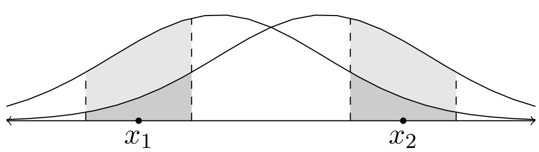

(d): High V, Low D^V, Max D^N

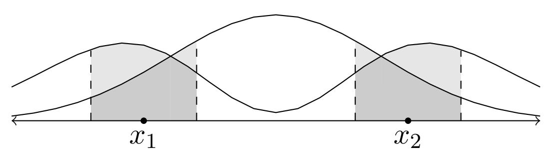

(d′): High V, Max D^V, High D^N

We combine the supply and demand sides to define equilibrium venue diversity and neighborhood diversity. Venue diversity in a neighborhood is measured by the negative Theil index of the gender composition of venues, i.e., $D_k^V = -T_k$, since higher levels of T_k correspond to less diversity. The overall amount of diversity in neighborhood k can be understood as how representative the gender composition of *actual* neighborhood visitors is relative to *potential* neighborhood visitors. Because the groups are of equal size, the latter is equal to $\frac{1}{2}$, so we can define neighborhood diversity as $D_k^N = -\left|\frac{f_1+f_2}{f_1+f_2+m_1+m_2} - \frac{1}{2}\right|$ where f_j and m_j are the numbers of female and male visitors to venue j respectively.

We use this simplified model to illustrate the relationship between neighborhood venue offerings and diversity in a series of diagrams. In Figure 5, we consider four different neighborhoods in order of increasing venue variety (a) - (d), which have counterparts (a′) - (d′) that are identical except for the demands that they face. In each of the neighborhoods in the first column, men's ideal points tend to be lower than women's ideal points. In neighborhood (a), there is no venue variety, as $x_1 = x_2$. As a result, there is no sorting, so the neighborhood exhibits maximal venue diversity (D^V). Also, since the venues attract equal numbers of men and women, there is maximal neighborhood diversity (D^N). In neighborhood (b), $x_1 = x_2$ as before, so there is still no venue variety or sorting, and hence maximal D^V. However, D^N is low since venue visitors are unrepresentative of the population at large. In neighborhood (c), $x_1 \neq x_2$, so this neighborhood has a moderate level of venue variety, which is accompanied by a moderate amount of sorting (and hence moderate levels of D^V). As a result, this neighborhood has a moderate overall level of D^N relative to neighborhoods (a) and (b). Finally, neighborhood (d) features a high level of venue variety, which is accompanied by a high level of sorting and hence low D^V. However, because the two venues cater to symmetric groups of consumers, an equal number of men and women go to one of the venues, and hence the neighborhood has maximal D^N.

The four analogous neighborhoods in the second column, (a′) - (d′) face different demands. In these hypothetical neighborhoods, women can be classified into two groups with fairly disparate taste for activities (say, hanging out in cafes and shopping) whereas men tend to be more homogenous in their tastes for activities (say, dining at restaurants). Mathematically, while men's and women's average ideal points are now both located at $\frac{1}{2}$, women prefer venues with low and high x_j's whereas men tend to prefer venues with moderate x_j's. The resulting levels of neighborhood and venue diversity as venue variety increases in neighborhoods (a′) - (d′) are quite different from the levels of diversity in their counterparts (a) - (d). For instance, an increase in venue variety from (b′) to (c′) *reduces* both D^V and D^N, whereas a further increase in venue variety from (c′) to (d′) *increases* both D^V and D^N.

We draw three important conclusions from this stylized analysis. First, sorting is made possible only by venue variety; it is trivial to note that there will be no sorting across venues with identical x_j's (and minimal sorting across venues with very similar x_j's). Accordingly, venue variety is an attractive candidate for a determinant of the wedge between venue diversity and neighborhood diversity that we have established in Section 3. Second, the relationship between venue variety and venue diversity is theoretically ambiguous. In the neighborhoods $a - d$, venue variety and venue diversity are inversely related to each other, but in neighborhoods $a' - d'$, this relationship no longer holds. Third, the relationship between venue variety and overall neighborhood diversity is also theoretically ambiguous. In neighborhoods $a - d$, venue variety and neighborhood diversity are directly related to each other, but in neighborhoods $a' - d'$, this relationship no longer holds. The latter two implications suggest that we must empirically determine the relationships between venue variety and venue and neighborhood diversity in order to determine the extent to which venue variety creates this wedge.

4.2 A Proxy for Venue Variety

In order to generalize the model and take it to the data, we need a measure of venue variety. Intuitively, venue variety should be lower in neighborhoods with more substitutable venues whose characteristics are more similar. One important characteristic of a venue is its location. All else constant, venues located farther from each other should be less substitutable. In addition, the offerings of a venue can be proxied for by its categorization in our data.

Because the subcategories of venues are so narrowly defined, we can interpret them as proxies for x_j provided that we compare venues only in narrow geographic areas (i.e., the same location). Thus, we can recast the first conclusion drawn above in terms of something that is measurable with our data: The proportion of sorting within a neighborhood that is due to sorting across subcategories should be high if neighborhoods are narrowly defined. In Table 3 we present the proportion of sorting within neighborhood that occurs across venue types for each neighborhood definition. Our findings are consistent with the model. The bulk of sorting within neighborhoods occurs across subcategories; for instance, about 90% of sorting within census blocks occurs across subcategories.[27] However, much less sorting within entire cities occurs across subcategories. This suggests that location is a better proxy for x_j when comparing venues that far from each other.

[27]Similarly, between 50% and 60% of the sorting within blocks occurs across categories depending on the city.

Table 3: Proportion of Within-Neighborhood Sorting by Gender Due to Sorting Across Subcategories:

	City	Tracts	Block Groups	Blocks
Atlanta	0.26	0.78	0.83	0.91
Chicago	0.26	0.84	0.89	0.94
Dallas	0.27	0.82	0.86	0.92
Los Angeles	0.20	0.83	0.87	0.92
New York City	0.31	0.70	0.81	0.90
Philadelphia	0.22	0.81	0.85	0.94
San Francisco	0.28	0.76	0.82	0.92
Washington, DC	0.26	0.75	0.82	0.91

Note: Subcategories (225) are defined in the appendix. Bootstrapped standard errors for all entries are less than 0.005 and are omitted for clarity.

Remark 2. There are two potential explanations for our findings of gender segregation in Section 3: (a) men (women) prefer to be in the company of other men (women) in venues (active segregation); and (b) men and women systematically prefer different types of venues (passive segregation). For simplicity, the model we consider here allows for only the second explanation, as the first explanation does not seem to be empirically important in our context. Indeed, the first explanation should generate a "social contagion" effect which would result in dynamic trends (and possibly tipping behavior) in the gender compositions of otherwise similar venues (Schelling (1971)); certain venues of a particular type would tend to become increasingly male while others of that same type would tend to become increasingly female. The results above are inconsistent with this explanation: nearby venues of the same type have very similar gender compositions, whereas nearby venues of different types have very different gender compositions. (Moreover, we do not find systematic trends in the gender composition of individual venues as shown in panel (c) of Figure 1.) This suggests that the reason we observe most individuals going to venues filled with others of the same gender is not because they actively seek similar people; instead, men and women just tend to prefer different activities.

4.3 Identifying The Causal Effects of Venue Variety on Venue and Neighborhood Diversity

The stylized model described above reaches ambiguous conclusions about the effects of venue variety on venue and neighborhood diversity, so we identify these causal effects empirically. Consider two small, nearby neighborhoods that are otherwise similar except for their levels of venue variety. For instance, one neighborhood may feature only restaurants, whereas another neighborhood may feature both restaurants and shops (compare to neighborhoods (a) and (c) in Figure 5). Given their small sizes and proximity, it is reasonable to consider their locations and the demands that they face to be approximately the same, except for their venue offerings. Thus, differences in venue and neighborhood diversity across these neighborhoods can be reasonably attributed to the difference in their venue variety. We implement an identification strategy that makes this comparison.

Following the model, the amount of local diversity in venues in block b can be measured by the negative Theil Index, $D_b^V = -T_b$, and the overall amount of neighborhood diversity in b can be measured by how representative the distribution of visitors in b are of the distribution of visitors in the whole city. As a generalization of the model, if f_{jb} and m_{jb} represent the total number of females and males in venue j in block b, and $s_b = \frac{\sum_{j \in b} f_{jb}}{\sum_{j \in b}(f_{jb}+m_{jb})}$, then we can define

$$D_b^N = -\left|s_b - s_c\right| \tag{3}$$

to be the overall amount of diversity in block b in city c. Finally, because we compare only small neighborhoods that are close to each other, we can take advantage of the classification of venues in our data to generalize the model above and define venue variety, V_b, as either the number of unique categories or subcategories of venues that are on offer in that block.

We estimate the regression equations:

$$D_b^V = \beta^V V_b + \alpha_g^V + X_b\delta^V + R_b\lambda^V + \epsilon_b^V \tag{4}$$

$$D_b^N = \beta^N V_b + \alpha_g^N + X_b\delta^N + R_b\lambda^N + \epsilon_b^N \tag{5}$$

where α_g are fixed effects at the block group level for $b \in g$, and X_b represents a set of block control variables that includes the total number of venues and the amount of checkin activity in b, R_b represents a set of residential control variables that includes the total number and the female share of residents in b, and ϵ_b^V represents an error term.[28] β^V and β^N are the coefficients of interest. For interpretation, we normalize all variables by their standard deviations, so β^V and β^N correspond to the effects of a one standard deviation increase in venue variety on venue and neighborhood diversity respectively (in units of their standard deviations).

In Figure 6, we present estimates of β^V (darker bars) and β^N (lighter bars) along with their corresponding 95% confidence interval for each city separately, and for V_b defined as either the number of unique categories or subcategories. We systematically find that $\hat{\beta}^V < 0$ and $\hat{\beta}^N > 0$. This implies that any increase in neighborhood diversity due to a an increase in venue variety will generate more intense sorting between venues within the neighborhood, thereby reducing the exposure to diversity at the venue level. Indeed, a one standard deviation increase in venue variety will lead to roughly a 0.2 standard deviation increase in neighborhood diversity and roughly a 0.4 standard deviation decrease in venue diversity.[29]

Figure 6: $\hat{\beta}^V$ and $\hat{\beta}^N$ By City

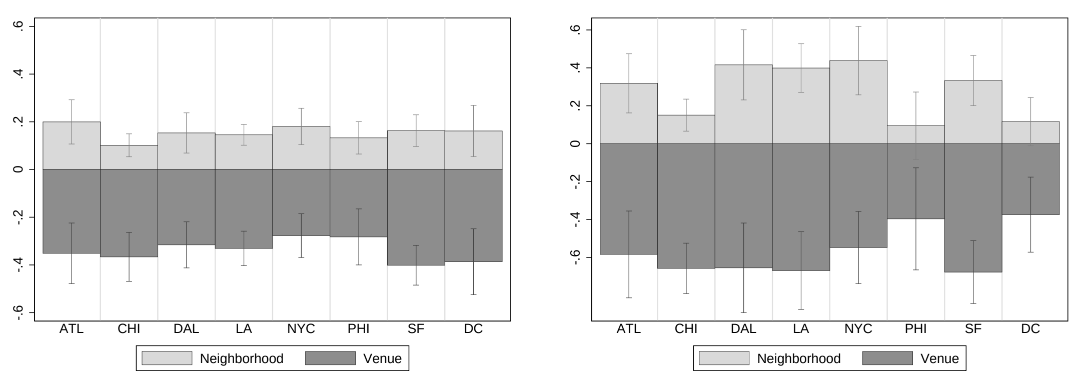

(a) V_b = Num. of Categories (b) V_b = Num. of Subcategories

Notes: The dark bars represent estimates of $\hat{\beta}^V$ from equation (4), and the light bars represent estimates of $\hat{\beta}^N$ from equation (5). 95% confidence intervals are also shown from robust standard errors clustered at the block group level. The number of observations for each of the 16 regressions is equal to the number of census blocks in each city (see Table 1), and the R^2 of each regression varies from 0.33 to 0.50.

[28]The residential control variables are obtained from the 2010 Census Summary File 1 (SF1).

[29]Our estimates of β^N corroborate the idea advanced by Glaeser et al. (2001) and Couture (2014) that the variety of venues and activities on offer is a primary amenity to urban consumers.

Can We Interpret $\hat{\beta}^V$ and $\hat{\beta}^N$ as Causal?

The causal parameters β^V and β^N are identified under the assumptions $\mathrm{Cov}\left(\epsilon_b^V, V_b|\alpha_g^V, X_b, R_b\right) = 0$ and $\mathrm{Cov}\left(\epsilon_b^N, V_b|\alpha_g^N, X_b, R_b\right) = 0$ respectively. Because we conduct our analysis at the block level, we explicitly consider small neighborhoods, and the inclusion of block group fixed effects α_g ensures that we only compare neighborhoods that are located near each other, which holds constant all determinants of the demand and supply that vary at the block group level. Still, certain neighborhood amenities that are correlated to venue variety might attract different groups of people to different nearby blocks, so we control for X_b to ensure that we compare blocks that have similar numbers of venues and levels of foot traffic, and we control for R_b to ensure the number of residents of each gender is similar across these blocks.

The remaining concern is that some unobserved neighborhood amenities that cannot be controlled for by these covariates may be correlated to venue variety. For instance, one might worry about simultaneity bias: different venues may decide to locate in neighborhoods that attract more diverse visitors, i.e. demand for venues causes supply of venues, rather than the other way around. The fact that neighborhoods are both small and close to each other in our context helps allay such concerns, as this could only be an issue if venues had control over and preferences for locating in specific blocks of a given block group. This seems implausible since locating in a particular block requires a commercial vacancy and the blocks are similar in venue density, foot traffic, and location.[30]

Nevertheless, we provide four robustness checks that address these and other concerns. The results of these four robustness checks are shown in Figure 7, where we compare the baseline estimates of β^V and β^N from equations (4) and (5) pooled over all eight cities with estimates from four alternative specifications.[31] In the first set of bars, we define venue variety as the number of distinct categories in a neighborhood, and in the second set of bars, we define venue variety as the number of distinct subcategories in a neighborhood.

[30] The motivation for this identifying assumption is analogous to the one made by Bayer et al. (2008) for residents. If the housing market is not too dense at all points in time (as appears to be case even in large metropolitan areas), then it is difficult for a venue owner to choose an exact census block in which to locate.

[31] We also conducted these robustness checks for each city separately and obtained similar results, which are reported in the appendix.

Figure 7: $\hat{\beta}^V$ and $\hat{\beta}^N$: Alternative Identification Strategies

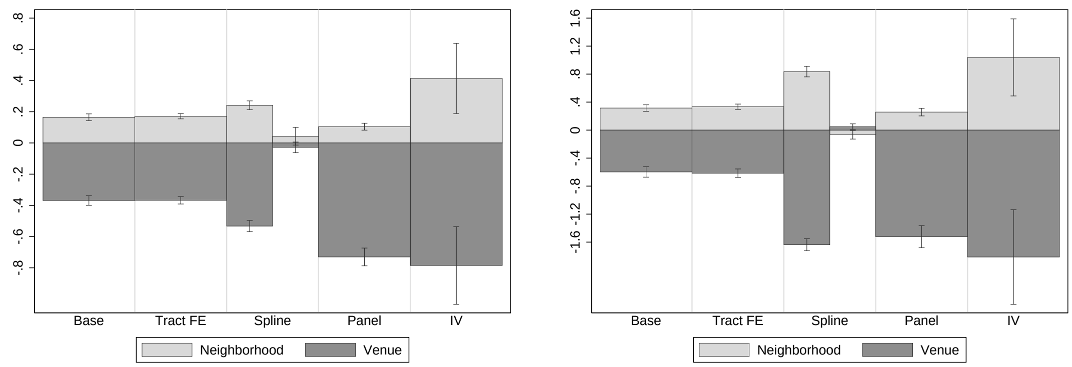

(a) V_b = Num. of Categories

(b) V_b = Num. of Subcategories

Notes: The dark shaded bars represent $\hat{\beta}^V$, and the light shaded bars represent $\hat{\beta}^N$. 95% confidence intervals are also shown from robust standard errors clustered at the block group level. The first bars correspond to baseline estimates from equations (4) and (5). The second bars replace the block group fixed effects in the baseline estimates with tract fixed effects. The third set of bars correspond to estimates of the parameters specified as a linear b-spline with a knot at 3 subcategories. The fourth bars correspond to estimates from equations (6) and (7) where the dataset is disaggregated to a monthly panel, and the block group fixed effects are replaced with block fixed effects. The fifth bars correspond to 2SLS estimates of the baseline regressions with zoning instruments.

In our first robustness check, we re-estimate equations (4) and (5) with tract fixed effects instead of block group fixed effects. Tracts typically encompass two or more block groups, so these fixed effects no longer control for unobserved amenities varying across block groups within tracts, which may confound our estimates. The results (denoted as "Tract FE") are virtually unchanged, which suggests that after controlling for X_b and R_b, amenities and local demand varying across block groups within tract are uncorrelated to V_b. It is difficult to conceive of unobservables that are correlated to V_b, that vary across blocks within block groups but do not vary across block groups within tracts.[32]

Second, we re-estimate equations (4) and (5) using linear B-splines in V_b, which allows us to estimate separate marginal effects of venue variety on diversity for neighborhoods with three or fewer subcategories and for neighborhoods with four or more subcategories. If $\hat{\beta}^V$ and $\hat{\beta}^N$ are causal estimates, then they will likely decline in magnitude as we compare nearby blocks with higher levels of venue variety.[33] In contrast, if these estimates reflect confounding factors that are

[32]For instance, simultaneity could only be a concern if venues had more control or preference over their choice of which block within a block group to locate relative to their choice of which block group within a tract to locate.

[33]Extending the intuition of the model presented above, in a block with greater number of venues with distinct

present irrespective of the level of V_b, then we should find that these effects do not decline for higher V_b. Indeed, we find that nearly all of these effects (denoted as "Spline") operate at low levels of venue variety in all cities in our sample.[34]

Third, we exploit the longitudinal variation in our data to estimate β^V and β^N using an alternative identification strategy. We re-specify equations (4) and (5) as

$$D_{bt}^V = \beta^V V_{bt} + \alpha_b^V + \alpha_{ct}^V + X_{bt}\delta^V + \epsilon_{bt}^V \quad (6)$$

$$D_{bt}^N = \beta^N V_{bt} + \alpha_b^N + \alpha_{ct}^N + X_{bt}\delta^N + \epsilon_{bt}^N, \quad (7)$$

respectively. The key difference is that all of our main explanatory variables and controls now vary by month. By doing so, we can identify β^V and β^N using only within-block variation in venue variety that arises due to the entry and exit of venues over time. We implement this identification strategy by including block fixed effects (α_b^V and α_b^N) that additionally control for all unobserved determinants of diversity that vary across blocks within block groups that were not controlled for in equations (4) and (5). The fixed effects α_{ct}^V and α_{ct}^N control for city level amenities that may vary by month in order to absorb any seasonality that varies across cities. Our results (denoted as "Panel") suggest that our baseline estimates of β^V are conservative, which is consistent with our sensitivity analysis in the appendix.

