Regional Science and Urban Economics 77 (2019) 253–263

Contents lists available at ScienceDirect

Regional Science and Urban Economics

journal homepage: www.elsevier.com/locate/regec

Gender segregation within neighborhoods☆

Gregorio Caetano[a,b,*], Vikram Maheshri[a,b]

[a] University of Georgia, USA
[b] University of Houston, USA

ARTICLE INFO

JEL:
R1
R2
R3
J1
J3

Keywords:
Gender segregation
Homophily
Peer groups
Urban sorting
Diversity

ABSTRACT

Homophily generates segregation, which reduces diversity in peer groups and leads to narrower social interactions. Using novel data from Foursquare, a popular mobile app that documents the activity of millions of people, we document robust, highly localized gender segregation within neighborhoods: most venues (e.g., shops, restaurants, parks, museums) in eight major US cities are highly gender segregated, and over half of the gender segregation across cities occurs within Census blocks. This segregation is mostly driven by venue offerings, not discriminatory preferences. A higher variety of venues on a block attracts more gender-balanced visitors, but, perversely, more intense sorting across those venues ultimately leads to more segregated venues. Hence top-down approaches to facilitate diverse interactions may be counterproductive. We find similar results for segregation by age.

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 widely studied in a number of important contexts such as residential neighborhoods, schools and workplaces (Card et al., 2008; 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 instance, neighbors may shop at different supermarkets, students may select different extracurricular activities, and coworkers may exercise at different gyms. Each of these mundane decisions may seem inconsequential in isolation, but in the aggregate, they shape our social experience. Unfortunately, they are difficult to observe at scale.

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 use this data to report a series of patterns of gender segregation in venues.

First, we document that interactions in venues are far more highly gender segregated than neighborhood-level measures would suggest. Indeed, 80–90 percent of gender segregation in venues (defined as the unevenness in the distribution of men and women across venues) is observed within Census tracts, and over half of that is observed within Census blocks. This suggests that actual social interactions may be substantially narrower than what researchers infer based on more readily available data that is, at best, disaggregated to the Census block level.

Next, we study how local diversity is shaped by the venues in a neighborhood. Specifically, we estimate the causal impact of an

☆ 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.

* Corresponding author. University of Georgia, USA.
E-mail address: gregorio.caetano@uga.edu (G. Caetano).

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

https://doi.org/10.1016/j.regsciurbeco.2019.05.004
Received 21 November 2018; Received in revised form 15 May 2019; Accepted 18 May 2019
Available online 31 May 2019

increase in the variety of venues on segregation. All else constant, a block with a larger variety of venues attracts a more diverse group of people, but once there, individuals sort more intensely across venues on the block. This perversely leads to an increase in gender segregation within venues. We term this result *the paradox of diversity*: as people are supplied with a more diverse set of options, their homophilic choices expose them to a less diverse set of peers. This paradox suggests denser urban areas may unwittingly foster narrower social interactions.

Many broad questions in social science are difficult to approach comprehensively because researchers cannot observe all of the choices that jointly determine individuals' social interactions. Our work joins a number of recent studies that have begun to use user-generated data to help close this gap, such as Davis et al. (2018) (Yelp!) and Couture (2014) (Google Places). Still, 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. To allay these concerns, we provide a detailed sensitivity analysis in an appendix that includes a Monte Carlo study of the role of measurement error in the context of our paper. We also replicate our analysis for age segregation and re-analyze segregation in each city separately. We obtain similar results in all these scenarios, which may speak to the general pull of homophily that permeates social interaction.[2]

The remainder of the paper is organized as follows. In Section 2, we describe our data set, and in Section 3, we show evidence of highly local gender segregation in venue choices. In Section 4, we explore the determinants of this phenomenon, and we show robust evidence for the paradox of diversity. In Section 5, we discuss the external validity of our findings to other environments and along dimensions other than gender. We conclude in Section 6.

2. Data

In order to analyze how people sort in venues within neighborhoods, we use 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 that allows us to extract a meaningful signal about the sorting of all individuals across venues from this novel dataset.[3] Ours is the first paper to use this large and highly detailed database of venue visitors to study diversity within neighborhoods.[4] We observe 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 sample regions are defined as the counties in which these cities are primarily located.[5] 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.[6] 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.[7] The 9 categories of venues are further classified into 225 narrow subcategories; detailed summary statistics disaggregated by subcategory can be found in the appendix.