Finally, we re-estimate β^V and β^N in equations (4) and (5) with a third, distinct identification strategy that uses variation in zoning laws across blocks within block groups as instrumental variables (IVs) for venue variety. We only use identifying variation in the variety of venues that stems from regulations that restrict the location of certain venues in certain blocks. This IV approach deals with any remaining simultaneity concerns and any remaining confounders that are uncorrelated to zoning laws such as most kinds of measurement error. Specifically, we use the share of lots in the block that are zoned to residential, commercial and mixed uses as instruments; hence, we effectively compare diversity in nearby blocks that are zoned differently and thus have different

x_j's, more of the support will be covered by venue visitors. As a result, a marginal increase in venue variety will have a smaller effect on both D_b^V and D_b^N since the additional venue will draw increasingly from individuals who were otherwise planning to go to another venue on the same block.

[34]These results are virtually unchanged when we place the knot at $2, ..., 5$ subcategories.

levels of venue variety (but a similar number of venues, overall traffic and number of residents of each gender).[35]

Differences in zoning laws are found to generate differences in the variety of venues in nearby census blocks. In Figure 8, we spatially illustrate the "first-stage" relationship between commercial zoning (here categorized in quartiles for visual clarity) and venue variety (number of unique venue subcategories) in Manhattan census blocks, which is clearly positive. More formally, a joint F-test of the significance of the three instruments for the number of unique venue categories and unique venue subcategories yields $F_{3,5779} = 25.00$ (0.00) and $F_{3,5779} = 12.27$ (0.00), respectively, where the p-values shown in parentheses are much smaller than 0.01. Our estimates (denoted as "IV") are, if anything, larger in magnitude than all OLS estimates, which suggests that the OLS estimates may be attenuated by measurement error. As a result, our findings that $\beta^V < 0$ and $\beta^N > 0$ should be understood to be conservative. This interpretation is consistent with our Monte Carlo study of measurement error described in the appendix.[36]

Figure 8: Commercial Zoning and Venue Variety (First Stage)

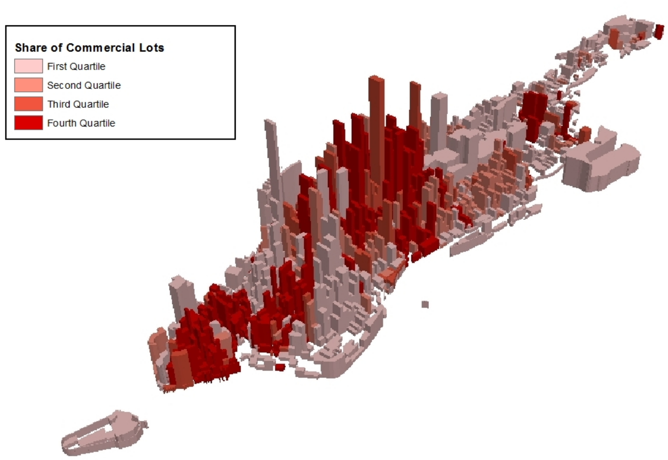

Notes: Each bar represents a census block in Manhattan. The height of each bar corresponds to V_b, the number of unique venue subcategories in b. Darker bars represent blocks with a greater proportion of commercially zoned lots.

[35]We obtained lot level data on zoning for each city from their respective planning offices. Lots can be zoned for other uses than the three that we use for IVs (e.g., manufacturing or parks), but our results were unchanged when using additional IVs.

[36]In order to ensure that $\hat{\beta}^V$ was not contaminated by the effect β^N and vice versa, we also implemented a robustness check where we added D_b^N as a control variable in the equation of D_b^V (equation (4)), and D_b^V as a control variable in the equation of D_b^N (equation (5)). Our results were unchanged.

5 Gender Homophily and the Labor Force Participation Gap

Social networks and the interactions that they facilitate have the potential to shape labor markets profoundly (Montgomery (1991)). One of the most studied features of the labor market is the persistent gender gap in labor force participation that is observed in many countries over most time periods. A considerable literature has analyzed the impacts of job referral networks on individual employment and wages (e.g., Ioannides and Loury (2004); Bayer et al. (2008); Schmutte (2015)) and found that social interactions play an important role in explaining labor market outcomes. In particular, Bayer et al. (2008) (hereafter, BRT) find that these networks can operate at a highly localized level: residents of the same block form stronger network ties than residents of nearby blocks, particularly when individuals are similar on observable characteristics (ostensibly due to homophily). We complement their analysis with causal evidence that their findings also extend to interactions between block residents and block visitors as mediated by the gender homophily studied in this paper. Specifically, we show that the gender diversity in venues is a determinant of the employment gap. This could be the case if people's job referral networks were developed at least partially through interactions in venues. Male and female residents of less venue-diverse blocks would tend to interact with each other less frequently. This would result in more disjoint referral networks by gender, which might in turn widen the gender gap.

Our identification strategy compares the gender gap in labor force participation among residents of otherwise identical, nearby census blocks that differ only in the diversity of their visitors at both the block and venue levels. While block diversity might proxy for the exposure to diversity outside of venues in the block, venue diversity proxies for the exposure to diversity inside venues. Hence, differences in this gender gap can be understood to be mediated through social interactions between block residents and venue visitors either inside or outside of venues. We obtain block-level employment statistics from the Longitudinal Employer-Household Dynamics (LEHD) Origin-Destination Employment Statistics, or LODES. For each block in our sample, we observe the numbers of male and female residents who are employed[37] in each year from 2012 to 2013.[38] These statistics are further disaggregated into low wage (less than \$1,250 per month), medium wage (between \$1,250 and \$3,333 per month) and high wage (at least \$3,333 per month) groups. For our measure of the gender gap in labor force participation, we construct $GAP_{wbt} = \frac{M_{wbt}}{MR_b} - \frac{F_{wbt}}{FR_b}$ where M_{wbt} and F_{wbt} are the numbers of male and female employees in wage group w in block b in year t from LODES, and MR_b and FR_b are the adult male and female populations of block b from the 2010 Census. We normalize GAP_{wbt} to have mean 0 and variance 1 across all blocks in the sample.[39]

We estimate the effects of venue and neighborhood diversity on the gender gap in the following regression equation:

$$GAP_{wbt} = \gamma_w^V D_b^V + \gamma_w^N D_b^N + \alpha_{wg} + X_b \delta_w + \epsilon_{wbt} \tag{8}$$

where D_b^V and D_b^N are defined as before, α_{wg} is a block group-wage group fixed effect, X_b is a vector of controls and ϵ_{wbt} is an error term. The parameters of interest, γ_w^V and γ_w^N, represent the effects of venue and neighborhood diversity, respectively, on the female labor force participation gap among workers in wage group w. The OLS estimates of these parameters can be interpreted as causal effects only if $\text{Cov}\left(\epsilon_{wbt}, D_b^V | D_b^N, \alpha_{wg}, X_b\right) = 0$ and $\text{Cov}\left(\epsilon_{wbt}, D_b^N | D_b^V, \alpha_{wg}, X_b\right) = 0$.

Our identification assumption is that residents do not sort by gender across blocks within a given block group. This assumption is inspired by BRT, who make a similar assumption and provide strong evidence of no residential sorting within block groups along any observable demographic dimension (including gender).[40] As BRT point out, housing markets are very thin at small geographic scales,

[37]The LODES data cover approximately 95 percent of wage and salary jobs, excluding a small number of employees in the military, security-related federal agencies, postal workers, employees at non-profit and religious institutions, informal workers and the self-employed (Graham et al. (2014)). Because we do not want to identify referral effects for part-time employment, which may operate along different social networks, we focus our analysis on the subsample of primary jobs; however, all results are qualitatively similar in the sample of all jobs, and in a subsample of only private-sector jobs.

[38]We use two years of data from LODES even though our measures of diversity do not vary over time because these measures are constructed from Foursquare check-in data that spans parts of both 2012 and 2013.

[39]Ideally, we would calculate the gender gap as $\widetilde{GAP}_{wbt} = \frac{M_{wbt}}{MR_{wbt}} - \frac{F_{wbt}}{FR_{wbt}}$, where MR_{wbt} and FR_{wbt} are the numbers of male and female residents in block b in year t who are a match for jobs of group w. Unfortunately, we cannot observe MR_{wbt} and FR_{wbt}. We conduct several robustness checks below to allay concerns that this measurement error (i.e., $\widetilde{GAP}_{wbt} \neq GAP_{wbt}$) biases our results.

[40]Our assumption is weaker than the one made in BRT as the outcome variable of our analysis is the gender *gap* in labor force participation rather than the labor force participation rate itself. For instance, our estimates are robust to residential sorting within block groups based on the propensity of employment provided that individuals of both

which restricts individuals' ability to choose to live in specific census blocks, and, even if they were able to, individuals may find it difficult to observe highly local (block level) amenities at the time of their residential decision. Under this identifying assumption, the parameters γ_w^V and γ_w^N can be interpreted as the causal effects of diversity among visitors to block b on the labor force participation gap among residents of block b. We posit that this effect occurs because residents are more likely to interact with visitors in their block than with visitors in nearby blocks. As discussed in BRT, to the extent that residents of a given block also interact with visitors of nearby blocks, our estimates of γ_w^V and γ_w^N will likely be attenuated.

We estimate equation (8) with 11 distinct specifications of controls (X_b), neighborhood fixed effects (α) and estimation subsamples to establish the robustness of our identification strategy. Results are presented in Table 4. Each row corresponds to a different specification. Before describing the results individually, we highlight two robust findings. First, greater gender diversity *inside venues* reduces the employment gap for low wage jobs; a one standard deviation increase in venue diversity shrinks this gap by roughly 1.5 percent of a standard deviation (see Remark 3 for interpretation). However, we do not find any evidence that venue diversity affects the employment gap for medium and high wage jobs despite obtaining consistently negative point estimates.[41] This is consistent with other findings that job referral effects are stronger for individuals who are less attached to the labor force (see Ioannides and Loury (2004) for a survey). Second, we do not find any statistically significant effects of gender diversity outside of venues on employment gaps for any type of job. Moreover, the point estimates for γ_w^N tend to be much smaller than the point estimates for γ_w^V. This suggests that social interactions inside venues are more important for reducing gender gaps in labor force participation than the social interactions that might take place on streets and sidewalks, at least among the lowest paid workers.

In specification (1), we specify only covariates from our Foursquare data in order to control for differences in venue density and check-in intensities of males and females across blocks. In specification (2), we control for the number of male and female residents on each block. Our estimates are essentially unchanged, which supports our assumption that there is no systematic sorting by gender

genders sort uniformly.

[41]Although we find relatively large estimates of γ_{high}^V and γ_{high}^N, these estimates are never statistically significant at the 10% level. As a robustness check that these can be plausibly interpreted as zero-estimates, when we estimate our most precise model (specification (11)) the point estimates decrease substantially. Further supporting evidence is provided in the appendix.

across blocks within block groups. In specification (3), we specify all control variables from specification (2) flexibly using cubic B-splines.[42] Doing so substantially improves the fit of the model, and we estimate slightly larger effects of venue diversity on low wage employment gaps. Furthermore, specifying MR_b and FR_b flexibly ensures that our use of them in constructing GAP_{wbt} does not introduce spurious correlation into the regression. The fact that as we better control for resident characteristics, our estimates of γ_{low}^V become larger in magnitude suggests that unobservables due to residential sorting might, if anything, bias our estimate of γ_{low}^V downward. In specification (4), we similarly control for the number of young and old residents on each block and find no changes in our estimates, which further supports our assumption of no residential sorting by gender within block groups.[43] In principle, our identifying assumption allows for sorting across blocks along other dimensions as long as this sorting occurs uniformly by gender. To check for whether sorting along other dimensions might bias our results, in specification (5) we control for other demographic characteristics of workers (race, Hispanic ethnicity and level of education). Although the fit of the model slightly improves, all coefficient estimates are unchanged, as expected. As a final test for the presence of sorting within block groups, we re-estimate the regression with *tract*-wage group fixed effects instead of block group-wage group fixed effects in specification (6). Any bias due to sorting within block groups in the baseline specification should be further exacerbated in this specification since sorting across block groups within the same tract is no longer controlled for. The fact that our estimates in specification (6) are smaller suggests that our baseline results might actually be conservative. This result corroborates our intuition when comparing the results from specifications (2) and (3).

Because of the linkages we found in Section 4, one might worry that we have only identified the effect of venue variety rather than gender diversity in venues. For instance, a lack of venue variety might generate more venue diversity and less neighborhood diversity in gender, age, race, and other demographic characteristics, all of which could plausibly influence the gender gap in employment. To address this concern, we include block level measures of age diversity in neighborhoods and venues (i.e., age analogs of D_b^N and D_b^V) as controls in specification (7). Our estimates decrease

[42]We use as many knots as possible for each control variable: 7 equally spaced knots for the numbers of male and female visitors and residents (24 covariates) and 3 equally spaced knots for the number of venues (2 covariates).

[43]This is also suggestive that measurement error in GAP_{wbt} does not bias our results, as older residents are likely a better match for jobs in higher wage groups. See footnote **??**.

only slightly, and these two variables do not improve the fit, which supports our claim that the effects we identify are due to gender sorting itself. Because we found that venue variety affects both venue and neighborhood diversity, we drop D_b^N from the model in specification (8). Our estimates are unchanged, which supports our claim that the effects we identify are mediated by gender diversity inside venues.

In specification (9), we restrict our sample by dropping all blocks in which at least one resident works on his or her residential block in order to gauge whether our results are primarily driven by residents who work extremely close to home (e.g., workers who live above the store that they own). If anything, our estimates slightly increase, suggesting that this is not the case. In specification (10), we re-estimate the model with block group-*year*-wage group fixed effects. Our results are unchanged, and the adjusted R^2 is substantially reduced, which suggest that pooling data from multiple years could be useful as it will only improve the precision of our estimates. Accordingly, in specification (11), we augment our sample by pooling LODES data on employment gaps from 2010-2013, which extends the sample period from two to four years. Although our explanatory variables do not vary over time, this should still increase our statistical power. Indeed, all standard errors decrease. Our main finding of $\hat{\gamma}^V_{\text{low}} < 0$ is unchanged, but we now obtain more precise zero-estimates for all effects on medium and high wage jobs.[44]

Remark 3. To better interpret our results, we find that a one standard deviation increase in venue diversity decreases the gender gap in labor force participation by 0.21 percentage points. For context, BRT find that a one standard deviation increase in what they define as the average "match quality" of neighborhood residents across several demographic variables reduces the gender gap in labor force participation by 0.9 percentage points. Our estimates and BRT estimates reflect the impact of similar social interactions that differ for only two reasons: (1) we allow for interactions between residents and visitors, where these visitors may or may not be residents while BRT allows only for interactions between residents, which should amplify the effect, and (2) we focus on interactions that happen only inside venues in the same block, while BRT focuses on interactions that might occur in any environment, which should diminish the effect.

[44]For completeness, we present estimates of specifications (1)-(10) using the 2010-2013 sample from LODES in the appendix.

Table 4: Effects of Diversity on the Labor Force Participation Gender Gap

	Specification	Low Wage Jobs γ^V_{low}	γ^N_{low}	Medium Wage Jobs $\gamma^V_{\text{med.}}$	$\gamma^N_{\text{med.}}$	High Wage Jobs γ^V_{high}	γ^N_{high}	N	Adj. R^2
(1)	Controls from Foursquare Data[1]	-0.013*	-0.001	-0.006	0.008	-0.016	0.020	76,236	0.159
		(0.007)	(0.009)	(0.009)	(0.009)	(0.015)	(0.015)		
(2)	(1) + residential gender controls[2]	-0.013*	-0.002	-0.006	0.007	-0.015	0.017	76,236	0.188
		(0.007)	(0.009)	(0.009)	(0.009)	(0.014)	(0.014)		
(3)	(2) + flexibly specified controls [3]	-0.015**	-0.002	-0.003	0.001	-0.016	0.013	76,236	0.325
		(0.007)	(0.007)	(0.008)	(0.008)	(0.012)	(0.012)		
(4)	(3) + flexible residential age controls[4]	-0.015**	-0.003	-0.003	0.000	-0.015	0.013	76,236	0.325
		(0.007)	(0.007)	(0.008)	(0.008)	(0.012)	(0.012)		
(5)	(4) + workers' demographics characteristics[5]	-0.015**	-0.002	-0.004	0.001	-0.016	0.012	76,236	0.330
		(0.007)	(0.007)	(0.008)	(0.008)	(0.012)	(0.012)		
(6)	(4) + tract-wage group FEs instead of block group-wage group FEs	-0.010**	-0.003	-0.005	0.006	-0.006	0.003	76,236	0.310
		(0.005)	(0.005)	(0.006)	(0.006)	(0.009)	(0.009)		
(7)	(4) + D^V_b and D^N_b for age[6]	-0.012**	-0.001	-0.001	-0.003	-0.013	0.011	76,236	0.325
		(0.007)	(0.007)	(0.008)	(0.009)	(0.012)	(0.013)		
(8)	(4), drop D^N_b	-0.015**	–	-0.003	–	-0.014	–	76,236	0.325
		(0.007)		(0.008)		(0.011)			
(9)	(4), drop blocks where residents also work[7]	-0.017**	-0.005	-0.009	0.003	-0.016	0.009	63,146	0.298
		(0.008)	(0.008)	(0.009)	(0.009)	(0.013)	(0.014)		
(10)	(4) + block group-year-wage group FEs instead of block group-wage group FEs	-0.016*	-0.003	0.000	0.001	-0.015	0.012	76,236	0.262
		(0.008)	(0.009)	(0.009)	(0.010)	(0.014)	(0.015)		
(11)	(4), $t \geq 2010$	-0.014**	-0.007	0.001	-0.000	-0.001	0.005	152,335	0.328
		(0.005)	(0.005)	(0.005)	(0.006)	(0.007)	(0.008)		

Notes: Low wage jobs pay less than $1,250 monthly, medium wage pay jobs pay between $1,250 and $3,333 monthly, and high wage pay jobs pay more than $3,333 monthly. All specifications include block group - wage group fixed effects, with the exceptions of (6) and (10). Robust standard errors clustered at the block level are presented in parentheses.
[1] : number of venues in block, and numbers of female and male block visitors (3 covariates for each group); [2] : add to (1) the numbers of female and male block residents (2 additional covariates for each group);[3] : Cubic B-spline (with as many knots as possible) of all controls in (2) (26 covariates for each group). [4] : add to (3) cubic B-spline of numbers of younger ($\leq$ 35) and older ($>$ 35) block residents (12 additional covariates for each group). [5] : Add to (4) the numbers of block workers who are White, Black, other (non-White and non-Black), Hispanic, non-Hispanic, college graduates, college non-graduates (7 additional covariates for each group). [6] : Add to (4) D^{Vy}_b and D^{Ny}_b, which are the analogous measures of D^V_b and D^N_b based on the proportion of younger visitors ($\leq$ 35), rather than based on the proportion of female visitors (2 additional covariates for each group). [7] : Drop observations from blocks with at least one resident who works in the same block. **: 5% significance level, *: 10% significance level

6 Discussion

Our main findings can be broadly summarized as follows: (1) gender homophily is quite pervasive and manifests mainly at levels that are unobserved in most datasets; (2) when different individuals are offered a greater variety of options to choose from, they sort across these options more intensely, which leads to greater gender segregation (lower exposure to gender diversity); and (3) gender homophily leads to greater gender gaps in the labor force participation rates among low wage residents. We discuss the external validity of these findings in three directions.

6.1 Do our findings extend to other environments?

The extensive gender segregation that we measure, which persists down to the venue level, is suggestive of further homophily *within* venues and activities. For instance, men and women may be inclined to sort to different tables within cafes and bars, or engage in different activities within gyms and parks. Hence, we believe that our findings should be understood as conservative estimates of the actual amount of gender segregation in the day-to-day activities of men and women. Moreover, the fact that gender segregation is observed across a wide variety of different recreational and commercial activities suggests that it may also pervade other social settings such as classrooms and workplaces.

Although patterns of venue visitors differ between weekdays and weekends and across cities, our general findings do not. All of our results are similar for each city-day of week combination (see appendix). Similarly, patterns of venue visitors exhibit seasonality, which is city dependent (e.g., many fewer people visit Chicago parks in the winter months relative to the summer, but this seasonal effect is much weaker in Los Angeles). However, all of our general findings are similar for each city-month of year combination (see appendix). We view the robustness of these results as suggestive that they might also hold in other urban and suburban environments. Of course, more research is needed to understand such additional sources of heterogeneity (e.g., in which venues do peer groups form most effectively, and what are the consequences of those peer groups?).

6.2 Do similar patterns of homophily operate along other demographic dimensions?

Our Foursquare data allows us to answer this question along only one additional dimension: age. For each venue in our sample, we observe the daily numbers of check-ins from users under 35 years of age and from users 35 years of age or older. With this information, we replicate our entire analysis, substituting for the proportion of females the proportion of youth. Our results are broadly similar to our results on gender, which is not a trivial finding given that gender and age are largely uncorrelated. Although we find roughly half as much age segregation as we do gender segregation, it occurs highly locally as shown in Table 5: from a third to about half of all age sorting in cities occurs within census blocks. As in the case of gender, we find that age homophily is primarily mediated by the fact that people of different ages prefer different activities. Finally, we also find that the causal effects of venue variety on venue and neighborhood age diversity are both qualitatively and quantitatively similar to the respective effects on gender diversity (Figure 9). A full reporting of all results from this replication is provided in the appendix.

Table 5: Venue Sorting Within Neighborhoods By Age

	Proportion of city-wide segregation attributable to homophily within:		
	Tracts	Block Groups	Blocks
Atlanta	0.75	0.68	0.45
Chicago	0.73	0.63	0.36
Dallas	0.76	0.68	0.45
Los Angeles	0.75	0.67	0.43
New York City	0.87	0.81	0.70
Philadelphia	0.68	0.61	0.34
San Francisco	0.83	0.76	0.53
Washington, DC	0.80	0.74	0.47

Note: Bootstrapped standard errors for all Theil indices in all cities are less than 0.005 and are omitted for clarity.

Figure 9: $\hat{\beta}^V$ and $\hat{\beta}^N$ For Age By City

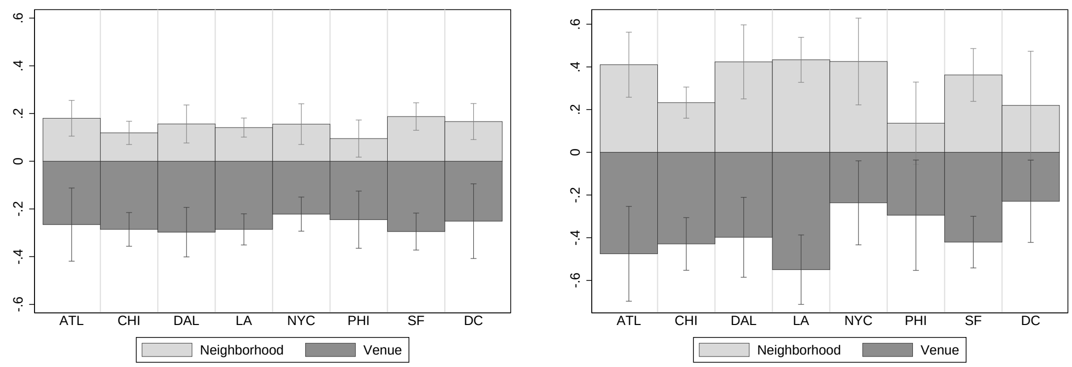

(a) V_b = Num. of Categories

(b) V_b = Num. of Subcategories

Notes: The dark bars represent estimates of $\hat{\beta}^V$ from equation (4), and the light bars represent estimates of $\hat{\beta}^N$ from equation (5). 95% confidence intervals are also shown from robust standard errors clustered at the block group level. The number of observations for each of the 16 regressions is equal to the number of census blocks in each city (see Table 1), and the R^2 of each regression varies from 0.23 to 0.52.

Although we cannot replicate our analysis along any other demographic dimensions, we conjecture that the robustness of our results across gender and age may be suggestive of similar patterns of homophily in day-to-day activities along other dimensions such as race and income. Indeed, there is reason to conjecture that hidden racial and income segregation may be even greater than the passive gender and age segregation that we observe in our data. For instance, active racial segregation is likely to contribute to a greater level of racial segregation over and above that which would be implied by passive racially homophilic forces, and venue prices are likely an additional contributor to income segregation.

In a broader sense, individuals sort along multiple demographic dimensions simultaneously. For instance, younger women are plausibly more likely to frequent the same venues as other young women than older women are. This suggests that the true amount of diversity to which individuals are exposed is even more attenuated by endogenous sorting than what we are able to observe, which is an additional reason why our findings on gender and age should be understood to be conservative estimates of the true levels of segregation in peer groups.[45]

Remark 4. Our findings of highly localized gender and age segregation complement the findings

[45] Gender and age are mostly uncorrelated to other characteristics, making them in some ways ideal candidates to provide a conservative conclusion on such inevitably incomplete analysis.

by other researchers of similarly localized segregation along other dimensions. For example, Carrell et al. (2013) document highly localized (within Air Force Academy squadron) segregation by student ability. Among entrepreneurs, Ruef et al. (2003) document segregation along a variety of "status-related dimensions" such as gender, ethnicity and professionalism. Kossinets and Watts (2009) analyze how segregation across a variety of demographic dimensions locally evolves in the university setting along different courses of study and residential choices. And Currarini et al. (2009) document substantial, highly localized (within school) segregation by ethnicity in high school friendships.

6.3 Do gender and age homophily affect other outcomes?

In Section 5 we attempted to identify an explicit causal link between gender local segregation in venues and local labor force participation gaps. This, however, is unlikely to be the only outcome that is impacted by gender homophily. A large body of research has found that exposure to female peers affects, for example, corporate governance and performance (e.g., Brown et al. (2002); Adams and Ferreira (2009)), student achievement (Hoxby (2000); Lavy and Schlosser (2011); Hill (2015)), substance abuse (Andrews et al. (2002)), the expression of political beliefs (Huckfeldt (1995)), and the level of intimacy in social networks (Verbrugge (1977)). Although peer effects with respect to age have not been widely studied, the systematically different beliefs that people of different ages may hold suggests that age homophily might play a role in the shaping of political preferences and the development of human capital by social interaction. Although we do not claim that all of these social interactions occur at all of the different types of venues that we are able to observe in our data, we believe it is plausible that repeated exposure to certain peers in venues may accumulate over time, in turn affecting peoples' beliefs, preferences, social norms, and actions. The identification of these various effects is a difficult proposition that carries heavy data demands and lies beyond the scope of this paper.

7 Conclusion

Peer groups shape the social environment in which we live. Homophily leads similar people to associate with one another, and we find that the amount of it that is commonly observed in datasets might only represent the tip of the iceberg when it comes to the actual extent of everyday homophily in people's lives. Using novel, user-generated data from Foursquare, a popular mobile app, we analyze how individuals sort into neighborhoods and further into venues in eight major US cities. We find that individuals sort by gender and by age across venues that are extremely close to each other and at a similar intensity in a variety of different city types, from the long established, dense, urban cores of New York City and Philadelphia to newer and more diffuse urban areas such as Los Angeles, Dallas and Atlanta. This lends some universality to the widespread, homophilic, endogenous peer group formation that we observe.

Our results echo the central themes of Jacobs (1961): individuals endogenously respond to the urban landscape around them, and it is the diversity of this landscape that gives rise to social interactions. However, they also invite a reassessment of whether mixed-use development in neighborhoods coupled with demographic density, which Jacobs and others have championed, are important ingredients for diversity to emerge. While we find that the resulting variety in the types of venues will lead to more overall diversity in neighborhoods, we also find that it will lead to *less* diversity at the venue level as similar individuals are able to more intensely segregate themselves into venues. Hence, strengthening the social interactions that form the basis for thriving communities, especially in dense, urban neighborhoods, may be a more complicated task for policymakers to achieve than previously thought.

These social interactions seem important. Gender homophily and the resulting decrease in diversity seems to contribute to the gender gap in labor force participation, which suggests that homophily at very fine levels in important dimensions such as gender, age, race, income, and political beliefs may impact society more broadly. For instance, our results relate closely to the recent debate on how well cities can offer exposure to a diversity of opinions that might be crucial for the formation of accurate and pro-social beliefs. If similar people tend to hold similar views, then homophily might impact the diversity of opinions to which they are exposed. On the one hand, Sunstein (2009) suggests that physical interactions in neighborhoods and in venues might be an important source of exposure to diverse views.[46] On the other hand, Gentzkow and Shapiro (2011) find that news media (both online and offline) offer more exposure to diverse opinions than neighbors, co-workers and family members do. Our findings help reconcile these two positions: physical interaction may well be a crucial source of exposure to diverse opinions, but most people choose not to be exposed

[46]"The diverse people who walk the streets and use the parks are likely to hear speakers' arguments; they might also learn about the nature and intensity of views held by their fellow citizens. (...) When you go to work or visit a park (...) it is possible that you will have a range of unexpected encounters" (p. 30).

to such diversity, even if inadvertently. They just tend to be drawn to the same activities as other, similar people.

More broadly, the formation of peer groups is a deeply personal choice. Although it is certainly affected by where people live, study and work, people make many smaller decisions on a daily basis that can shape their social environments in profound ways. These might revolve around seemingly insignificant actions such as frequenting a specific venue, making an acquaintance, or joining a conversation, any of which may turn out to be memorable and impactful. While the informal and personal nature of these decisions makes them difficult to observe in standard data sets, the proliferation of user-generated data sets has the potential to offer researchers a window into this rich source of socialization. We view this work as an early step along that path.

References

Adams, R. B., Ferreira, D., 2009. Women in the boardroom and their impact on governance and performance. Journal of financial economics 94 (2), 291–309.

Altonji, J. G., Blank, R. M., 1999. Race and gender in the labor market. Handbook of labor economics 3, 3143–3259.

Andrews, J. A., Tildesley, E., Hops, H., Li, F., 2002. The influence of peers on young adult substance use. Health psychology 21 (4), 349.

Arcidiacono, P., Nicholson, S., 2005. Peer effects in medical school. Journal of public Economics 89 (2), 327–350.

Arribas-Bel, D., Bakens, J., 2014. "the magic's in the recipe": Urban diversity and popular amenities. mimeo.

Atkinson, A. B., 1970. On the measurement of inequality. Journal of economic theory 2 (3), 244–263.

Banzhaf, H. S., Walsh, R. P., 2013. Segregation and tiebout sorting: The link between place-based investments and neighborhood tipping. Journal of Urban Economics 74 (0), 83 – 98.

Bayer, P., Ross, S. L., Topa, G., 2008. Place of work and place of residence: Informal hiring networks and labor market outcomes. Journal of Political Economy 116 (6), 1150–1196.

Becker, G. S., 1985. Human capital, effort, and the sexual division of labor. Journal of labor economics, S33–S58.

Bobo, L., Charles, C., Krysan, M., Simmons, A., Fredrickson, G., 2012. The real record on racial attitudes. Social Trends in American Life: Findings from the General Social Survey since 1972, 38.

Bourguignon, F., 1979. Decomposable income inequality measures. Econometrica 47 (4), 901–920.

Boustan, L., 2012. School desegregation and urban change: Evidence from city boundaries. American Economic Journal: Applied Economics 4 (1), 85–108.

Brown, D. A., Brown, D. L., Anastasopoulos, V., 2002. Women on boards: Not just the right thing... but the" bright" thing. Conference Board of Canada.

Bruch, E., Mare, R., 2006. Neighborhood Choice and Neighborhood Change. American Journal of Sociology 112 (3), 667–708.

Caetano, G., Maheshri, V., 2014. School segregation and the identification of tipping behavior. Mimeo.

Card, D., Mas, A., Rothstein, J., 2008a. Tipping and the Dynamics of Segregation. Quarterly Journal of Economics 123 (1), 177–218.

Carrell, S. E., Sacerdote, B. I., West, J. E., 2013. From natural variation to optimal policy? the importance of endogenous peer group formation. Econometrica 81 (3), 855–882.

Castilla, E. J., 2005. Social networks and employee performance in a call center1. American journal of sociology 110 (5), 1243–1283.

Chetty, R., Hendren, N., 2015. The impacts of neighborhoods on intergenerational mobility: Childhood exposure effects and county-level estimates. Tech. rep., Working Paper, May 20.

Couture, V., 2014. Valuing the consumption benefits of urban diversity. mimeo.

Currarini, S., Jackson, M. O., Pin, P., 2009. An economic model of friendship: Homophily, minorities, and segregation. Econometrica 77 (4), 1003–1045.

Davis, D. R., Dingel, J. I., Monras, J., Morales, E., 2014. Spatial and social frictions in the city: Evidence from yelp. mimeo.

Echenique, F., Fryer Jr, R., Kaufman, A., 2006. Is school segregation good or bad? The American economic review, 265–269.

Ellison, G., Glaeser, E. L., 1997. Geographic concentration in us manufacturing industries: A dartboard approach. Journal of Political Economy 105 (5), 889–927.

Even, W. E., Macpherson, D. A., 1993. The decline of private-sector unionism and the gender wage gap. Journal of Human Resources, 279–296.

Fernández, R., Fogli, A., Olivetti, C., 2004. Mothers and sons: Preference formation and female labor force dynamics. The Quarterly Journal of Economics, 1249–1299.

Fernandez, R. M., Castilla, E. J., Moore, P., 2000. Social capital at work: Networks and employment at a phone center. American journal of sociology, 1288–1356.

Gentzkow, M., Shapiro, J. M., 2011. Ideological segregation online and offline. The Quarterly Journal of Economics 126 (4), 1799–1839.

Glaeser, E., Sacerdote, B., Scheinkman, J., 1996. Crime and social interactions. The Quarterly Journal of Economics, 507–548.

Glaeser, E. L., Kolko, J., Saiz, A., 2001. Consumer city. Journal of economic geography 1 (1), 27–50.

Goldin, C., 2014. A grand gender convergence: Its last chapter. The American Economic Review 104 (4), 1091–1119.

Goldin, C., Polachek, S., 1987. Residual differences by sex: Perspectives on the gender gap in earnings. The American Economic Review, 143–151.

Graham, M. R., Kutzbach, M. J., McKenzie, B., Oct. 2014. Design Comparison of LODES and ACS Commuting Data Products. Working Papers 14-38, Center for Economic Studies, U.S. Census Bureau.

Hill, A. J., 2015. The girl next door: The effect of opposite gender friends on high school achievement. American Economic Journal: Applied Economics 7, 147–77.

Hotelling, H., 1929. Stability in competition. The Economic Journal 39 (153), 41–57.

Hoxby, C., 2000. Peer effects in the classroom: Learning from gender and race variation. Tech. rep., National Bureau of Economic Research.

Huckfeldt, R. R., 1995. Citizens, politics and social communication: Information and influence in an election campaign. Cambridge University Press.

Ioannides, Y. M., Loury, L. D., 2004. Job information networks, neighborhood effects, and inequality. Journal of Economic Literature 42 (4), pp. 1056–1093.

Jacobs, J., 1961. The death and life of great American cities. Random House LLC.

Kling, J. R., Liebman, J. B., Katz, L. F., 2007. Experimental analysis of neighborhood effects. Econometrica 75 (1), 83–119.

Kossinets, G., Watts, D. J., 2009. Origins of homophily in an evolving social network1. American journal of sociology 115 (2), 405–450.

Ladd, G. W., 1999. Peer relationships and social competence during early and middle childhood. Annual review of psychology 50 (1), 333–359.

Lavy, V., Schlosser, A., 2011. Mechanisms and impacts of gender peer effects at school. American Economic Journal: Applied Economics 3 (2), 1–33.

Mammen, K., Paxson, C., 2000. Women's work and economic development. The Journal of Economic Perspectives, 141–164.

Mas, A., Moretti, E., 2009. Peers at work. The American economic review 99 (1), 112–145.

McPherson, M., Smith-Lovin, L., Cook, J. M., 2001. Birds of a feather: Homophily in social networks. Annual review of sociology, 415–444.

Montgomery, J. D., 1991. Social networks and labor-market outcomes: Toward an economic analysis. The American economic review, 1408–1418.

Montgomery, J. D., 1992. Job search and network composition: Implications of the strength-of-weak-ties hypothesis. American Sociological Review, 586–596.

Oreopoulos, P., 2003. The long-run consequences of living in a poor neighborhood. The quarterly journal of economics, 1533–1575.

Pan, J., 2015. Gender segregation in occupations: The role of tipping and social interactions. Journal of Labor Economics 33 (2), 365–408.

Ruef, M., Aldrich, H. E., Carter, N. M., 2003. The structure of founding teams: Homophily, strong ties, and isolation among us entrepreneurs. American sociological review, 195–222.

Sacerdote, B., 2001. Peer effects with random assignment: Results for dartmouth roommates. Quarterly Journal of Economics 116 (2), 681–704.

Schelling, T. C., 1969. Models of Segregation. The American Economic Review 59 (2), 488–493.

Schelling, T. C., 1971. Dynamic Models of Segregation. Journal of Mathematical Sociology 1, 143–186.

Schmutte, I. M., 2015. Job referral networks and the determination of earnings in local labor markets. Journal of Labor Economics 33 (1).

Shorrocks, A. F., 1980. The class of additively decomposable inequality measures. Econometrica: Journal of the Econometric Society, 613–625.

Sunstein, C. R., 2009. Republic. com 2.0. Princeton University Press.

Theil, H., 1967. Economics and information theory. North-Holland.

Verbrugge, L. M., 1977. The structure of adult friendship choices. Social forces 56 (2), 576–597.

Weinberg, B. A., 2007. Social interactions with endogenous associations. Tech. rep., National Bureau of Economic Research.

Weitzman, M. L., 1992. On diversity. The Quarterly Journal of Economics 107 (2), 363–405.

Appendix

This appendix is divided into three sections. In Section A, we provide robustness checks for our analysis of sorting by gender. In Section B, we present the results of a full analysis of sorting by age instead of by gender. In Section C, we present more detailed venue summary statistics stratified by subcategory.[47]

A Robustness Checks: Sorting by Gender

A.1 Measurement Error

Although user-generated datasets offers much promise, they are accompanied by several potential concerns regarding measurement error. In this section, we present evidence that our main results are qualitatively robust to many reasonable forms of measurement error.