Because we observe daily check-ins at each venue, we can check for systematic trends in our data over the sample period. Although there is substantial day-of-week variation in the number of check-ins (Panel (a) of Fig. 1), the gender composition of check-ins is nearly constant. This suggests that we can aggregate the data at least to the weekly level. There is no systematic weekly variation in check-in frequency, no discernible seasonality or aggregate trend, and the gender composition of check-ins is roughly constant throughout the sample period (Panel (b) of Fig. 1). This suggests that we can aggregate the data set to the annual level.[8]

3. Measuring gender segregation within neighborhoods

The widely used Theil (1967) index of segregation is particularly well-suited to measure segregation within neighborhoods.[9] 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.[10] If the neighborhood is fully integrated, then all of its venues will have the same gender composition as the neighborhood overall, and

[2] These sensitivity analyses are included in the supplementary appendix available at https://bit.ly/2C7Vf1b.

[3] We provide evidence that our main results are qualitatively robust to men and women having different propensities to "check-in" in the appendix.

[4] A small number of studies (e.g., Arribas-Bel and Bakens, 2014) have used Foursquare data obtained indirectly via the Foursquare API (application programming interface). Foursquare data obtained via the API does not disaggregate check-ins along any demographic dimension.

[5] 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.

[6] We also only consider check-ins from users who have specified their gender (97.6% of check-ins have a specified gender). None of the restrictions described here seem to bias our results (see appendix).

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

[8] 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).

[9] 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.

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

Table 1
Summary statistics.

City	Venues	Check-ins	μ	σ	$p_{75}-p_{25}$	Tracts	B. Groups	Blocks
Atlanta	4115	2.84	0.46	0.17	0.19	180	361	1307
Chicago	13,665	8.11	0.49	0.16	0.19	1100	2235	6237
Dallas	5065	2.40	0.45	0.16	0.19	421	774	1986
Los Angeles	23,108	10.2	0.46	0.15	0.18	1902	3584	9182
New York City	16,203	16.2	0.49	0.17	0.19	282	945	2501
Philadelphia	3933	2.10	0.47	0.16	0.19	301	568	1757
San Francisco	6601	4.78	0.42	0.15	0.16	182	440	1898
Washington, DC	3687	2.98	0.43	0.16	0.17	152	272	1069
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	6441	6.52	0.44	0.12	0.13	20		
Outdoors	4795	4.62	0.44	0.16	0.21	22		
Cafes	4483	3.88	0.47	0.14	0.18	3		
Entertainment	4189	4.08	0.46	0.13	0.15	29		
Hotels	1798	2.24	0.40	0.11	0.13	5		
Gyms	1625	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 25th and 75th percentiles of the proportion of females in venues.

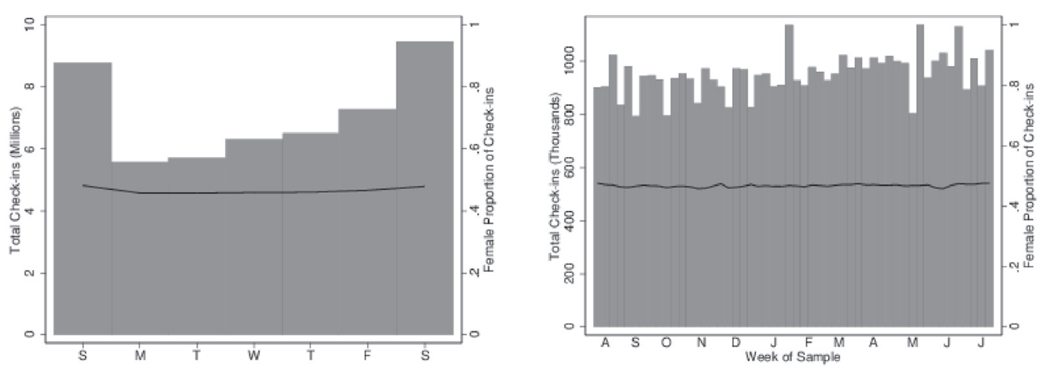

(a) Day-of-Week Variation (b) Week-of-Year Variation

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.

Fig. 1. Check-ins and gender composition over time.

$T_k = 0$. Neighborhoods with less diverse venues have larger values of T_k.[11] 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.

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 Fig. 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, $T_b \le T_g \le T_t$ 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 segregated 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.[12] Formally, we can decompose the total Theil index of city c into within- and

[11] 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.

[12] 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. In addition, as Shorrocks (1980) points out, many other commonly used measures of segregation, diversity, exposure or inequality with other attractive properties are not additively decomposable, so they are less useful and appropriate in our context.