A.1.1 Checkins Are Not Representative of Venue Visits

Our primary concern is that the proportion of females that we observe in a venue may be systematically different from the proportion of females that actually visit the venue. We argue that this likely does not confound our analysis, and, in any case, we show empirically that our results are qualitatively robust to the extent that it does. Indeed, in the presence of such measurement error, our results should actually be understood as conservative estimates of the amount of sorting within neighborhoods and the effects of this sorting on neighborhood and venue diversity.

To fix ideas, let $\tilde{f}_{jk}$ and $\tilde{m}_{jk}$ represent the actual numbers of females and males who visit venue j in neighborhood k. We can write the relationships between the observed and actual variables as

$$f_{jk} = \gamma^{f}_{jk} \cdot \tilde{f}_{jk} \tag{9}$$

$$m_{jk} = \gamma^{m}_{jk} \cdot \tilde{m}_{jk} \tag{10}$$

where the γ_{jk} parameters represent gender and venue specific check-in rates. All observed variables previously defined in terms of f_{jk} and m_{jk} have an actual, unobserved counterpart denoted with a tilde.

[47] This appendix is available online at http://bit.ly/1KzNf2X.

When mismeasurement is not gender specific, i.e., $\gamma_{jk}^{f} = \gamma_{jk}^{m}$, the female shares of check-ins at venues are unchanged, so all of our results are unaffected. This is a particularly nice feature, as it ensures our results are robust to any basic form of measurement error due to the fact that not all venue customers use the Foursquare app. Moreover, if mismeasurement is gender specific, but the mismeasurement in the female share of venues is only neighborhood specific (i.e., $s_{jk} = \gamma_k^s \cdot \tilde{s}_{jk}$), then our estimates of neighborhood Theil indices and their geographic decompositions are unchanged. This ensures that our results are robust to neighborhood specific sources of measurement error such as those correlated to unobserved neighborhood amenities.

In general, measurement error may be not only gender and neighborhood specific but also venue specific. We check the sensitivity of our main results to a general form of measurement error by conducting a Monte Carlo simulation. Without loss of generality, we define $\omega_{jk} = \frac{\gamma_{jk}^{m}}{\gamma_{jk}^{f}}$ to be the relative oversampling of males in venue j. For each iteration l, we randomly draw ω_{kj}^{l} for each venue from a uniform distribution $[\underline{\omega}, \overline{\omega}]$. We then calculate the "true" values of $\tilde{s}_{jk}^{l}$, $\tilde{T}_k^l$ for that iteration. Using these "true" values, we can simulate the main results of the paper, and the variation of the results across iterations allows us to construct confidence intervals. Although ω_{jk}^{l} is randomly drawn, it is positively correlated to $\tilde{s}_{jk}^{l}$ by construction.[48]

We conduct the Monte Carlo simulation under three separate parametrizations to capture qualitatively different types of measurement errors. In the first parametrization, we set $\underline{\omega} = 0.5, \overline{\omega} = 1.5$, which allows males to check in up to 50% less or more frequently than females, though they check in at the same rate on average. In the second parametrization, we set $\underline{\omega} = 2, \overline{\omega} = 4$. This increases the measurement error in two ways: it assumes that on average males check in three times more than females do, and it allows for greater dispersion of γ_{jk} across venues. In the third parametrization, we set $\underline{\omega} = 1, \overline{\omega} = 5$ which further worsens measurement error by allowing for even greater dispersion of ω_{jk} across venues.[49]

We report the Monte Carlo ($N = 500$) results for each of the main estimates of the paper in Figure 10 and in Table 6. Each panel in Figure 10 contains 24 bars, which represent the three different sets of parameters for each city in our sample. For each set of parameters, the bars represent

[48] We also performed alternative Monte Carlo simulations where we allowed ω_{jk}^{l} to be positively (or negatively) correlated to s_{jk} instead and obtained qualitatively similar results.

[49] We also performed analogous Monte Carlo simulations assuming females check in more rather than less frequently than males on average and found analogous results.

the average estimate of that result across all 500 iterations. We also show 95% confidence intervals for these estimates along with the previously presented value of that result under the assumption of no measurement error denoted with an "x".

In the first panel of Figure 10, it is clear that our estimate of the fraction of the city sorting that occurs within census blocks is robust to various amounts of measurement error; if anything we underestimate the amount of sorting that occurs locally.[50] Even though the actual estimates under the assumption of no measurement error may fall outside of the confidence interval, they are qualitatively the same. A large proportion of sorting happens within blocks under all reasonable assumptions on measurement error. In second and third panels, we show how our regression results are affected by different kinds of measurement errors. If anything, measurement error leads to attenuation bias, mainly in $\hat{\beta}^V$. This is consistent with the results of our panel and IV identification strategies and suggests that our conclusion that $\beta^V < 0$ and $\beta^N > 0$ may be conservative. Overall, these simulations suggest that our results are generally robust to measurement error. Even though erroneously assuming away measurement error might lead us to estimate parameters that would fall outside of the true confidence intervals in some cases, our qualitative conclusions should not be affected even by very extreme forms of measurement error.

In Table 6, we present 95% confidence intervals for our estimates of γ_w^V and γ_w^N from analogous Monte Carlo simulations of row (4) of Table 8. Consistent with our main findings, our estimates of γ_{low}^V are robustly significantly different from zero at the 90% level.

[50]Because the measurement error that we introduce in the Monte Carlo simulation is correlated to the female share of venues, the across-neighborhood component of city sorting tends to be magnified more than the within-neighborhood component (see equation (2)).

Figure 10: Robustness: Monte Carlo Results

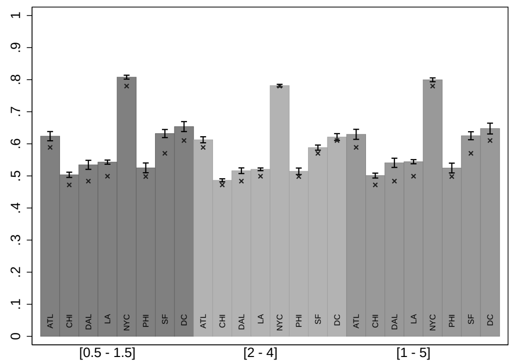

(a) Proportion of City Sorting due to Within Census Blocks Sorting

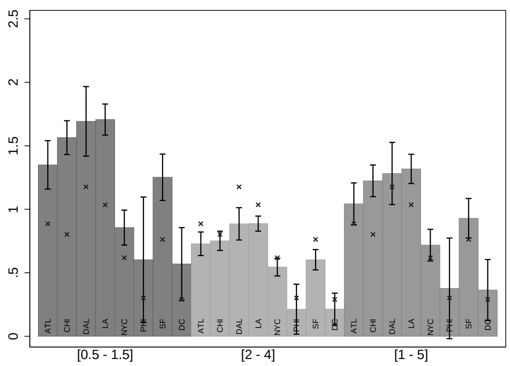

(b) $\hat{\beta}^N$

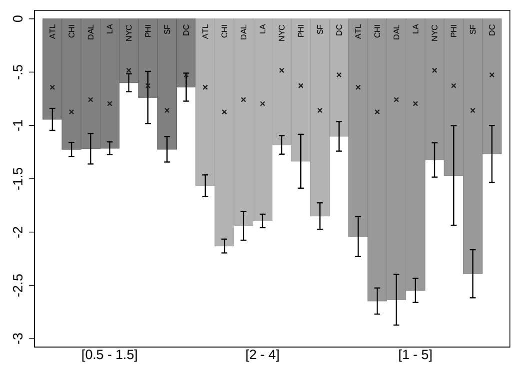

(c) $\hat{\beta}^V$

Notes: Each panel presents Monte Carlo results for three different set of parameters $[\underline{\omega}, \overline{\omega}]$, which represent the interval of the uniform distribution from which ω_{jk} is drawn: $[0.5, 1.5]$, $[2, 4]$ and $[1, 5]$. The bars represent the estimates of the Monte Carlo with 95% confidence intervals, and "x" represents the estimates under the assumption of no measurement error, which are reported in the paper.

Table 6: Monte Carlo Estimates of Effects of Gender Diversity on Labor Force Participation Gaps

	Low Wage Jobs		Medium Wage Jobs		High Wage Jobs	
$[\underline{\omega}, \overline{\omega}]$:	γ^V_{low}	γ^N_{low}	$\gamma^V_{\text{med.}}$	$\gamma^N_{\text{med.}}$	γ^V_{high}	γ^N_{high}
$[0.5, 1.5]$	-0.011**	0.000	-0.002	0.001	-0.010	0.006
$[2, 4]$	-0.011***	0.002	0.002	-0.001	-0.016**	-0.009
$[1, 5]$	-0.009*	0.001	0.001	-0.000	-0.012	0.006

Notes: Low wage jobs pay less than $1,250 monthly, medium wage pay jobs pay between $1,250 and $3,333 monthly, and high wage pay jobs pay more than $3,333 monthly. All specifications include block group - wage group fixed effects, cubic B-spline (with as many knots as possible) for number of venues in block, numbers of female and male block visitors, numbers of female and male block residents, numbers of younger ($\leq$ 35) and older ($>$ 35) block residents (38 covariates for each group). ***: 1% significance level, **: 5% significance level, *: 10% significance level computed from Monte Carlo simulations with 500 draws.

A.1.2 Selected Venue Coverage

There may be some venues that do not experience any check-in activity during the sample period, so it is useful to consider the implications of this form of measurement error on our results. Given the vast size of our data set, the number of unobserved venues is likely to be small in the well traveled urban areas that comprise our sample. In Figure 11, we present heat maps of the density of venues in our sample for each of our eight cities. Borders correspond to census tracts, and more darkly shaded tracts contain more venues. In all of the sample cities, we find a concentration of venues in the central business district, and some reduction in venue density in more residential surrounding areas. This is anecdotally consistent with the structure of these cities and indicates the density of venues in our sample is spatially consistent with the density of venues in the overall population of venues.

Although we do not have information on unsampled venues by definition, we conjecture that unsampled venues would tend to be more similar to "barely sampled" venues (i.e., those that experience only a small number of check-ins) than to the more robustly sampled venues that comprise the bulk of our data set. This suggests an empirical robustness check that we can perform to see if hypothetically observing unsampled venues would dramatically alter our results. A venue is included in our sample if it experiences at least 10 check-ins over the one year sample period. As a robustness check, we increase this threshold in increments of 5 check-ins and replicate our entire analysis using these diminishing subsamples. If our results do not change much near the 10 check-in threshold, then it is reasonable to assume that the exclusion of unsampled venues would also have a small effect on our results.[51]

In the first two panels of Figure 12, we present our three main results – the fraction of sorting in each city due to sorting within census blocks and the estimates of β^V and β^N from our baseline regressions – replicated on subsamples with inclusion thresholds varying from 10 check-ins to 365 check-ins during our sample period. In the first panel, the fractions of sorting in each city that are due to sorting within blocks are quite flat near the 10 check-in threshold, which indicates that measurement error due to unsampled venues is not likely to affect our evidence of the intensity of homophily and highly local sorting. In the second panel, the estimates of β^V and β^N are also flat

[51]It is inadvisable to include venues that experience fewer than 10 check-ins in our sample because then we would be unable to obtain sufficiently fine estimates of the gender compositions of those venues.

near the 10 check-in threshold. To the extent that they trend away from zero as we include venues with fewer check-ins suggests that this form of measurement error attenuates our results. Hence, if anything our reported estimates are conservative.[52]

[52]The trends away from zero of our regression coefficient estimates as we include venues with fewer check-ins are consistent with our finding that the effects of venue variety on venue and neighborhood diversity are largest in neighborhoods with low levels of venue variety (b-spline specification). This serves as additional evidence that, if anything, our regression estimates are conservative.

Figure 11: Venue Coverage Maps of Sample Cities

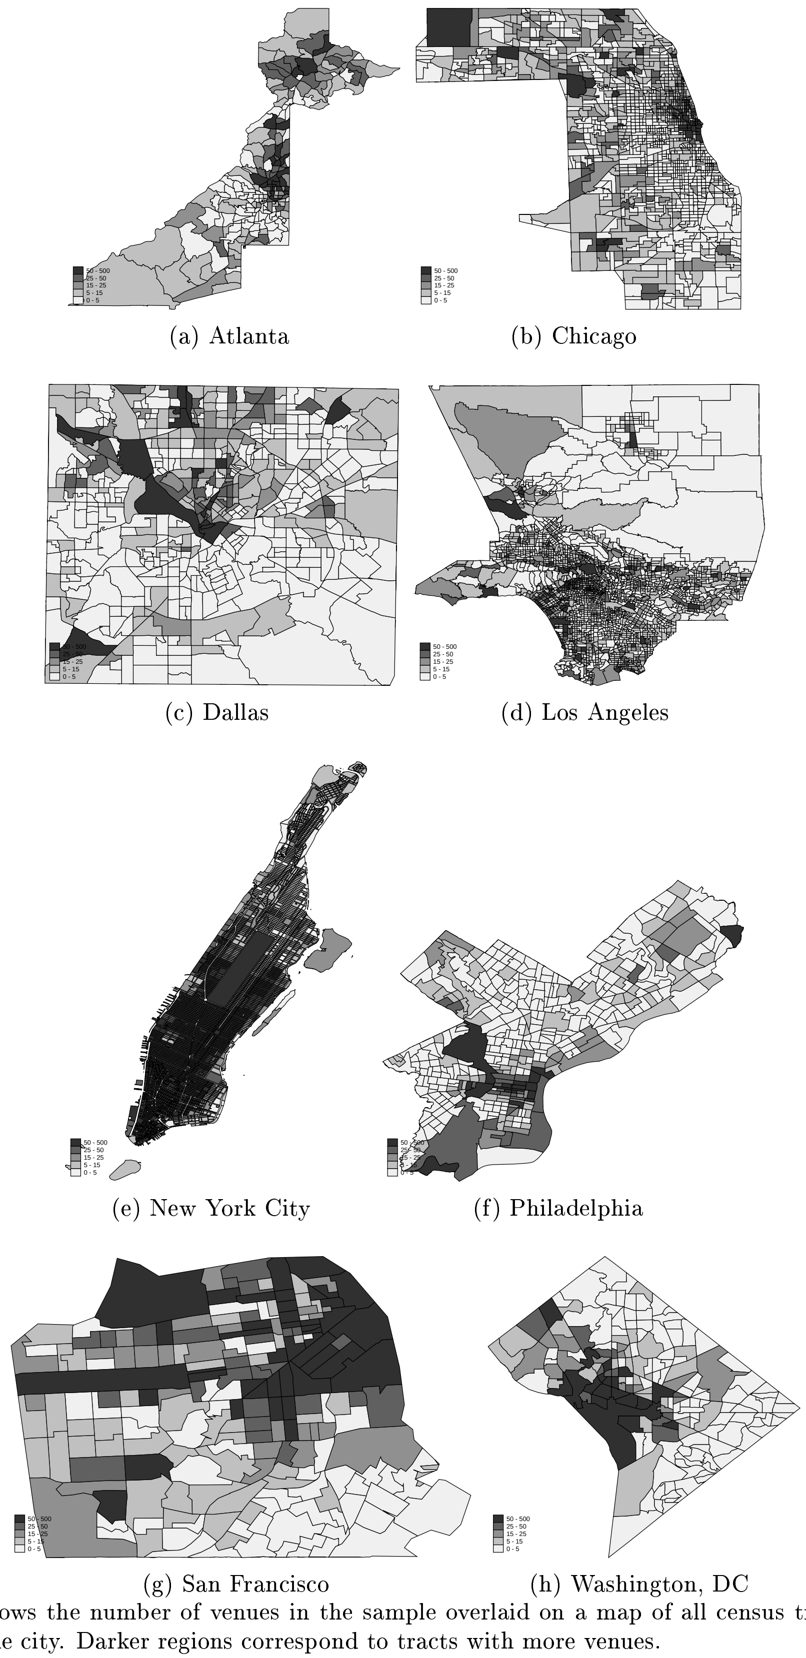

(a) Atlanta (b) Chicago

(c) Dallas (d) Los Angeles

(e) New York City (f) Philadelphia

(g) San Francisco (h) Washington, DC

Notes: Each map shows the number of venues in the sample overlaid on a map of all census tracts in the primary county of each sample city. Darker regions correspond to tracts with more venues.

Figure 12: Robustness: Selected Venue Coverage

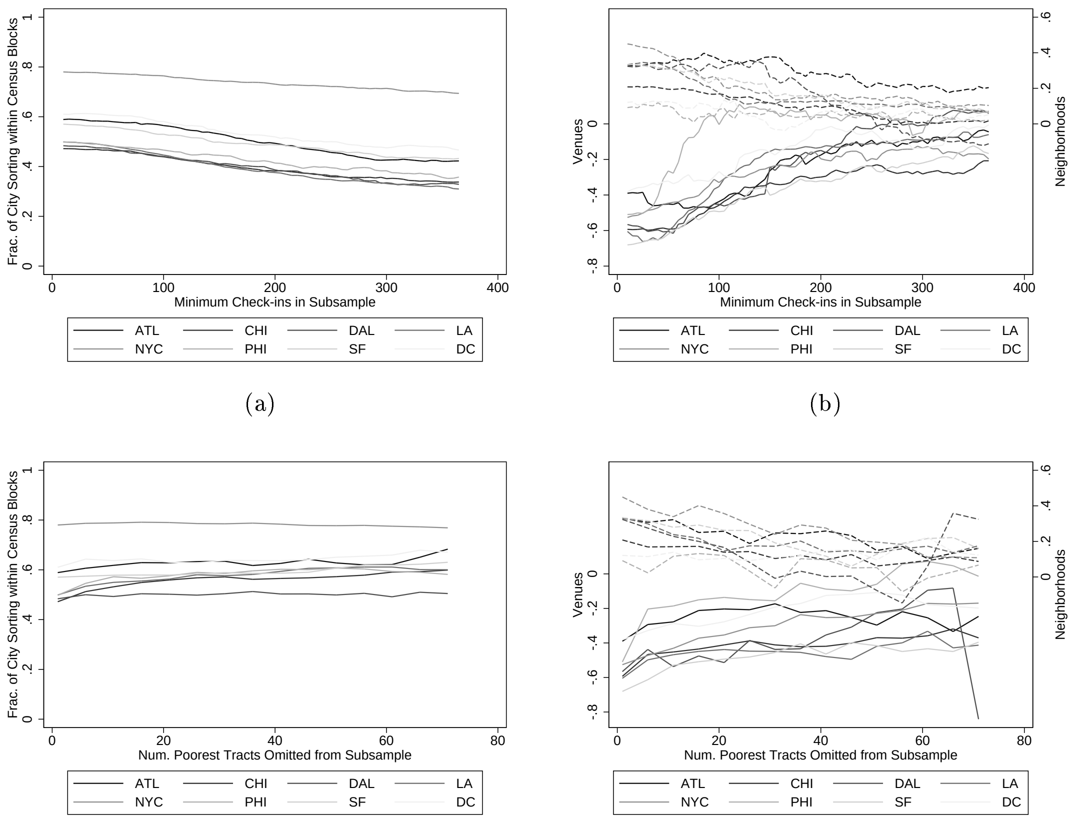

(a) (b)

(c) (d)

Notes: Solid lines refer to the y-axis on the left, and dashed lines refer to the y-axis on the right. The measures in the first two panels are recalculated using subsamples that include only venues that experience at least a given number of check-ins during our sample period. The measures in the last two panels are recalculated using subsamples that include only venues in tracts with sufficiently high median income ranks according to the 2013 American Communities Survey.

Because checking in on Foursquare requires the use of a "smart" mobile device, Foursquare users likely tend to be wealthier, and hence they might disproportionately frequent more expensive venues. We can assess the extent to which the potential selection of venues in our sample due to this effect biases our results in a similar exercise to the one above. In each city, we rank all tracts by their median household income according to the 2013 American Community Survey. We incrementally eliminate all venues in the 5 poorest tracts, 10 poorest tracts, 15 poorest tracts, etc. and replicate our entire analysis using these diminishing subsamples. If our results do not change much as we are changing the poorest tracts of the sample, then it is reasonable to assume that any selection of the venues in our sample due to users being wealthier would also have a small effect on our results. In the third and fourth panels of Figure 12, we present the same three results replicated on subsamples with the 5 to 75 poorest tracts in each city omitted. The results are highly similar to their counterparts in the first two panels, which suggests that this form of measurement error does not qualitatively affect our main results.

In sum, these results allows us to conclude that the coverage of the venues in our sample is quite comprehensive, and to the extent that there may be selection in the sample then our reported results will be conservative.

A.1.3 Sampling Error: A Falsification Test

Consider the extreme situation in which all venues in neighborhood k have the same true female share $\tilde{s}_{jk}$ but we observe variation in s_{jk} across venues purely because of sampling error. Under this falsification exercise, how would our main results differ? To answer this, we simulate a counterfactual in which the individuals in a city sort across tracts, block groups and blocks according to the data, but they do not sort within blocks. This provides an intuitive falsification test of our interpretation of our main findings: if the block level Theil indices constructed under this counterfactual are similar to their analogs as constructed with our data, then our results should not be interpreted as evidence of local sorting.

We implement this test by randomly assigning individuals to venues in a particular block in proportion to the overall gender distribution that we observe in that block. If we observe venue i in block b with $f_{ib} + m_{ib}$ check-ins in our data, we recreate the gender composition of venue i by taking $f_{ib} + m_{ib}$ independent draws from a Bernoulli(p_b) distribution with replacement, where

$$p_b = \frac{\sum_{i \in b} f_{ib}}{\sum_{i \in b} f_{ib} + m_{ib}} \tag{11}$$

is the overall proportion of female check-ins in block b (i.e., across all venues). For each 1 that is drawn, we add a female to venue i, and for each 0 that is drawn, we add a male to venue i. The variation in the gender composition of venues within blocks in this simulated sample is fully attributable to measurement error.

Table 7: Placebo Tests: No Active Sorting Within Census Blocks

Placebo for:	Proportion of city-wide sorting due to sorting within:			$\hat{\beta}^V$	$\hat{\beta}^N$
	Tracts	Block Groups	Blocks		
Atlanta	0.73	0.59	0.03	0.00	0.39
Chicago	0.68	0.53	0.02	0.00	0.26
Dallas	0.62	0.46	0.03	0.00	0.33
Los Angeles	0.68	0.51	0.04	0.00	0.32
New York City	0.70	0.51	0.05	0.00	0.71
Philadelphia	0.72	0.60	0.03	0.00	0.26
San Francisco	0.65	0.54	0.04	0.00	0.44
Washington, DC	0.70	0.61	0.03	0.00	0.40

Notes: All results are calculated under the placebo assumption that individuals do not actively sort within census blocks. Bootstrapped standard errors for all entries in all cities are less than 0.005 and are omitted for clarity.

For each simulated sample of venues, we can re-estimate our results. We repeat this exercise 500 times and report the mean and standard deviation of these counterfactual results across all repetitions. In Table 7, we present the fraction of sorting within each city that is due to sorting within neighborhood types, and baseline estimates of β^V and β^Nunder this counterfactual assumption of no sorting within blocks. The results are as expected. The proportion of venue sorting within tracts and within block groups decreases slightly as expected, and the proportion of sorting within blocks is reduced to nearly zero, as such sorting can only be due to measurement error. In addition, our estimates of β^V decrease to zero as expected (with no sorting within blocks, venue diversity should be unaffected by venue variety) while our estimates of β^N remain positive and of the same order of magnitude as before, as sorting across neighborhoods is unchanged under the counterfactual.

From this exercise, we find that our main results differ completely from their counterfactual counterparts, which constitutes further evidence that our main results are not artifacts of measurement error.

A.2 Dynamic Misaggregation

Per the discussion in the data section, we aggregated check-ins in our sample annually to reduce any potential measurement error. However, if there are strong dynamic components to gender sorting, this aggregation could potentially obscure interesting longitudinal variation in venue sorting. For example, this could happen if venues varied in substitutability by season (e.g., people may not enjoy parks as much in the winter, especially in cold weather cities), or by day of the week (e.g., people may prefer downtown venues on weekdays). We replicate our analysis disaggregated by day of week and by month, and present the main results in Figure 13.

Figure 13: Robustness: Dynamic Aggregation

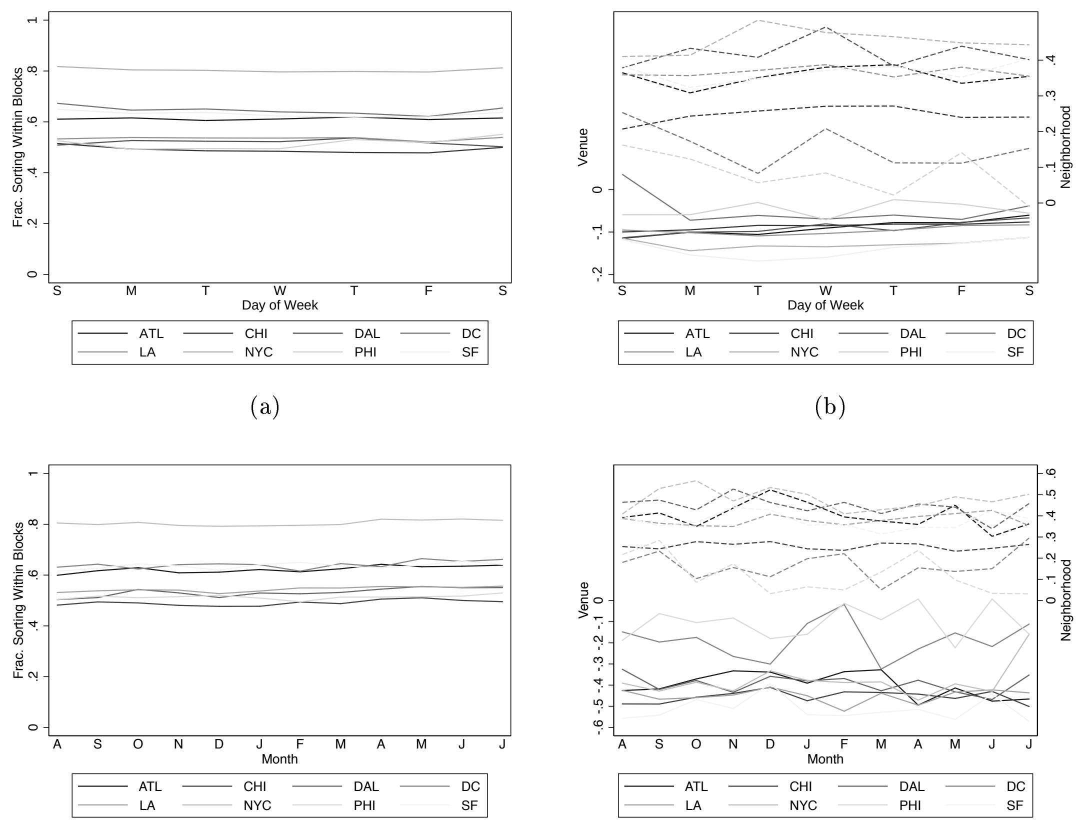

(a) (b) (c) (d)

Notes: Solid lines refer to the y-axis on the left, and dashed lines refer to the y-axis on the right. All measures are calculated by replicating the analysis by day of week (first two panels) or by month (last two panels).

Looking at the first and third panels, it is immediate that there is nearly zero dynamic variation in the fraction of sorting in each city due to sorting within blocks. There is markedly more dynamic variation in our estimates of β^V and β^N for each city, as depicted in the second and fourth panels (left and right axes, respectively). However, this variation does not follow any systematic trend, and we infer that $\beta^V < 0$ and $\beta^N > 0$ for all days of the week and months of the year. These exercises suggest that our main results are unaffected by our choice of annual aggregation.

A.3 Robustness Checks for $\hat{\beta}^V$ and $\hat{\beta}^N$ by City

The following plots show estimates of β^V and β^N by city for the baseline specification (block group FEs), the specification replacing block group FEs for tract FEs, the b-spline specification and the panel data specification replacing block FEs and city-month FEs for block group FEs.[53]

The dark shaded bars represent $\hat{\beta}^V$, and the light shared bars represent $\hat{\beta}^N$. Venue variety is defined as the number of unique venue categories in the first column and the number of unique venue subcategories in the second column. The first bars correspond to baseline estimates. The second bars replace the block group fixed effects in the baseline estimates with tract fixed effects. The third set of bars correspond to estimates of the parameters specified as a linear b-spline with a knot at 3 three categories or subcategories. The fourth bars correspond to estimates where the dataset is disaggregated to a monthly panel and the block group fixed effects are replaced with block fixed effects. As can be seen, the results are similar to the ones reported in the paper.

[53]We were unable to obtain IV estimates disaggregated by city due to their lack of precision.

Figure 14: $\hat{\beta}^V$ and $\hat{\beta}^N$: Alternative Identification Strategies For Gender Sorting by City (1 of 2)

V_b = Category $\quad\quad$ V_b = Subcategory

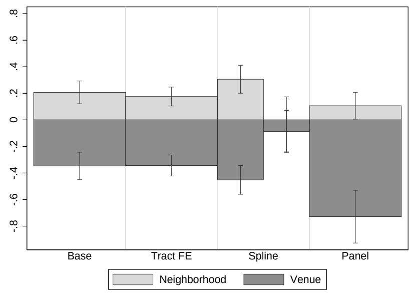

(a) Atlanta

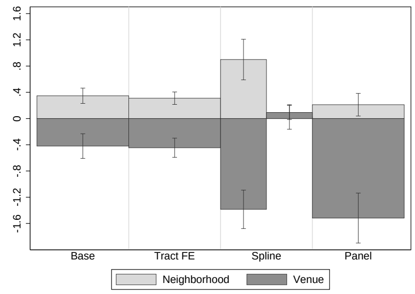

(b) Atlanta

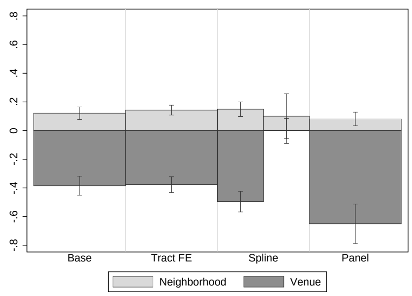

(c) Chicago

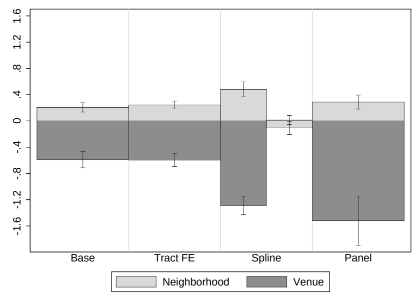

(d) Chicago

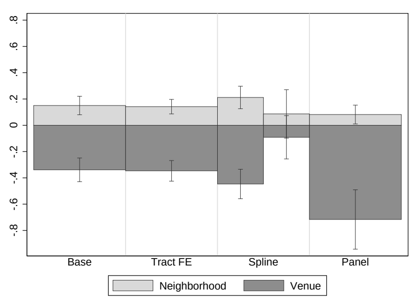

(e) Dallas

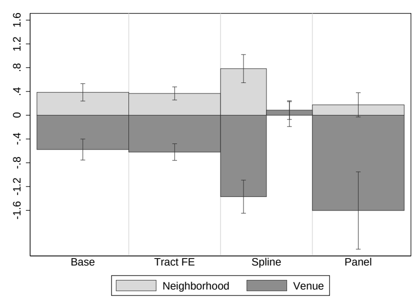

(f) Dallas

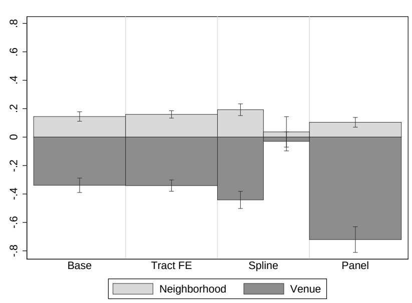

(g) Los Angeles

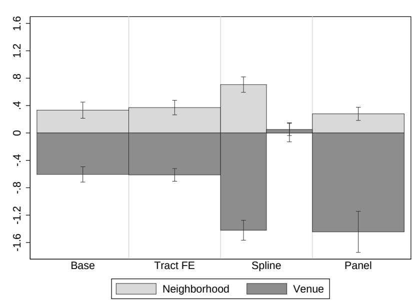

(h) Los Angeles

Note: See next page.

Figure 15: $\hat{\beta}^V$ and $\hat{\beta}^N$: Alternative Identification Strategies For Gender Sorting by City (2 of 2)

V_b = Category $\qquad\qquad$ V_b = Subcategory

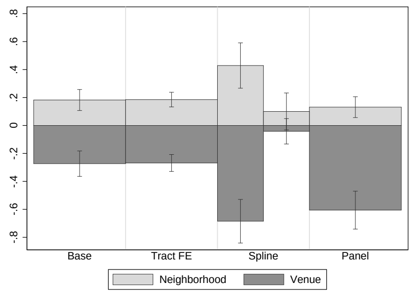

(a) New York City

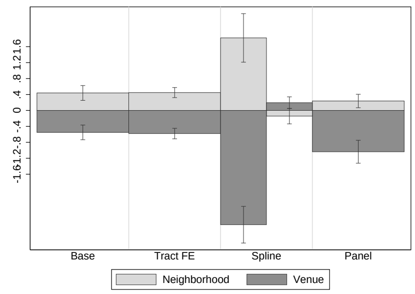

(b) New York City

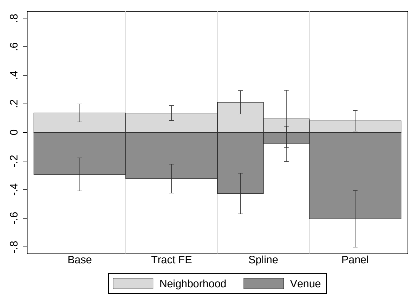

(c) Philadelphia

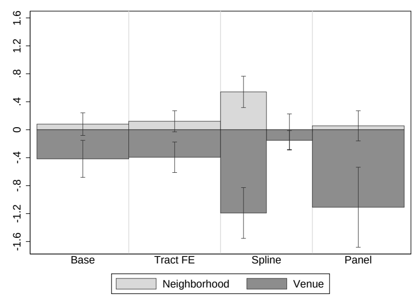

(d) Philadelphia

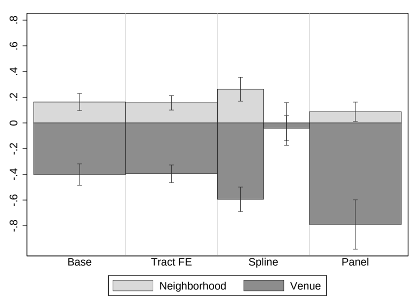

(e) San Francisco

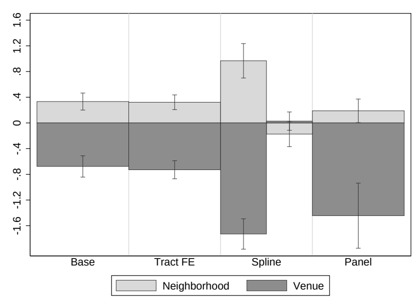

(f) San Francisco

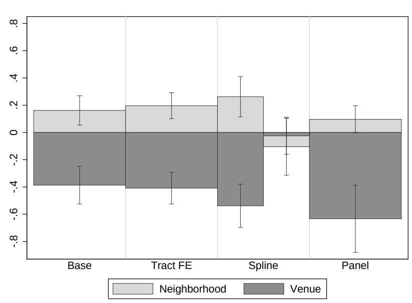

(g) Washington, DC

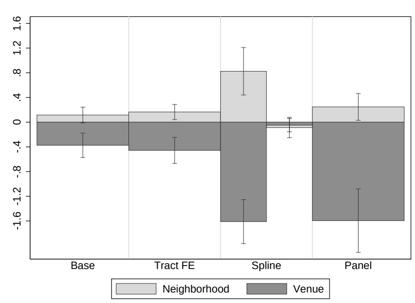

(h) Washington, DC

A.4 Gender Homophily and the Labor Force Participation Gap: Further Robustness Checks

Table 8: Effects of Diversity on the Labor Force Participation Gender Gap: $t \geq 2010$

	Specification	Low Wage Jobs γ^V_{low}	γ^N_{low}	Medium Wage Jobs $\gamma^V_{\text{med.}}$	$\gamma^N_{\text{med.}}$	High Wage Jobs γ^V_{high}	γ^N_{high}	N	Adj. R^2
(1)	Controls from Foursquare Data[1]	-0.014**	-0.004	-0.004	0.006	0.001	0.008	152,335	0.195
		(0.006)	(0.006)	(0.007)	(0.007)	(0.010)	(0.011)		
(2)	(1) + residential gender controls[2]	-0.014**	-0.004	-0.004	0.005	0.002	0.005	152,335	0.218
		(0.005)	(0.006)	(0.006)	(0.007)	(0.010)	(0.010)		
(3)	(2) + flexibly specified controls [3]	-0.014**	-0.007	0.001	-0.000	-0.002	0.005	152,335	0.328
		(0.005)	(0.005)	(0.005)	(0.006)	(0.008)	(0.008)		
(4)	(3) + flexible residential age controls[4]	-0.014**	-0.007	0.001	-0.000	-0.001	0.005	152,335	0.328
		(0.005)	(0.005)	(0.005)	(0.006)	(0.007)	(0.008)		
(5)	(4) + workers' demographics characteristics[5]	-0.013**	-0.006	0.000	-0.000	-0.002	0.004	152,335	0.334
		(0.005)	(0.005)	(0.005)	(0.006)	(0.007)	(0.008)		
(6)	(4) + tract-wage group FEs instead of block group-wage group FEs	-0.010**	-0.006	-0.002	0.006	0.002	0.000	152,335	0.303
		(0.004)	(0.004)	(0.004)	(0.005)	(0.006)	(0.006)		
(7)	(4) + D_b^V and D_b^N for age[6]	-0.011**	-0.004	0.003	-0.002	-0.001	0.003	152,335	0.328
		(0.005)	(0.005)	(0.006)	(0.006)	(0.008)	(0.009)		
(8)	(4), drop D_b^N	-0.014**	–	0.001	–	-0.000	–	152,335	0.328
		(0.005)		(0.005)		(0.007)			
(9)	(4), drop blocks where residents also work[7]	-0.014**	-0.007	-0.001	-0.000	-0.001	0.000	127,024	0.309
		(0.006)	(0.006)	(0.006)	(0.006)	(0.008)	(0.009)		
(10)	(4) + block group-year-wage group FEs instead of block group-wage group FEs	-0.013**	-0.007	0.003	0.000	-0.002	0.005	152,335	0.267
		(0.006)	(0.007)	(0.007)	(0.007)	(0.010)	(0.011)		

Notes: This Table is analogous to Table 4 in the paper, however it is estimated on a subsample of LODES data from 2010-2013. Low wage jobs pay less than \$1,250 monthly, medium wage pay jobs pay between \$1,250 and \$3,333 monthly, and high wage pay jobs pay more than \$3,333 monthly. All specifications include block group - wage group fixed effects, with the exceptions of (6) and (10). Robust standard errors clustered at the block level are presented in parentheses.