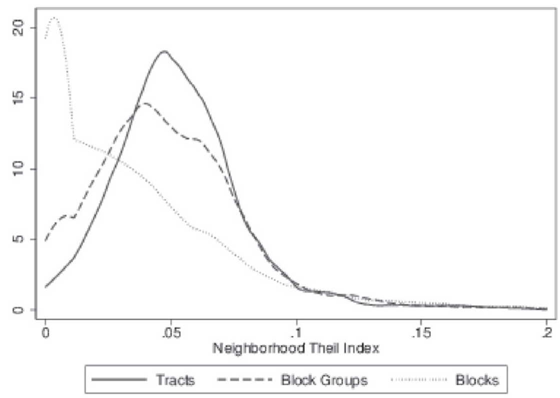

Notes: All densities are estimated using a bandwidth of 0.005 and an Epanechnikov kernel. For clarity, we present the density only for values of the domain less than 0.2; fewer than 1% of neighborhoods of any type have a Theil Index in excess of 0.2. Theil Indices are pooled across neighborhoods in all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

Fig. 2. Densities of Theil indices for various neighborhood definitions.

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.

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 segregation within neighborhoods independently of how individuals sort across neighborhoods. In Table 2, we present the proportion of city-wide gender segregation in venues within neighborhoods, i.e., the contribution of the first term of equation (2).[13] Intuitively, this captures how much of the variation in the gender composition of city venues is "local." It is immediate that the majority of gender segregation in city venues is highly localized.

To better interpret the measures in Table 2, we benchmark the observed gender compositions of venues against the gender compositions of venues that would be hypothetically observed in the absence of within-neighborhood segregation.[14] This exercise reveals how much additional segregation can be measured because we observe sorting across venues as opposed to only sorting across neighborhoods (as in the vast majority of studies). By observing sorting at the more disaggregated venue level, we measure 2–4 times more segregation than in data aggregated to the block level, and 4–12 times more segregation than in data aggregated to the tract level.[15] For Manhattan, these numbers are on the higher end: we measure 4 (12) times more segregation than we would have with data aggregated to the block (tract) level.[16]

[13] 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.

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

[15] 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).

[16] The amount of gender segregation in venues that we find is comparable to the amount of residential segregation along other demographic dimensions found in Fischer et al. (2004). After appropriately rescaling all measures for comparison, we find that in a representative tract with a Theil of 0.05, the extent to which women are segregated in venues is the following percentage of the extent to which these various demographic groups are segregated residentially on average: 65% for Blacks, 93% for Whites, 125% for Hispanics, 224% for foreign born individuals, 154% for top quintile earners, 181% for bottom quintile earners, 107% for homeowners, 330% for married households, 801% for households with children under 15, 362% of households headed by somebody aged 18–29, and 374% for households headed by somebody older than 64.

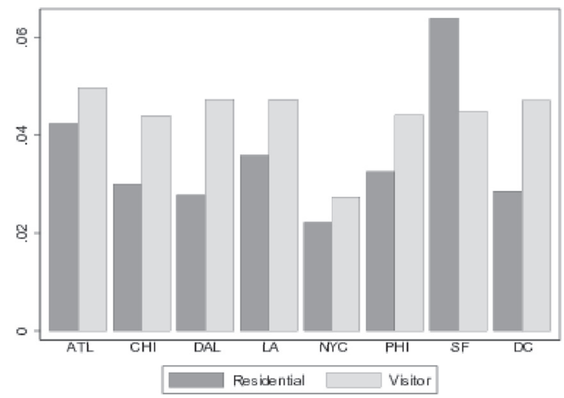

Note: Residential segregation is calculated as the Theil index of the gender composition of block residents from the 2010 Census. For comparability, visitor segregation is calculated as the Theil index of the gender composition of check-ins in blocks. Bootstrapped standard errors for all estimates are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

Fig. 3. Residential Segregation vs. Visitor Segregation.

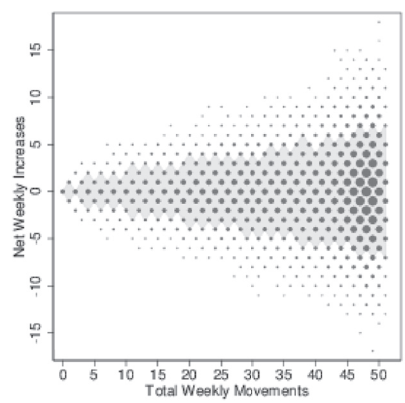

Notes: In this scatter plot of venues in our data, larger dots correspond to a greater numbers of venues. A venue experiences a weekly increase (decrease) in gender composition if the proportion of female check-ins rises (falls) by at least one percentage point.

Fig. 4. Segregation dynamics.

Of course, researchers typically observe only broader location choices that individuals make such as the neighborhoods where they reside. In Fig. 3, we compare how residents sort across blocks with how visitors sort across blocks for each city in our sample. Residential segregation is calculated as the Theil index of the gender composition of block residents for each city from the 2010 Census. Visitor segregation 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 segregation than visitor segregation.[17] As a result, our findings suggest that studies relying on residential data alone may overstate individuals' exposure to diversity. We have reason to believe that we also underestimate the extent to which gender segregation within blocks mitigates exposure to diversity because 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. This leads us to speculate that men and women may also 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.

3.1. Segregation dynamics

There are two explanations for our findings: men (women) prefer to be in the company of other men (women) in venues ("active" segregation); and men and women systematically prefer different types of venues ("passive" segregation). Active segregation can lead to "tipping" and other dynamics made famous in the seminal Schelling (1971) model. We find little evidence for active segregation, which we summarize in Fig. 4. Intuitively, if there is active segregation, we should expect certain venues to attract increasing shares of men and other venues to attract increasing shares of women. However, when we analyze trends in gender composition for venues over time, that is not what

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

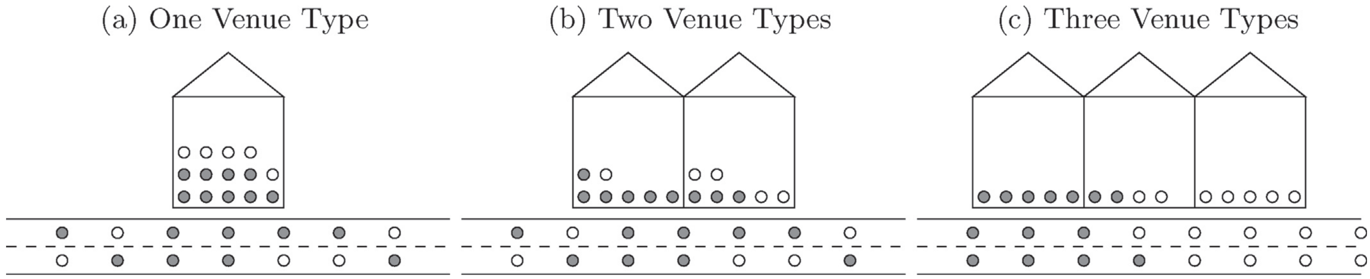

Fig. 5. The paradox of diversity.

we find. 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 Fig. 1.[18] 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 over our sample period.[19] This suggests that most venues' gender composition simply bounces around a stable level. We conclude that the bulk of the gender segregation that we observe in venues is passive.

4. Venue variety and diversity

It is often suggested that flexibly designed neighborhoods with a variety of different offerings encourage diverse interactions.[20] However, this need not be the case –a greater variety of venue offerings may actually discourage diverse interactions between people. We term this the *paradox of diversity*. We provide some intuition for this result graphically in Fig. 5. In the first panel, we present a neighborhood with one type of venue. The people in the neighborhood (shown on the street) have no choice of where to go, so they all end up in the same venue (shown in the building). In the second panel, a second type of venue opens. Now the people in the neighborhood can choose where to go. Since the venue offerings are different and different types of people have a tendency to like different things, there will be some sorting across venues and everybody will be exposed to less diverse peers. In the third panel, several more types of venues open up. Now, the neighborhood becomes attractive to a more diverse set of people, so there is more diversity on the street; however, with a greater ability to sort to their ideal venue, men and women segregate themselves in different venues, leading them to be exposed to less diversity.

What is critical as we move from Panel (a) to (c) is not that the number of venues change, but rather that their offerings differ. It is the variety of venues that gives rise to sorting (i.e., passive segregation). There are two effects of increasing variety that are illustrated in Fig. 5. Greater variety attracts a more diverse set of people to the neighborhood (as seen in the streets of the figure, we move from 9 dark, 5 light in panel (a) to 7 dark, 7 light in panel (b)). Greater variety also leads to a decrease in diversity in the venues through greater sorting (as seen in the buildings of the figure). Hence, it is unclear whether larger cities actually offer more exposure to diversity in practice. A cosmopolitan city may be filled with many different types of people, but if there is a different restaurant for each different type of person, then it might be just as hard to experience diverse interactions.[21]

4.1. Identifying the causal effects of venue variety on venue and neighborhood diversity

Our data offer us an opportunity to test for the paradox of diversity directly. Our empirical strategy proceeds as follows: 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. 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.