[1] : number of venues in block, and numbers of female and male block visitors (3 covariates for each group); [2] : add to (1) the numbers of female and male block residents (2 additional covariates for each group);[3] : Cubic B-spline (with as many knots as possible) of all controls in (2) (26 covariates for each group). [4] : add to (3) cubic B-spline of numbers of younger (≤ 35) and older (> 35) block residents (12 additional covariates for each group). [5] : Add to (4) the numbers of block workers who are White, Black, other (non-White and non-Black), Hispanic, non-Hispanic, college graduates, college non-graduates (7 additional covariates for each group). [6] : Add to (4) D_b^{Vy} and D_b^{Ny}, which are the analogous measures of D_b^V and D_b^N based on the proportion of younger visitors (≤ 35), rather than based on the proportion of female visitors (2 additional covariate for each group). [7] : Drop observations from blocks with at least one resident who works in the same block. **: 5% significance level, *: 10% significance level

B Results of Analysis by Age

We replicate all tables and figures for sorting by age, including those reported in the previous section.

Figure 16: Proportion of Youth in Venues by Subcategory

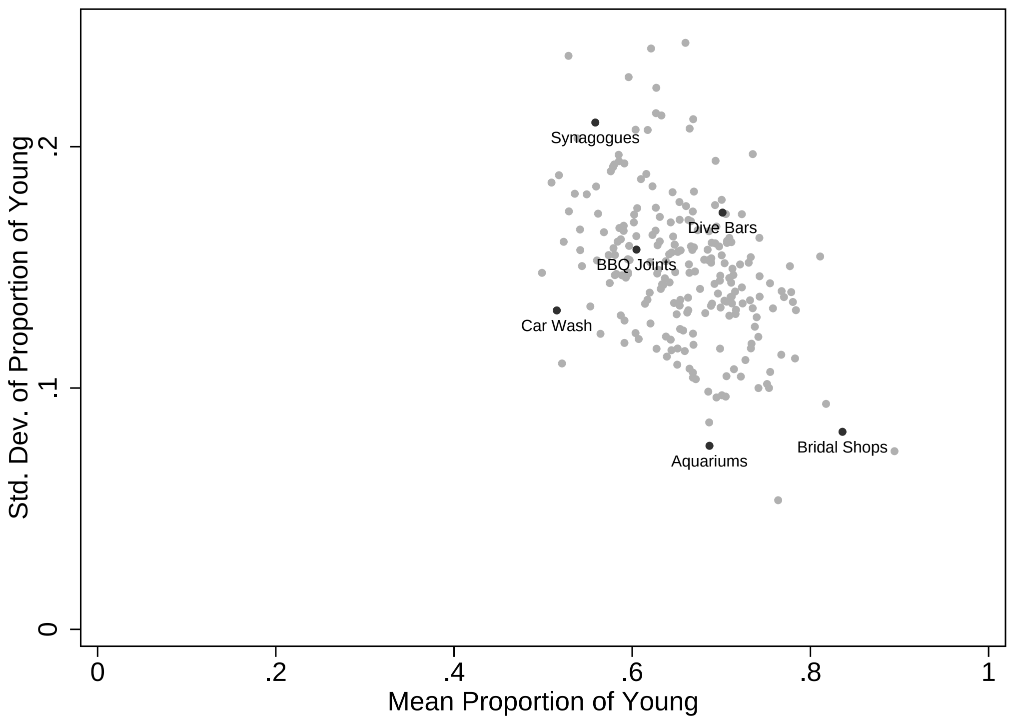

Note: This scatter plot pools venues from all cities in the sample. Each dot represents all of the venues within a subcategory.

Figure 17: Check-ins and Age Composition Over Time

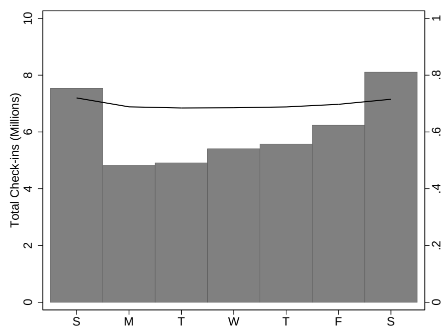

(a) Day-of-Week Variation

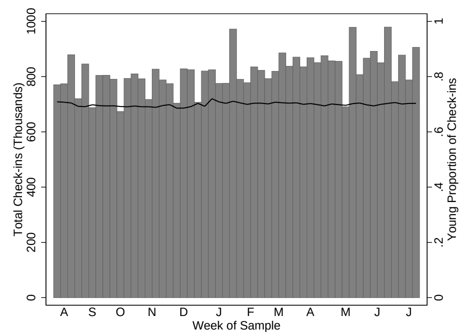

(b) Week-of-Year Variation

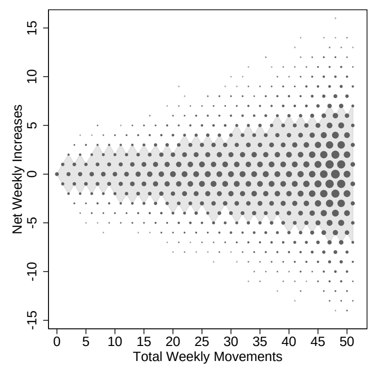

(c) Within Venue Variation

Notes: (a), (b): Bars represent total check-ins, lines represent age composition of aggregate check-ins. The 53rd week of the sample is omitted because it only contains a single day. (c): In this scatter plot of venues in our data, larger dots correspond to a greater numbers of venues. A venue experiences a weekly increase (decrease) in gender composition if the proportion of female check-ins rises (falls) by at least one percentage point.

Figure 18: Densities of Age Theil Indices for Various Neighborhood Definitions

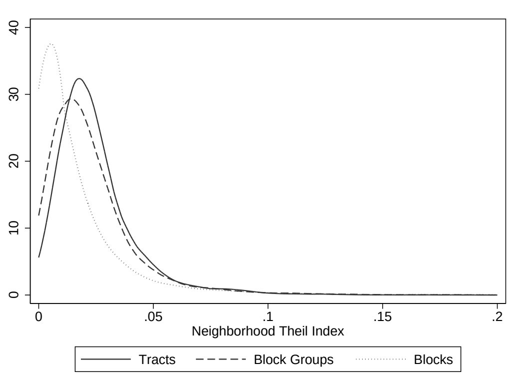

Notes: All densities are estimated using a bandwidth of 0.005 and an Epanechnikov kernel. For clarity, we present the density only for values of the domain less than 0.2; fewer than 1% of neighborhoods of any type have a Theil Index in excess of 0.2. Theil Indices are pooled across neighborhoods in all cities.

Figure 19: Sorting of Residents vs. Sorting of Visitors by Age

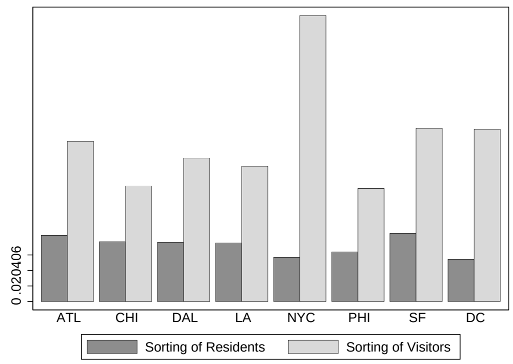

Note: “Sorting of Residents” is calculated as the Theil index of the gender composition of block residents from the 2010 Census. For comparability, “Sorting of Visitors” is calculated as the Theil index of the gender composition of check-ins in blocks. Bootstrapped standard errors for all estimates are below 0.005 and are omitted for clarity.

Figure 20: Excess Predicted Homophily in Venue Data

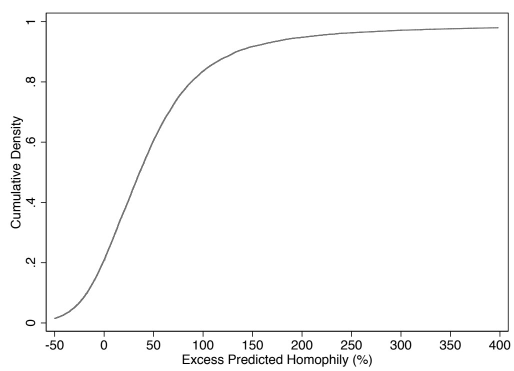

Note: In this figure, we present the empirical cumulative distribution of how much more likely we would predict that a youth would encounter another youth in a census block using venue level data than if we used residential data. Note that in some neighborhoods, we are more likely to predict homophily in residential data than in venue level data (negative excess predicted homophily). However, in most neighborhoods the opposite is true.

Table 9: Proportion of Within-Neighborhood Sorting By Age Due to Sorting Across Subcategories:

	City	Tracts	B. Groups	Blocks
Atlanta	0.14	0.75	0.82	0.91
Chicago	0.12	0.81	0.86	0.93
Dallas	0.14	0.79	0.83	0.92
Los Angeles	0.09	0.80	0.85	0.91
New York City	0.15	0.67	0.77	0.88
Philadelphia	0.16	0.79	0.83	0.93
San Francisco	0.16	0.71	0.77	0.91
Washington, DC	0.20	0.71	0.77	0.90

Note: Subcategories (225) are defined in the next section. Bootstrapped standard errors for all entries are less than 0.005 and are omitted for clarity.

Figure 21: $\hat{\beta}^V$ and $\hat{\beta}^N$: Alternative Identification Strategies for Age Sorting

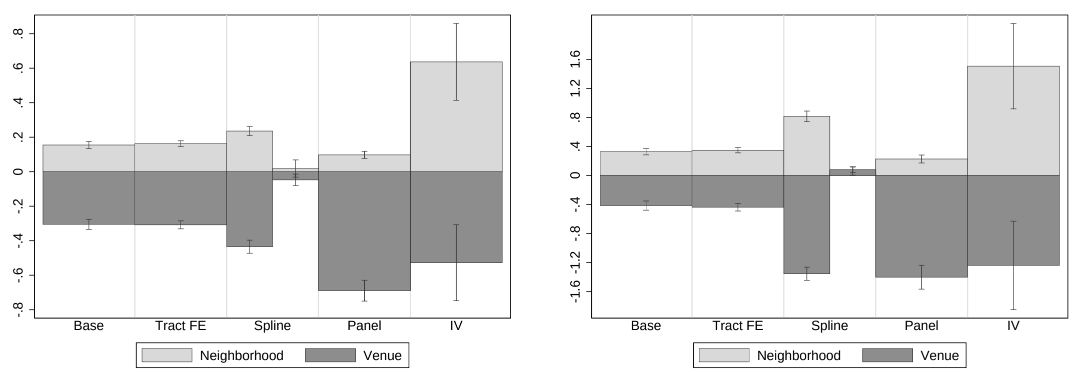

(a) V_b = Number of Categories (b) V_b= Number of Subcategories

Notes: The dark shaded bars represent $\hat{\beta}^V$, and the light shared bars represent $\hat{\beta}^N$. The first bars correspond to baseline estimates (block group FEs). The second bars replace the block group fixed effects in the baseline estimates with tract fixed effects. The third bars correspond to estimates from where the dataset is disaggregated to a monthly panel and the block group fixed effects are replaced with block fixed effects. The fourth bars correspond to 2SLS estimates of the baseline regressions with zoning instruments.

Figure 22: Robustness (Age): Monte Carlo Results

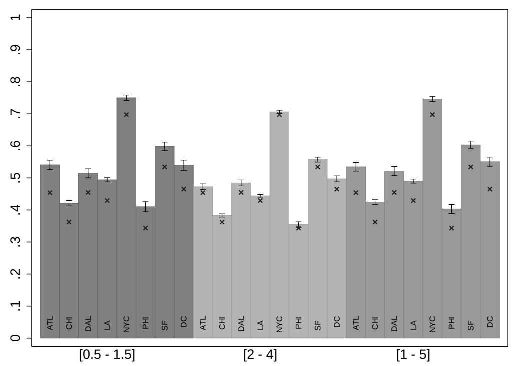

(a) Proportion of City Sorting due to Within Census Blocks Sorting

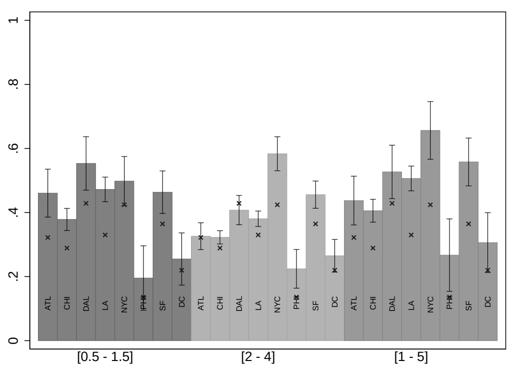

(b) $\hat{\beta}^N$

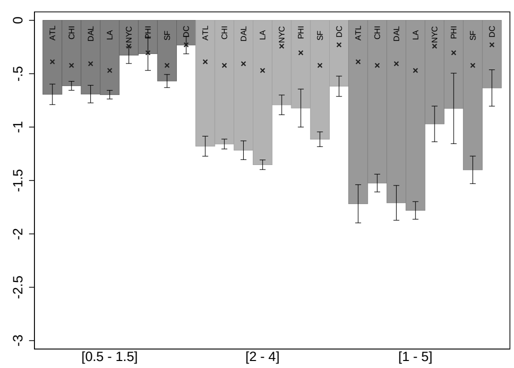

(c) $\hat{\beta}^V$

Notes: Each panel presents Monte Carlo results for three different set of parameters $[\underline{\omega}, \overline{\omega}]$, which represent the interval of the uniform distribution from which ω_{jk} is drawn: $[0.5, 1.5]$, $[2, 4]$ and $[1, 5]$. The bars represent the estimates of the Monte Carlo with 95% confidence intervals, and "x" represents the estimates under the assumption of no measurement error, which are reported in the paper.

Figure 23: Robustness (Age): Selected Venue Coverage

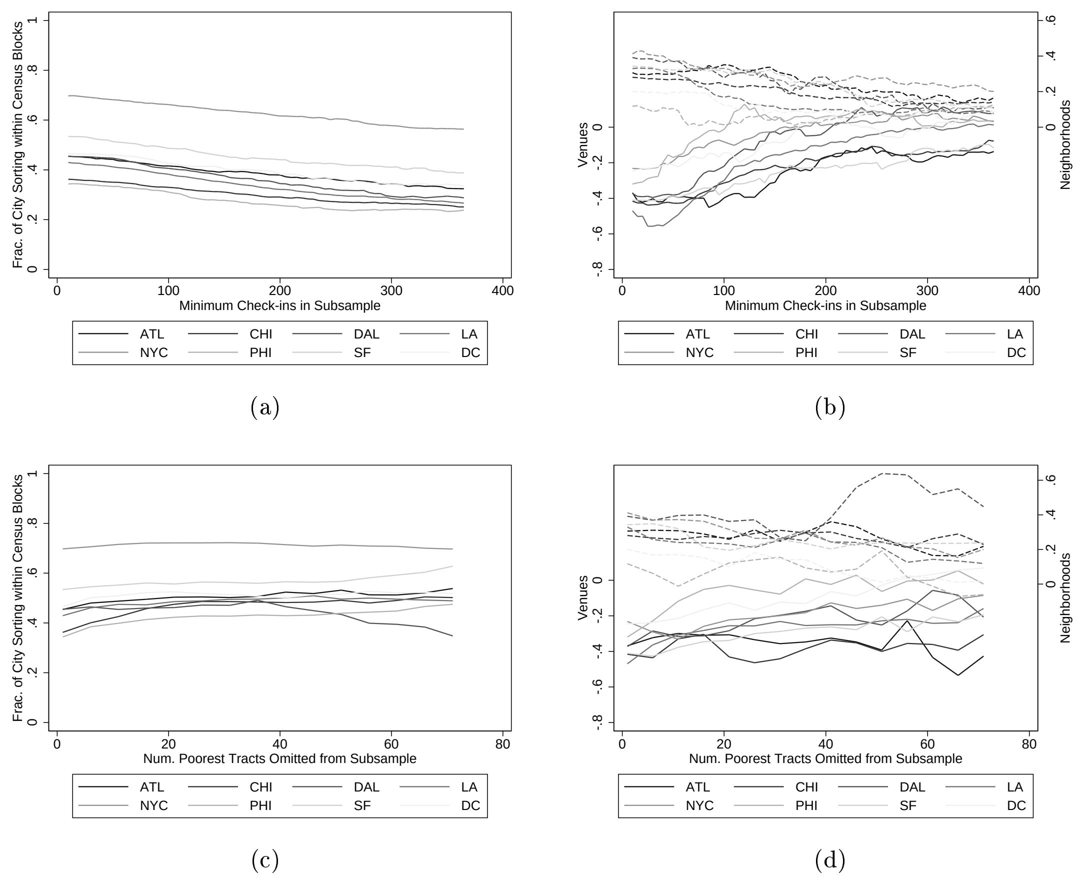

Notes: Solid lines refer to the y-axis on the left, and dashed lines refer to the y-axis on the right. The measures in the first two panels are recalculated using subsamples that include only venues that experience at least a given number of check-ins during our sample period. The measures in the last two panels are recalculated using subsamples that include only venues in tracts with sufficiently high median income ranks according to the 2013 American Communities Survey.

Table 10: Placebo Tests: No Active Age Sorting Within Census Blocks

Placebo for:	Proportion of city-wide sorting due to sorting within: Tracts	Block Groups	Blocks	$\hat{\beta}^V$	$\hat{\beta}^N$
Atlanta	0.54	0.41	0.02	0.00	0.38
Chicago	0.60	0.44	0.02	0.00	0.25
Dallas	0.59	0.44	0.03	0.00	0.40
Los Angeles	0.59	0.44	0.03	0.00	0.35
New York City	0.60	0.42	0.04	0.00	0.48
Philadelphia	0.53	0.42	0.02	0.00	0.15
San Francisco	0.67	0.54	0.03	0.00	0.37
Washington, DC	0.65	0.53	0.02	0.00	0.19

Notes: All results are calculated under the placebo assumption that individuals do not actively sort within census blocks. Bootstrapped standard errors for all entries in all cities are less than 0.005 and are omitted for clarity.