In order to implement this strategy, we need measures of venue diversity, neighborhood diversity, and venue variety. We measure the amount of local diversity in venues in block b with the negative Theil Index, $D_b^V = -T_b$. The overall amount of neighborhood diversity in b should capture how representative the distribution of visitors in b are of the distribution of visitors in the whole city. Specifically, 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, we leverage the classification of venues in our data and define venue variety, V_b, as either the number of unique categories or subcategories of

[18] 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 Fig. 1 with alternative thresholds of 5, 10 and 15 percentage points and obtained qualitatively similar results.

[19] These conclusions do not change if we restrict our sample to the first or the second half of the year in the sample. We find a few exceptions to this rule, but they are infrequent and follow no salient systematic pattern.

[20] See, for example, the influential work of Jacobs (1961), Sandercock and Bridgman (1999) and Florida (2002). Fainstein (2005) provides a useful discussion of how the relationship between urban diversity and variety is viewed by urban planners and social scientists.

[21] The paradox of diversity sets up an interesting trade-off between serving the narrower needs of consumers and enriching society more broadly by increasing their exposure to diversity. Waldfogel (2009) introduces the concept of the "tyranny of the market" in which markets with large fixed costs can fail to serve individuals with niche preferences. While sufficiently "thick" markets do not suffer from the tyranny of the market, our analysis suggests that they will instead suffer from a lack of exposure to diversity at venues. On the other hand, "thin" markets that fall prey to the tyranny of the market are less vulnerable to a lack of exposure to diversity at venues.

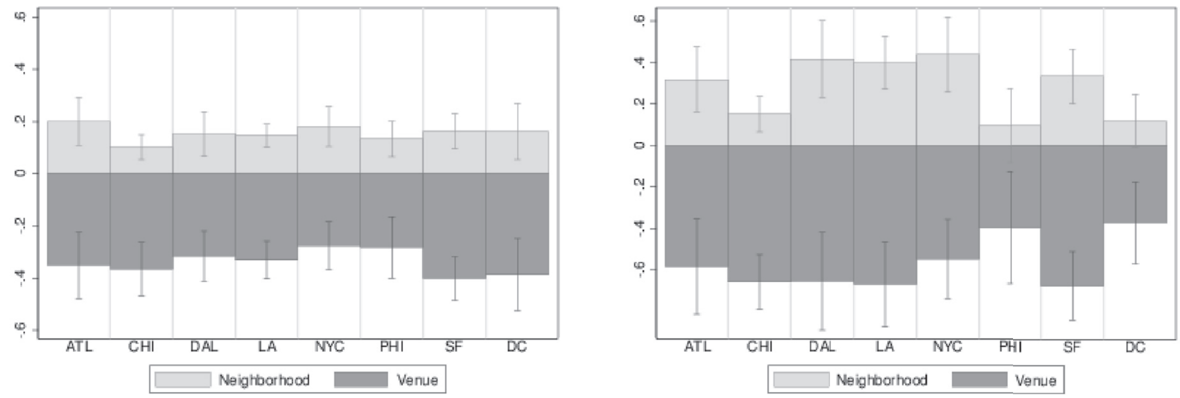

(a) V_b = Num. of Categories (b) V_b = Num. of Subcategories

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.

Fig. 6. $\hat{\beta}^V$ and $\hat{\beta}^N$ by city.

venues that are on offer in that block. We estimate the regression equations:

$$D_b^V = \beta^V V_b + \alpha_g^V + X_b \delta^V + R_b \lambda^V + \epsilon_b^V \quad (4)$$

$$D_b^N = \beta^N V_b + \alpha_g^N + X_b \delta^N + R_b \lambda^N + \epsilon_b^N \quad (5)$$

where α_g are fixed effects at the block group level for $b \in g$, and X_b represents a set of block level 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.[22] 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 Fig. 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 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.[23]

4.1.1. *Can we interpret $\hat{\beta}^V$ and $\hat{\beta}^N$ as causal?*

β^V and β^N are identified under the assumptions $\mathrm{Cov}\left(\epsilon_b^V, V_b \mid \alpha_g^V, X_b, R_b\right) = 0$ and $\mathrm{Cov}\left(\epsilon_b^N, V_b \mid \alpha_g^N, X_b, R_b\right) = 0$ 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. Thus, 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.[24]

Nevertheless, we provide four robustness checks that address these and other concerns. The results of these four robustness checks are shown in Fig. 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.[25] 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.

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

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

[23] 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.

[24] 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.

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

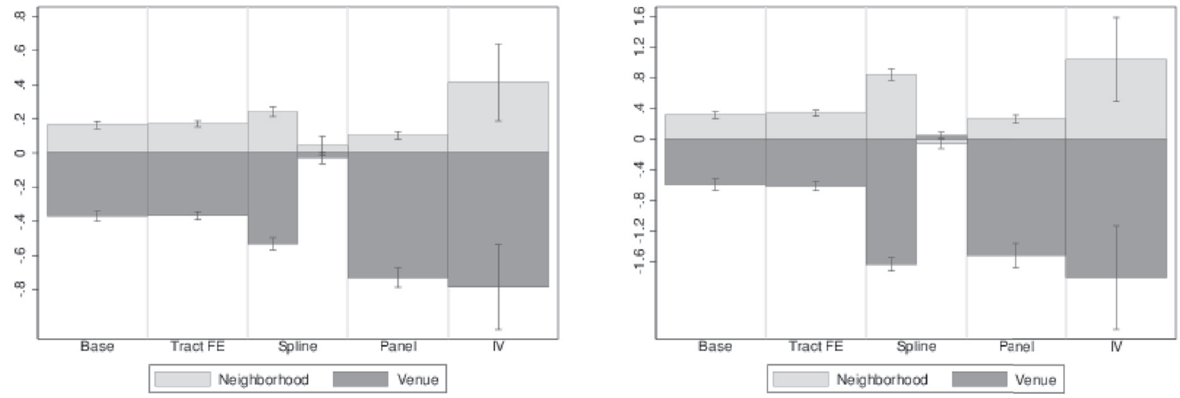

(a) V_b = Num. of Categories (b) V_b = Num. of Subcategories

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.

Fig. 7. $\widehat{\beta}^V$ and $\widehat{\beta}^N$: Alternative identification strategies.

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.[26]

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 $\widehat{\beta}^V$ and $\widehat{\beta}^N$ are causal estimates, then they will likely decline in magnitude as we compare nearby blocks with higher levels of venue variety. In contrast, if these estimates reflect confounding factors that are 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.[27]

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

$$D^V_{bt} = \beta^V V_{bt} + \alpha^V_b + \alpha^V_{ct} + X_{bt}\delta^V + \epsilon^V_{bt} \tag{6}$$

$$D^N_{bt} = \beta^N V_{bt} + \alpha^N_b + \alpha^N_{ct} + X_{bt}\delta^N + \epsilon^N_{bt}, \tag{7}$$

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^V_b and α^N_b) 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^V_{ct} and α^N_{ct} 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 levels of venue variety (but a similar number of venues, overall traffic and number of residents of each gender).[28]

Differences in zoning laws are found to generate differences in the variety of venues in nearby Census blocks. In Fig. 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

[26] 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.

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

[28] 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.

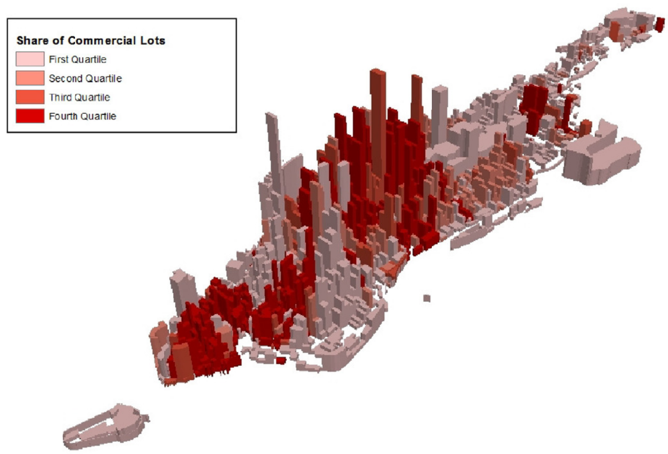

Notes: Each bar represents a Census block in Manhattan. The height of each bar corresponds to V_b, the number of unique venue subcategories in b. Darker bars represent blocks with a greater proportion of commercially zoned lots.

Fig. 8. Commercial zoning and venue variety (first stage).