Figure 24: Robustness (Age): Dynamic Aggregation

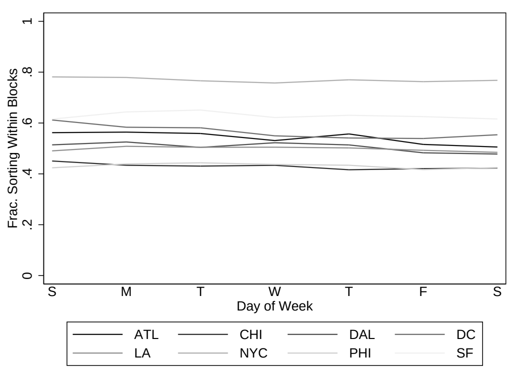

(a)

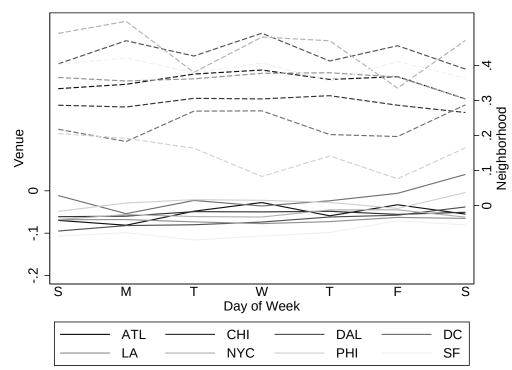

(b)

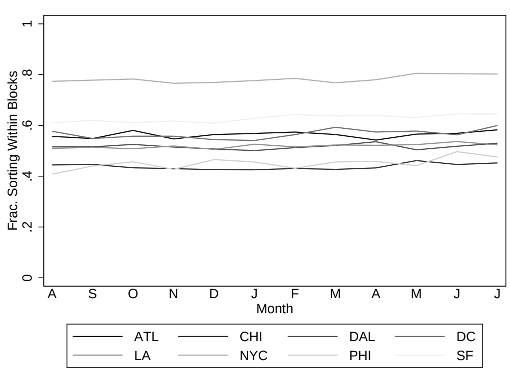

(c)

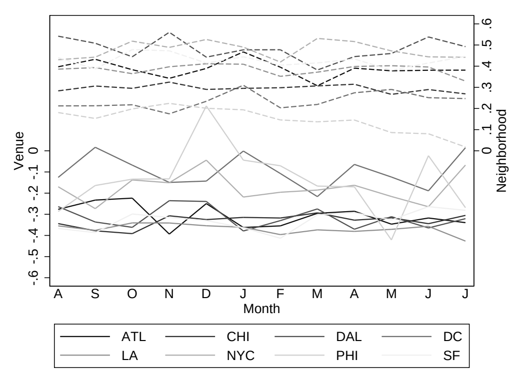

(d)

Notes: Solid lines refer to the y-axis on the left, and dashed lines refer to the y-axis on the right. All measures are calculated by replicating the analysis by day of week (first two panels) or by month (last two panels).

Figure 25: $\hat{\beta}^V$ and $\hat{\beta}^N$: Alternative Identification Strategies For Age Sorting by City (1 of 2)

V_b = Category $\qquad$ V_b = Subcategory

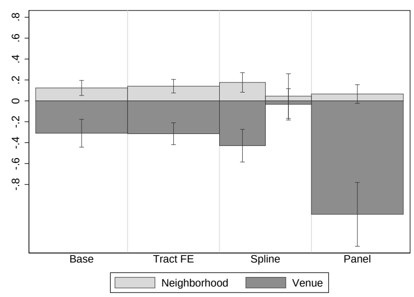

(a) Atlanta

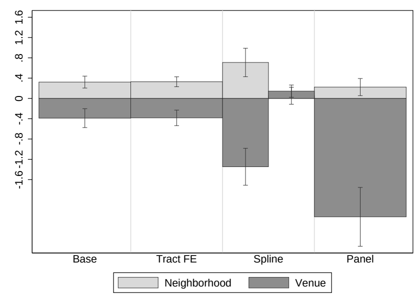

(b) Atlanta

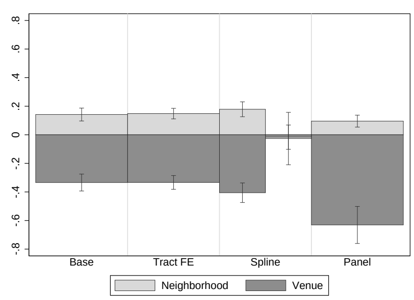

(c) Chicago

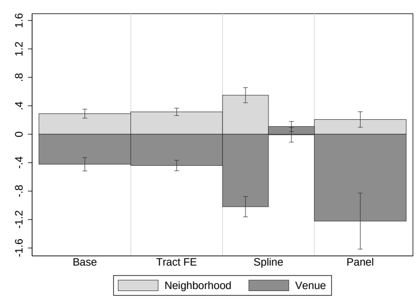

(d) Chicago

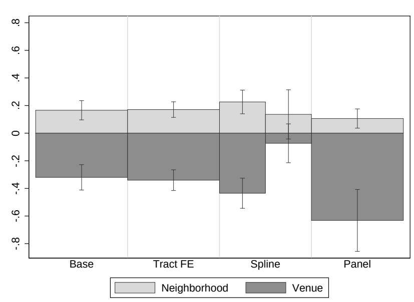

(e) Dallas

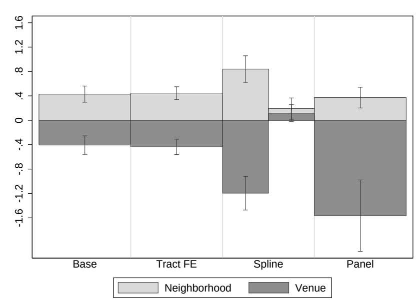

(f) Dallas

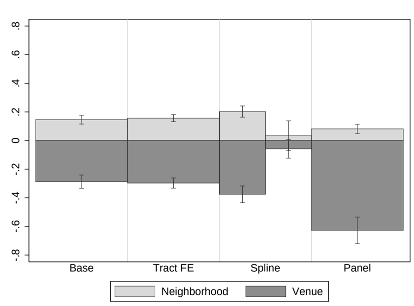

(g) Los Angeles

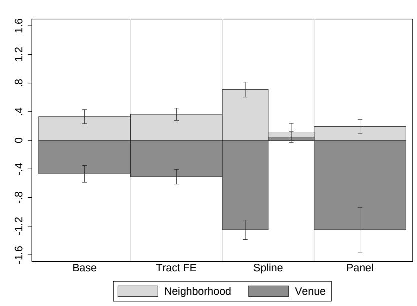

(h) Los Angeles

Note: See next page.

Figure 26: $\hat{\beta}^V$ and $\hat{\beta}^N$: Alternative Identification Strategies For Age Sorting by City (2 of 2)

V_b = Category $\qquad$ V_b = Subcategory

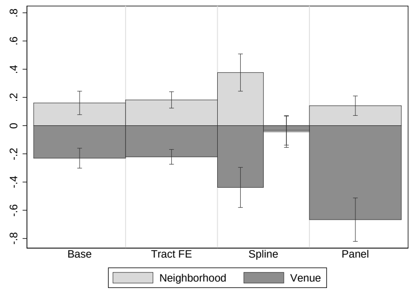

(a) New York City

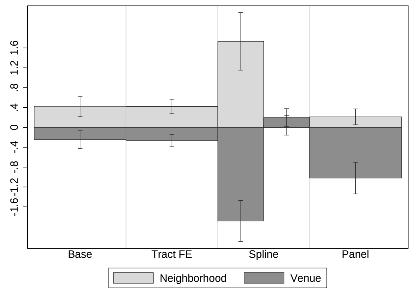

(b) New York City

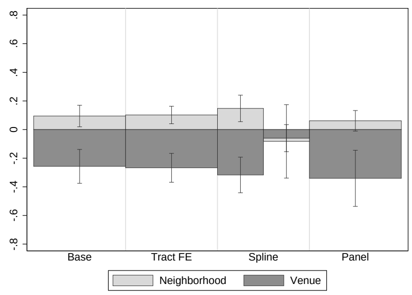

(c) Philadelphia

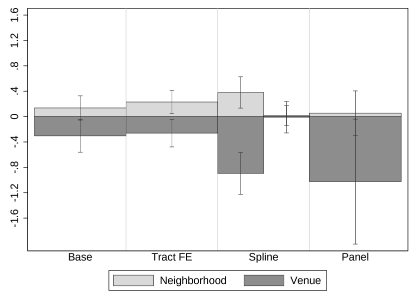

(d) Philadelphia

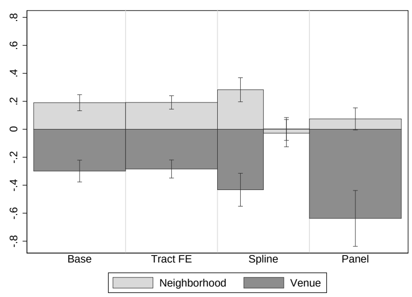

(e) San Francisco

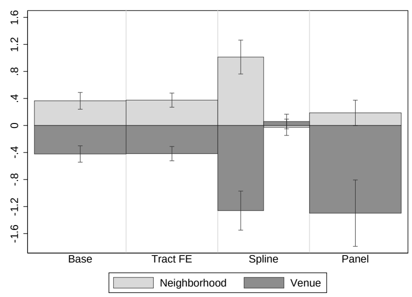

(f) San Francisco

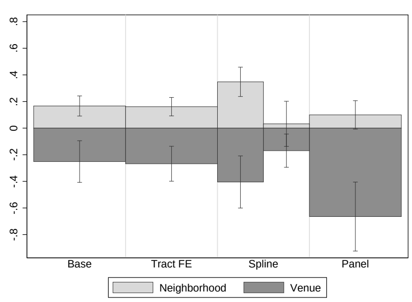

(g) Washington, DC

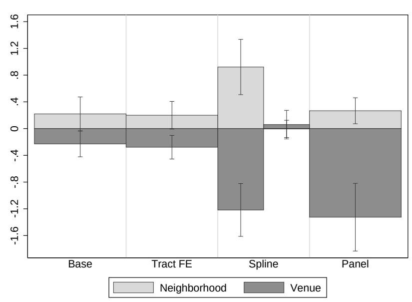

(h) Washington, DC

Notes: The dark shaded bars represent $\hat{\beta}^V$, and the light shared bars represent $\hat{\beta}^N$. Venue variety is defined as the number of unique venue categories in the first column and the number of unique venue subcategories in the second column. The first bars correspond to baseline estimates. The second bars replace the block group fixed effects in the baseline estimates with tract fixed effects. The third set of bars correspond to estimates of the parameters specified as a linear b-spline with a knot at 3 three categories or subcategories. The fourth bars correspond to estimates where the dataset is disaggregated to a monthly panel and the block group fixed effects are replaced with block fixed effects.

C Summary Statistics by Subcategory

The 9 categories of venues are further classified into 225 narrow subcategories. Foursquare users very actively check into even surprising types of venues such as *Banks*, *Cemeteries*, *Pharmacies, Synagogues*, and *Dog Runs*. In Figure 27, we present a scatter plot of the mean and standard deviation of the gender composition of venues for each subcategory throughout our entire sample. In general, the pattern of gender compositions of venues across subcategories looks intuitive and reasonable. For example, *Men's Stores*, not surprisingly, cater to mostly men, and this is fairly consistent across stores; conversely, *Nail Salons* cater to mostly women across all stores. *Hair Salons/Barbershops* cater to a mixed customer base in the aggregate; however the high standard deviation of the gender composition of these venues suggests that they may serve very different clientele – either predominantly male or predominantly female. In contrast, *Wine Bars*, which exhibit a similarly mixed clientele in the aggregate seem to also exhibit this mixed gender composition at the venue level. Although there are some small differences in the gender compositions of subcategories for different cities, their relative means and standard deviations tend to be stable.

Figure 27: Proportion of Females in Venues by Subcategory

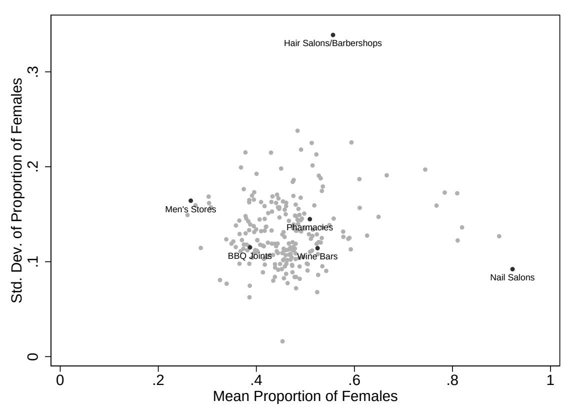

Note: This scatter plot pools venues from all cities in the sample. Each dot represents all of the venues within a subcategory.

Full summary statistics disaggregated by subcategory can be found in the table below.

		Proportion of Females			Proportion of Youth				
Category	Subcategory	μ	σ	$p_{75} - p_{25}$	μ	σ	$p_{75} - p_{25}$	Venues	Check-ins
Bars	Bars	0.44	0.11	0.12	0.73	0.15	0.19	2114	2,379,893
Bars	Beer Gardens	0.42	0.10	0.14	0.71	0.16	0.22	104	163,102
Bars	Breweries	0.39	0.10	0.13	0.63	0.16	0.21	127	193,472
Bars	Cocktail Bars	0.47	0.10	0.12	0.72	0.15	0.18	322	368,735
Bars	Dive Bars	0.40	0.11	0.12	0.70	0.17	0.22	360	271,981
Bars	Gastropubs	0.46	0.08	0.09	0.71	0.11	0.12	212	332,930
Bars	Hookah Bars	0.44	0.11	0.16	0.89	0.07	0.09	104	48,287
Bars	Hotel Bars	0.42	0.12	0.14	0.58	0.16	0.23	243	118,877
Bars	Karaoke Bars	0.47	0.12	0.13	0.77	0.14	0.20	179	108,432
Bars	Lounges	0.46	0.13	0.14	0.70	0.18	0.22	530	394,834
Bars	Nightclubs	0.43	0.13	0.16	0.78	0.14	0.16	463	380,946
Bars	Other Nightlife	0.48	0.16	0.20	0.67	0.18	0.24	69	31,027
Bars	Pubs	0.44	0.09	0.10	0.71	0.15	0.16	517	693,208
Bars	Sake Bars	0.47	0.10	0.10	0.72	0.14	0.18	19	12,780
Bars	Speakeasies	0.47	0.11	0.13	0.75	0.14	0.16	90	100,912
Bars	Sports Bars	0.43	0.11	0.13	0.71	0.16	0.23	404	545,632
Bars	Strip Clubs	0.30	0.17	0.22	0.65	0.18	0.22	145	38,122
Bars	Whisky Bars	0.43	0.08	0.10	0.74	0.13	0.12	45	75,513
Bars	Wine Bars	0.52	0.11	0.13	0.66	0.15	0.17	347	243,684
Bars	Wineries	0.51	0.13	0.17	0.60	0.16	0.23	47	14,709
Cafes	Cafeterias	0.37	0.20	0.36	0.60	0.23	0.36	41	17,696
Cafes	Cafes	0.48	0.14	0.18	0.69	0.17	0.21	1280	699,162
Cafes	Coffee Shops	0.48	0.14	0.18	0.67	0.16	0.22	3162	3,160,881
Entertainment	Aquariums	0.54	0.09	0.12	0.69	0.08	0.11	17	35,862
Entertainment	Arcades	0.44	0.17	0.21	0.71	0.16	0.18	78	38,214
Entertainment	Art Galleries	0.49	0.15	0.17	0.69	0.16	0.18	233	41,570
Entertainment	Art Museums	0.47	0.11	0.12	0.65	0.12	0.10	133	296,219

		Proportion of Females			Proportion of Youth				
Category	Subcategory	μ	σ	$p_{75}-p_{25}$	μ	σ	$p_{75}-p_{25}$	Venues	Check-ins
Entertainment	Bowling Alleys	0.44	0.12	0.15	0.63	0.14	0.20	144	106,817
Entertainment	Casinos	0.34	0.12	0.18	0.59	0.12	0.18	12	15,262
Entertainment	Comedy Clubs	0.47	0.12	0.12	0.74	0.12	0.14	115	87,982
Entertainment	Concert Halls	0.46	0.11	0.12	0.67	0.17	0.23	116	166,265
Entertainment	General Entertainment	0.46	0.15	0.19	0.65	0.16	0.21	822	367,113
Entertainment	Historic Sites	0.39	0.14	0.18	0.59	0.17	0.21	154	135,739
Entertainment	History Museums	0.45	0.13	0.14	0.57	0.14	0.16	165	116,896
Entertainment	Indie Movie Theaters	0.47	0.11	0.11	0.65	0.13	0.16	92	85,124
Entertainment	Indie Theaters	0.48	0.13	0.15	0.70	0.15	0.18	70	36,177
Entertainment	Jazz Clubs	0.42	0.10	0.11	0.60	0.12	0.20	80	55,955
Entertainment	Movie Theaters	0.45	0.10	0.11	0.67	0.12	0.14	180	275,574
Entertainment	Multiplexes	0.49	0.08	0.08	0.69	0.09	0.11	106	432,959
Entertainment	Museums	0.47	0.11	0.12	0.61	0.12	0.17	174	152,003
Entertainment	Music Venues	0.42	0.12	0.13	0.69	0.16	0.21	295	371,347
Entertainment	Performing Arts Venues	0.49	0.14	0.14	0.67	0.16	0.23	189	149,840
Entertainment	Piano Bars	0.52	0.07	0.09	0.74	0.15	0.32	15	14,778
Entertainment	Pool Halls	0.42	0.16	0.18	0.74	0.20	0.22	39	23,962
Entertainment	Public Art	0.39	0.17	0.23	0.59	0.19	0.29	52	48,019
Entertainment	Racetracks	0.40	0.17	0.21	0.54	0.18	0.19	49	21,480
Entertainment	Rock Clubs	0.45	0.09	0.10	0.73	0.12	0.15	106	172,387
Entertainment	Science Museums	0.46	0.10	0.09	0.63	0.12	0.17	74	153,446
Entertainment	Stadiums	0.40	0.11	0.12	0.59	0.13	0.17	31	252,616
Entertainment	Theaters	0.48	0.13	0.15	0.66	0.15	0.19	482	333,304
Entertainment	Water Parks	0.49	0.15	0.17	0.59	0.19	0.25	9	2,936