Table 3
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.

error. As a result, our findings that $\beta^V < 0$ and $\beta^N > 0$ should be understood to be conservative.[29]

5. Discussion

5.1. Is segregation similar along other demographic dimensions?

Our data allows us to answer this question along exactly one additional dimension: age. For each venue, we observe the daily numbers of check-ins from users under 35 years of age and from users 35 years of age or older. We replicate our entire analysis substituting the proportion of youth for the proportion of females. Our results are broadly similar, which is not a trivial finding given that gender and age are largely uncorrelated. Although we find roughly half as much age segregation as gender segregation, it is highly localized as shown in Table 3: roughly a third to a half of all venue sorting by age in cities occurs within Census blocks. We also find that age segregation is primarily passive: people of different ages simply 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 (Fig. 9). A full reporting of all results from this replication is provided in the appendix.

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 passive segregation patterns along other dimensions such as race and income.

In a broader sense, individuals sort along multiple demographic dimensions simultaneously. For instance, younger women probably visit venues filled with other young women (rather than old women or young men). This suggests that sorting along multiple dimensions would further attenuate our measures of venue diversity.[30] Indeed, our findings complement a growing body of research that finds highly 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.

[29] In order to ensure that $\widehat{\beta}^V$ 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.

[30] Gender and age are mostly uncorrelated to other characteristics. This makes them ideal candidates to reach a conservative conclusion for our analysis that is inevitably incomplete.

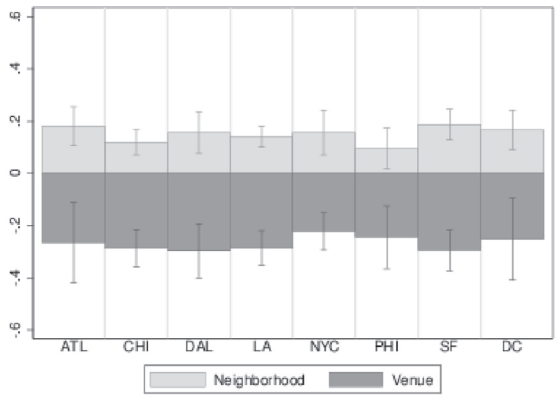

(a) V_b = Num. of Categories

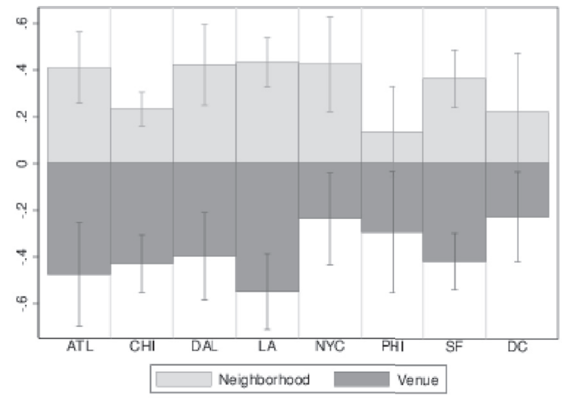

(b) V_b = Num. of Subcategories

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.

Fig. 9. $\hat{\beta}^V$ and $\hat{\beta}^N$ for age by city.

5.2. Do gender and age homophily affect other outcomes?

A large body of research has found that exposure to female peers affects 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 all of these social interactions don't occur at all of the different venues in our data, it is conceivable that repeated, informal exposure to peers in venues may accumulate over time, in turn affecting peoples' beliefs, preferences, social norms, and actions. Moreover, interactions between distantly related individuals may be quite valuable as suggested by the theory of the "strength of weak ties" (Granovetter, 1973): an interaction with a stranger may increase one's exposure to diversity by much more on the margin than an interaction with a close friend with a common social network.

Indeed, Akerlof and Kranton (2010) offer a theory of gender identity, developed in part through casual encounters, by which, "…society-wide changes are necessary to change gender norms…. The model predicts many implications of such changes. Women's participation in the labor force will increase. Occupational segregation will decrease…" (p. 90). Loury (2006) finds that female informal contacts have a lower impact on employment outcomes than male informal contacts, implying that gender segregated referral networks may contribute to the gender gap. And Bayer et al. (2008) show that interactions among neighborhood residents are gender segregated, and they seem to contribute to the gender gap in labor force participation. 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, though this is admittedly speculative. Identifying these various effects is a difficult proposition that carries heavy data demands and lies beyond the scope of this paper.

6. Conclusion

Peer groups shape our social environment. Homophily leads similar people to associate with one another, so the amount of segregation that is commonly observed in datasets might only represent the tip of the iceberg when it comes to the lack of diversity 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, our findings 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 may be a more complicated task for policymakers to achieve than previously thought.

Our analysis contributes to the ongoing debate on the ability of cities to 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 segregation 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.[31] On the other hand, Gentzkow and Shapiro (2011) find that news media (both online and offline) offer more exposure to diverse opinions than neighbors, coworkers 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 to such diversity, even if inadvertently. They just tend to be drawn to the same activities as similar people.

[31] "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).

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.

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.regsciurbeco.2019.05.004.

References

Adams, R.B., Ferreira, D., 2009. Women in the boardroom and their impact on governance and performance. J. Financ. Econ. 94 (2), 291–309.
Akerlof, G.A., Kranton, R.E., 2010. Identity Economics: How Identities Shape Our Work, Wages, and Well-Being. Princeton University Press, Princeton, NJ.
Andrews, J.A., Tildesley, E., Hops, H., Li, F., 2002. The influence of peers on young adult substance use. Health Psychol. 21 (4), 349.
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. J. Econ. Theory 2 (3), 244–263.
Bayer, P., Ross, S.L., Topa, G., 2008. Place of work and place of residence: informal hiring networks and labor market outcomes. J. Political Econ. 116 (6), 1150–1196.
Bourguignon, F., 1979. Decomposable income inequality measures. Econometrica 47 (4), 901–920.
Boustan, L., 2012. School desegregation and urban change: evidence from city boundaries. Am. Econ. J. Appl. Econ. 4 (1), 85–108.
Brown, D.A., Brown, D.L., Anastasopoulos, V., 2002. Women on Boards: Not Just the Right Thingbut the Bright Thing. Conference Board of Canada.
Card, D., Mas, A., Rothstein, J., 2008. Tipping and the dynamics of segregation. Q. J. Econ. 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.
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., 2018. How segregated is urban consumption? J. Political Econ.. 0 (ja), null https://doi.org/10.1086/701680.
Echenique, F., Fryer Jr., R., Kaufman, A., 2006. Is school segregation good or bad? Am. Econ. Rev. 265–269.
Fainstein, S.S., 2005. Cities and diversity: should we want it? can we plan for it? Urban Aff. Rev. 41 (1), 3–19.
Fernandez, R.M., Castilla, E.J., Moore, P., 2000. Social capital at work: networks and employment at a phone center. Am. J. Sociol. 1288–1356.
Fischer, C.S., Stockmayer, G., Stiles, J., Hout, M., 2004. Distinguishing the geographic levels and social dimensions of u.s. metropolitan segregation, 19602000. Demography. 41 (1), 37–59. https://doi.org/10.1353/dem.2004.0002.
Florida, R., 2002. The Rise of the Creative Class, vol. 9. Basic books, New York.
Gentzkow, M., Shapiro, J.M., 2011. Ideological segregation online and offline. Q. J. Econ. 126 (4), 1799–1839.
Glaeser, E.L., Kolko, J., Saiz, A., 2001. Consumer city. J. Econ. Geogr. 1 (1), 27–50.
Granovetter, M.S., 1973. The strength of weak ties. Am. J. Sociol. 78 (6), 1360–1380.
Hill, A.J., 2015. The girl next door: the effect of opposite gender friends on high school achievement. Am. Econ. J. Appl. Econ. 7, 147–177.
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.
Jacobs, J., 1961. The Death and Life of Great American Cities. Random House LLC.
Kossinets, G., Watts, D.J., 2009. Origins of homophily in an evolving social network1. Am. J. Sociol. 115 (2), 405–450.
Lavy, V., Schlosser, A., 2011. Mechanisms and impacts of gender peer effects at school. Am. Econ. J. Appl. Econ. 3 (2), 1–33.
Loury, L.D., 2006. Some contacts are more equal than others: informal networks, job tenure, and wages. J. Labor Econ. 24 (2), 299–318.
McPherson, M., Smith-Lovin, L., Cook, J.M., 2001. Birds of a feather: homophily in social networks. Annu. Rev. Sociol. 415–444.
Ruef, M., Aldrich, H.E., Carter, N.M., 2003. The structure of founding teams: homophily, strong ties, and isolation among us entrepreneurs. Am. Sociol. Rev. 195–222.
Sandercock, L., Bridgman, R., 1999. Towards cosmopolis: planning for multicultural cities. Can. J. Urban Res. 8 (1), 108.
Schelling, T.C., 1971. Dynamic models of segregation. J. Math. Sociol. 1, 143–186.
Shorrocks, A.F., 1980. The class of additively decomposable inequality measures. Econometrica: J. Econ. Soc. 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. Soc. Forces 56 (2), 576–597.
Waldfogel, J., 2009. The Tyranny of the Market: Why You Can't Always Get what You Want. Harvard University Press.
Weitzman, M.L., 1992. On diversity. Q. J. Econ. 107 (2), 363–405.