		Proportion of Females			Proportion of Youth				
Category	Subcategory	μ	σ	$p_{75}-p_{25}$	μ	σ	$p_{75}-p_{25}$	Venues	Check-ins
Entertainment	Zoos	0.47	0.09	0.10	0.52	0.11	0.13	157	90,529
Food	African Restaurants	0.50	0.13	0.16	0.68	0.15	0.25	25	7,774
Food	American Restaurants	0.47	0.11	0.13	0.63	0.15	0.20	2444	1,845,916
Food	Argentinian Restaurants	0.44	0.08	0.11	0.65	0.11	0.12	37	14,020
Food	Asian Restaurants	0.46	0.11	0.14	0.70	0.14	0.19	779	357,845
Food	Australian Restaurants	0.53	0.11	0.11	0.82	0.09	0.13	14	14,714
Food	Bakeries	0.53	0.12	0.14	0.68	0.14	0.18	910	555,341
Food	BBQ Joints	0.39	0.12	0.16	0.60	0.16	0.23	453	297,253
Food	Brazilian Restaurants	0.41	0.09	0.11	0.67	0.10	0.15	66	40,859
Food	Breakfast Spots	0.48	0.11	0.13	0.64	0.14	0.19	661	405,581
Food	Burger Joints	0.40	0.11	0.13	0.65	0.14	0.19	1283	913,425
Food	Burrito Places	0.38	0.11	0.14	0.71	0.15	0.20	175	117,843
Food	Caribbean Restaurants	0.48	0.12	0.17	0.64	0.12	0.16	93	32,959
Food	Cuban Restaurants	0.46	0.10	0.14	0.66	0.14	0.15	106	86,381
Food	Cupcake Shops	0.59	0.13	0.15	0.71	0.14	0.14	146	112,410
Food	Delis / Bodegas	0.38	0.15	0.20	0.67	0.17	0.24	824	343,274
Food	Dim Sum Restaurants	0.46	0.08	0.09	0.71	0.10	0.14	79	62,700
Food	Diners	0.42	0.10	0.13	0.63	0.15	0.20	654	447,090

		Proportion of Females			Proportion of Youth				
Category	Subcategory	μ	σ	$p_{75} - p_{25}$	μ	σ	$p_{75} - p_{25}$	Venues	Check-ins
Food	Donut Shops	0.39	0.14	0.18	0.64	0.17	0.23	180	91,355
Food	Eastern European Restaurants	0.46	0.10	0.12	0.69	0.10	0.16	65	28,409
Food	Ethiopian Restaurants	0.47	0.10	0.11	0.74	0.10	0.12	50	15,807
Food	Falafel Restaurants	0.38	0.12	0.15	0.73	0.15	0.17	93	44,572
Food	Fast Food Restaurants	0.42	0.14	0.18	0.64	0.15	0.19	2422	647,960
Food	Filipino Restaurants	0.46	0.10	0.10	0.67	0.11	0.13	16	11,933
Food	Food Trucks	0.40	0.13	0.17	0.72	0.14	0.17	62 4	210,581
Food	French Restaurants	0.50	0.09	0.12	0.65	0.13	0.16	514	330,423
Food	Fried Chicken Joints	0.38	0.12	0.15	0.62	0.14	0.20	339	80,861
Food	German Restaurants	0.39	0.06	0.07	0.66	0.11	0.13	60	102,679
Food	Greek Restaurants	0.45	0.11	0.17	0.65	0.14	0.18	233	99,018
Food	Hot Dog Joints	0.37	0.11	0.14	0.63	0.14	0.19	272	172,987
Food	Ice Cream Shops	0.53	0.12	0.15	0.69	0.15	0.21	750	409,897
Food	Indian Restaurants	0.39	0.10	0.14	0.68	0.13	0.17	525	208,403
Food	Italian Restaurants	0.50	0.11	0.13	0.63	0.14	0.18	1828	897,952

		Proportion of Females			Proportion of Youth				
Category	Subcategory	μ	σ	$p_{75} - p_{25}$	μ	σ	$p_{75} - p_{25}$	Venues	Check-ins
Food	Japanese Restaurants	0.47	0.11	0.14	0.69	0.14	0.19	779	335,856
Food	Juice Bars	0.52	0.13	0.16	0.74	0.13	0.17	322	159,643
Food	Korean Restaurants	0.48	0.10	0.13	0.78	0.11	0.14	417	235,319
Food	Latin American Restaurants	0.46	0.11	0.14	0.71	0.14	0.17	207	100,725
Food	Malaysian Restaurants	0.48	0.12	0.11	0.75	0.10	0.14	22	14,836
Food	Mediterranean Restaurants	0.44	0.12	0.17	0.69	0.13	0.18	371	177,166
Food	Mexican Restaurants	0.44	0.12	0.16	0.64	0.16	0.22	2361	1,301,614
Food	Middle Eastern Restaurants	0.41	0.12	0.16	0.70	0.13	0.18	236	73,273
Food	Mongolian Restaurants	0.48	0.11	0.15	0.66	0.12	0.22	9	2,290
Food	Moroccan Restaurants	0.47	0.11	0.11	0.73	0.11	0.14	28	5,678
Food	New American Restaurants	0.49	0.09	0.12	0.66	0.13	0.17	366	393,351
Food	Peruvian Restaurants	0.48	0.08	0.07	0.70	0.10	0.13	21	22,046
Food	Pizza Places	0.40	0.12	0.16	0.69	0.15	0.20	1993	841,333
Food	Portuguese Restaurants	0.53	0.09	0.15	0.76	0.05	0.09	6	7,238

		Proportion of Females			Proportion of Youth				
Category	Subcategory	μ	σ	$p_{75} - p_{25}$	μ	σ	$p_{75} - p_{25}$	Venues	Check-ins
Food	Ramen / Noodle House	0.46	0.09	0.10	0.75	0.11	0.13	228	203,034
Food	Salad Shop	0.51	0.11	0.16	0.74	0.13	0.17	184	151,994
Food	Sandwich Places	0.38	0.13	0.17	0.70	0.15	0.19	2265	888,057
Food	Scandinavian Restaurants	0.48	0.08	0.11	0.70	0.10	0.18	20	18,902
Food	Seafood Restaurants	0.47	0.10	0.11	0.62	0.13	0.18	550	403,785
Food	Soup Places	0.51	0.11	0.16	0.75	0.10	0.15	63	45,932
Food	South American Restaurants	0.44	0.09	0.12	0.69	0.10	0.13	63	18,152
Food	Southern / Soul Food Restaurants	0.48	0.11	0.12	0.60	0.15	0.18	172	143,823
Food	Spanish Restaurants	0.49	0.10	0.13	0.67	0.12	0.15	76	38,048
Food	Steakhouses	0.44	0.10	0.12	0.56	0.12	0.16	458	309,250
Food	Sushi Restaurants	0.50	0.10	0.12	0.71	0.14	0.18	1207	512,255
Food	Tapas Restaurants	0.53	0.10	0.11	0.72	0.13	0.14	139	119,025
Food	Tea Rooms	0.58	0.13	0.15	0.78	0.15	0.17	253	187,002
Food	Thai Restaurants	0.46	0.11	0.13	0.71	0.13	0.17	815	316,483
Food	Turkish Restaurants	0.46	0.10	0.14	0.72	0.10	0.10	27	13,034
Food	Vegetarian / Vegan Restaurants	0.52	0.12	0.13	0.70	0.12	0.13	329	194,472

		Proportion of Females			Proportion of Youth				
Category	Subcategory	μ	σ	$p_{75} - p_{25}$	μ	σ	$p_{75} - p_{25}$	Venues	Check-ins
Food	Vietnamese Restaurants	0.45	0.10	0.12	0.73	0.12	0.15	392	173,671
Food	Wings Joints	0.43	0.11	0.14	0.72	0.13	0.17	226	139,954
Food	Yogurt	0.59	0.11	0.13	0.77	0.11	0.13	73	37,867
Gyms	Baseball Fields	0.41	0.14	0.17	0.58	0.20	0.26	121	25,378
Gyms	Baseball Courts	0.39	0.16	0.22	0.62	0.21	0.27	38	7,981
Gyms	Dance Studios	0.67	0.19	0.28	0.78	0.14	0.18	85	45,293
Gyms	Golf Courses	0.31	0.16	0.24	0.56	0.17	0.24	295	83,751
Gyms	Gyms	0.49	0.22	0.27	0.69	0.19	0.27	640	984,359
Gyms	Martial Arts Dojos	0.48	0.24	0.12	0.67	0.21	0.28	18	13,145
Gyms	Skate Parks	0.30	0.16	0.22	0.63	0.17	0.19	25	4,948
Gyms	Skating Rinks	0.45	0.16	0.18	0.59	0.17	0.21	65	33,786
Gyms	Soccer Fields	0.40	0.13	0.19	0.62	0.24	0.29	40	10,549
Gyms	Tennis Courts	0.42	0.13	0.15	0.60	0.16	0.18	45	15,974
Gyms	Tracks	0.48	0.14	0.16	0.67	0.16	0.25	27	34,870
Gyms	Yoga Studios	0.77	0.16	0.16	0.72	0.17	0.25	226	154,706
Hotels	Bed & Breakfasts	0.39	0.07	0.10	0.60	0.17	0.28	10	1,083
Hotels	Hotels Pools	0.43	0.15	0.20	0.60	0.15	0.23	24	4,453
Hotels	Hotels	0.40	0.11	0.12	0.59	0.15	0.16	1637	2,203,596
Hotels	Motels	0.36	0.12	0.15	0.65	0.16	0.19	111	22,408
Hotels	Resorts	0.39	0.17	0.16	0.55	0.18	0.35	16	9,713
Outdoors	Beaches	0.45	0.15	0.17	0.58	0.16	0.19	187	120,023
Outdoors	Cemeteries	0.50	0.15	0.23	0.52	0.16	0.23	88	26,666
Outdoors	Cities	0.49	0.14	0.18	0.60	0.15	0.21	255	567,323
Outdoors	Dog Runs	0.48	0.19	0.23	0.61	0.19	0.26	167	81,925
Outdoors	Farms	0.48	0.14	0.27	0.57	0.16	0.17	28	6,522

		Proportion of Females			Proportion of Youth				
Category	Subcategory	μ	σ	$p_{75} - p_{25}$	μ	σ	$p_{75} - p_{25}$	Venues	Check-ins
Outdoors	Fields	0.42	0.15	0.16	0.60	0.21	0.29	72	27,714
Outdoors	Gardens	0.47	0.14	0.14	0.57	0.16	0.20	131	63,519
Outdoors	Harbors / Marinas	0.43	0.12	0.16	0.59	0.17	0.23	175	98,020
Outdoors	Lakes	0.42	0.14	0.18	0.52	0.19	0.29	93	61,825
Outdoors	Monuments / Landmarks	0.37	0.12	0.14	0.55	0.13	0.19	110	233,864
Outdoors	Mountains	0.46	0.16	0.26	0.58	0.15	0.23	15	2,798
Outdoors	Neighborhoods	0.45	0.16	0.19	0.61	0.17	0.23	558	1,069,034
Outdoors	Other Great Outdoors	0.45	0.17	0.22	0.58	0.19	0.27	377	189,642
Outdoors	Parks	0.43	0.17	0.22	0.58	0.19	0.26	1329	1,176,399
Outdoors	Playgrounds	0.45	0.16	0.19	0.53	0.17	0.23	348	74,294
Outdoors	Plazas	0.39	0.16	0.21	0.58	0.19	0.24	317	532,148
Outdoors	Pools	0.47	0.18	0.22	0.63	0.21	0.30	132	33,801
Outdoors	Rivers	0.41	0.12	0.16	0.56	0.15	0.22	16	12,557
Outdoors	Scenic Lookouts	0.41	0.16	0.17	0.56	0.18	0.23	247	158,037
Outdoors	Sculpture Gardens	0.37	0.18	0.23	0.51	0.19	0.27	145	76,976
Outdoors	Ski Areas	0.45	0.02	0.03	0.61	0.13	0.27	3	1,792
Outdoors	Vineyards	0.61	0.16	0.22	0.63	0.22	0.32	2	181
Shops/Services	Accessories Stores	0.52	0.21	0.30	0.70	0.14	0.17	189	29,879
Shops/Services	Arts & Crafts Stores	0.65	0.15	0.21	0.66	0.13	0.19	295	89,275
Shops/Services	Automotive Shops	0.39	0.13	0.17	0.58	0.16	0.22	838	129,367

		Proportion of Females			Proportion of Youth				
Category	Subcategory	μ	σ	$p_{75} - p_{25}$	μ	σ	$p_{75} - p_{25}$	Venues	Check-ins
Shops/Services	Banks	0.43	0.16	0.23	0.65	0.18	0.26	1546	328,903
Shops/Services	Bike Shops	0.35	0.12	0.14	0.65	0.17	0.23	199	34,493
Shops/Services	Board Shops	0.37	0.12	0.18	0.71	0.17	0.26	60	8,419
Shops/Services	Bookstores	0.44	0.16	0.20	0.65	0.16	0.19	339	193,775
Shops/Services	Boutiques	0.59	0.23	0.38	0.74	0.14	0.19	487	108,881
Shops/Services	Bridal Shops	0.90	0.13	0.08	0.84	0.08	0.10	56	10,276
Shops/Services	Butchers	0.37	0.10	0.12	0.54	0.17	0.19	41	13,392
Shops/Services	Camera Stores	0.35	0.12	0.14	0.67	0.1	0.12	26	9,480
Shops/Services	Candy Stores	0.53	0.12	0.15	0.65	0.15	0.15	146	67,379
Shops/Services	Car Dealerships	0.38	0.14	0.20	0.59	0.15	0.19	182	35,351
Shops/Services	Car Wash	0.40	0.11	0.15	0.52	0.13	0.17	92	23,370
Shops/Services	Cheese Shops	0.50	0.10	0.11	0.66	0.12	0.19	32	23,550
Shops/Services	Clothing Stores	0.53	0.19	0.26	0.73	0.14	0.17	1498	644,653
Shops/Services	Cosmetics Shops	0.81	0.17	0.19	0.72	0.14	0.18	684	156,441
Shops/Services	Department Stores	0.59	0.12	0.15	0.64	0.11	0.14	757	997,837
Shops/Services	Design Studios	0.46	0.16	0.19	0.68	0.16	0.19	221	42,758
Shops/Services	Drugstores / Pharmacies	0.51	0.14	0.20	0.62	0.16	0.24	1525	577,272
Shops/Services	Electronics Stores	0.36	0.14	0.17	0.64	0.16	0.19	508	381,047
Shops/Services	Farmers Markets	0.51	0.12	0.16	0.58	0.15	0.20	199	147,509
Shops/Services	Financial or Legal Services	0.38	0.22	0.22	0.63	0.21	0.30	26	15,416
Shops/Services	Flea Markets	0.53	0.14	0.16	0.69	0.13	0.17	91	30,724
Shops/Services	Flower Shops	0.49	0.17	0.24	0.62	0.19	0.24	52	8,407
Shops/Services	Food & Drink Shops	0.46	0.15	0.21	0.62	0.18	0.28	122	40,589

		Proportion of Females			Proportion of Youth				
Category	Subcategory	μ	σ	$p_{75} - p_{25}$	μ	σ	$p_{75} - p_{25}$	Venues	Check-ins
Shops/Services	Food Courts	0.42	0.15	0.22	0.67	0.17	0.20	168	81,914
Shops/Services	Gaming Cafes	0.28	0.16	0.19	0.66	0.24	0.19	14	5,084
Shops/Services	Garden Centers	0.52	0.16	0.19	0.54	0.16	0.26	16	3,338
Shops/Services	Gift Shops	0.53	0.17	0.24	0.64	0.15	0.20	384	74,114
Shops/Services	Gourmet Shops	0.51	0.13	0.13	0.65	0.16	0.22	140	121,390
Shops/Services	Grocery Stores	0.49	0.13	0.17	0.62	0.15	0.21	1903	1,749,221
Shops/Services	Gyms or Fitness Centers	0.51	0.23	0.27	0.69	0.18	0.23	463	1,005,051
Shops/Services	Hardware Stores	0.37	0.11	0.12	0.54	0.15	0.22	321	142,828
Shops/Services	Health Food Stores	0.43	0.21	0.36	0.70	0.14	0.20	20	2,558
Shops/Services	Hobby Shops	0.40	0.19	0.31	0.65	0.16	0.21	69	22,991
Shops/Services	Jewelry Stores	0.61	0.19	0.27	0.69	0.15	0.23	178	39,630
Shops/Services	Kids Stores	0.63	0.13	0.21	0.59	0.13	0.15	59	8,974
Shops/Services	Lingerie Stores	0.81	0.12	0.14	0.76	0.13	0.14	145	46,643
Shops/Services	Liquor Stores	0.38	0.15	0.19	0.66	0.17	0.23	381	109,533
Shops/Services	Malls	0.49	0.16	0.17	0.63	0.17	0.19	302	779,691
Shops/Services	Markets	0.48	0.07	0.13	0.64	0.12	0.20	19	67,220
Shops/Services	Men's Stores	0.27	0.16	0.17	0.70	0.15	0.19	244	50,261
Shops/Services	Miscellaneous Shops	0.54	0.18	0.25	0.63	0.16	0.20	750	213,273
Shops/Services	Mobile Phone Shops	0.40	0.12	0.16	0.67	0.15	0.20	148	22,030
Shops/Services	Motorcycle Shops	0.33	0.08	0.13	0.50	0.15	0.14	10	1,718
Shops/Services	Music Stores	0.37	0.13	0.13	0.65	0.12	0.18	62	14,940
Shops/Services	Nail Salons	0.92	0.09	0.06	0.77	0.14	0.18	163	32,808
Shops/Services	Newsstands	0.37	0.14	0.14	0.53	0.24	0.37	24	3,743

		Proportion of Females			Proportion of Youth				
Category	Subcategory	μ	σ	$p_{75} - p_{25}$	μ	σ	$p_{75} - p_{25}$	Venues	Check-ins
Shops/Services	Optical Shops	0.52	0.13	0.19	0.70	0.16	0.19	76	14,180
Shops/Services	Paper / Office Supplies Stores	0.48	0.15	0.18	0.60	0.17	0.25	365	78,112
Shops/Services	Pet Stores	0.58	0.13	0.17	0.59	0.16	0.22	362	99,574
Shops/Services	Record Shops	0.34	0.08	0.09	0.64	0.12	0.14	123	49,108
Shops/Services	Salons / Barbershops	0.56	0.34	0.67	0.71	0.16	0.21	860	187,652
Shops/Services	Shoe Stores	0.51	0.20	0.32	0.71	0.14	0.18	500	117,650
Shops/Services	Smoke Shops	0.26	0.15	0.23	0.54	0.20	0.35	73	19,403
Shops/Services	Spas / Massages	0.78	0.17	0.23	0.71	0.14	0.19	526	114,510
Shops/Services	Sporting Goods Shops	0.41	0.13	0.15	0.62	0.14	0.18	375	155,592
Shops/Services	Tanning Salons	0.74	0.20	0.29	0.81	0.15	0.22	80	18,677
Shops/Services	Tattoo Parlors	0.55	0.14	0.17	0.74	0.16	0.17	138	19,843
Shops/Services	Thrift / Vintage Stores	0.56	0.15	0.19	0.69	0.16	0.23	319	62,505
Shops/Services	Toy / Game Stores	0.48	0.15	0.20	0.60	0.15	0.20	204	112,059
Shops/Services	Video Game Stores	0.29	0.11	0.13	0.71	0.15	0.19	147	27,089
Shops/Services	Video Stores	0.45	0.20	0.23	0.66	0.21	0.23	48	8,928
Shops/Services	Wine Shops	0.46	0.13	0.16	0.66	0.18	0.23	191	56,436
Shops/Services	Women's Stores	0.82	0.14	0.15	0.78	0.13	0.16	358	80,636
Spiritual	Churches	0.48	0.17	0.23	0.58	0.19	0.28	671	264,923
Spiritual	Synagogues	0.53	0.19	0.28	0.56	0.21	0.30	43	12,120
Spiritual	Temples	0.44	0.16	0.17	0.63	0.17	0.21	31	8,202