# Do Street Fairs Boost Local Businesses? A Quasi-Experimental Analysis Using Social Network Data

Ke Zhang, Konstantinos Pelechrinis
School of Information Sciences, University of Pittsburgh, Pittsburgh, PA, USA
* E-mail: Corresponding kpele@pitt.edu

## Abstract

Local businesses and retail stores are a crucial part of local economy. Local governments design policies for facilitating the growth of these businesses that can consequently have positive externalities on the local community. However, many times these policies have completely opposite from the expected results (e.g., free curb parking instead of helping businesses has been illustrated to actually hurt them due to the small turnover per spot). Hence, it is important to evaluate the outcome of such policies in order to provide educated decisions for the future. In the era of social and ubiquitous computing, mobile social media, such as Foursquare, form a platform that can help towards this goal. Data from these platforms capture semantic information of human mobility from which we can distill the potential economic activities taking place. In this paper we focus on street fairs (e.g., arts festivals) and evaluate their ability to boost economic activities in their vicinity. In particular, we collected data from Foursquare for the three month period between June 2015 and August 2015 from the city of Pittsburgh. During this period several street fairs took place. Using these events as our case study we analyzed the data utilizing propensity score matching and a quasi-experimental technique inspired by the difference-in-differences method. Our results indicate that street fairs provide positive externalities to nearby businesses. We further analyzed the spatial reach of this impact and we find that it can extend up to 0.6 miles from the epicenter of the event.

**Keywords:** Quasi-experimental design, Difference-in-differences, Social Media, Urban Informatics, Local Businesses

## 1   Introduction

A healthy local business sector is important for the prosperity of the surrounding community. City governments design policies and community organizations take actions that aim in boosting the growth of such businesses. This growth can have rippling positive externalities, such as, reducing local unemployment rates, keeping the local economy alive[1] and facilitating regional resilience to name just a few. These are even more important during periods of economic crises and recession, similar to the recent one in 2008 that US is just getting itself out of.

However, these efforts might not have the results expected. For example, many local governments during the "Small Business Saturday" (last Saturday of November) offer free curb parking. The rationale behind this policy is to give incentives to city dwellers (i.e., reduced trip cost to the business) to shop locally. However, the outcome is in many cases radically different. The underpricing of curb parking creates latent incentives for

---

[1]As per the New Economics Foundation "local purchases are twice as efficient in terms of keeping the local economy alive".

drivers to keep their cars parked for longer than normal periods of times. This leads to low turnover per parking spot and hence, ultimately to fewer number of customers in the local stores [1]. Therefore it is crucial to evaluate the efficiency of similar *interventions*. Knowing what boosts the local economy and what not, can allow the involved parties to make educated decisions for their future actions and ultimately lead to *urban intelligence* through data-driven decisions and policy making. In this study we are interested in a specific question and in particular, we are studying a research hypothesis related with the **impact of street fairs on neighboring local businesses.**

The golden standard for evaluating public policies is randomized experiments. However, in many cases designing and running the experiment is impossible from a practical point of view. Hence, quasi-experimental techniques [2] have been developed to analyze observational data in such a way that resembles a field experiment. To complicate things more with respect to our specific research hypothesis, evaluating the economic impact of street fairs requires access to the appropriate revenue data. While a city government office can obtain access to information such as sales tax revenue, local business advocates and citizens organizations will certainly face obstacles in obtaining such kind of data. This type of information is not part of the Open Data released by local governments and are accessible (if at all) in a very limited form through pay-per-request APIs (e.g., `http://zip-tax.com/pricing`). This lack of transparency can be compensated to a certain extend by utilizing information from social networks and social media. While similar types of data can potentially suffer from well-documented biases (e.g., demographic biases), they form an open platform that can be easily accessed and analyzed by citizens themselves to facilitate further investigation of issues, leading to a grassroots approach to urban governance.

In our case, given that we do not have actual revenue data for the businesses in the area of Pittsburgh as aforementioned, we collect Foursquare *check-ins* from the city of Pittsburgh over a three-month period (June-August 2015) and evaluate the effect of summer street fairs on local economy. The check-in information can serve as a proxy - even though not perfect - for the revenue $\rho$ generated [3]. We would like to emphasize here that, our study aims in evaluating the impact of street fairs on the brick-and-mortar stores that are adjacent to the event location and not that on the participating entities – which is expected to be positive in order for them to participate.

In order to analyze our data we rely on two quasi-experimental techniques. First, an increase in the check-ins for the venues near the street fair does not necessarily mean that this was due to the event. One or more control areas need to be used for comparison. However, our data are not generated through a randomized experiment but they are purely observational. For our analysis, this essentially means that we cannot assume that the area hosting a street fair event is chosen at random. Consequently, we cannot assume that the areas that do not host street fairs exhibit the same characteristics with respect to unobserved confounding features and hence, we cannot compare the revenue in the treated area with any untreated area. For overcoming this problem, we rely on quasi-experimental design techniques that identify appropriate control areas. In particular, we rely on propensity score matching [4], adopted in our setting by utilizing expert domain knowledge, in order to pick a set of *matched* areas $\mathcal{A}_m$ with the treated area $\alpha$ that will serve as our control subjects. Second, once the matched areas for comparison are chosen, we adopt the difference-in-differences method [5] in our setting in order to quantify the impact of the street fairs on local businesses. In a nutshell, the difference-in-differences is a regression model that examines the average change of the treatment group once the treatment has been applied and compares it with the control group. The implicit assumption is that this difference would be zero if the treatment had not been applied. We elaborate further on these two methods in the following section.

The main contributions of our work can be summarized as follows:

- We provide quantifiable evidence that support the positive impact of street fairs on local businesses.

- We show how social media data - despite their potential biases - can be useful to public policy makers and local governments since they are transparent, accessible and are able to provide good evidence when analyzed properly.

**Scope of our work:** While in the current study we are focusing on the effect of street fairs on local businesses the method can be applied in a variety of scenarios that include an external event/stimulant. For example, one can use our framework to quantify the effect of short-term road closures and/or constructions on the local economy. This is especially important during the bidding phase of a construction project since these effects should be included in the calculation of liquidated damages [6]. However, they are not currently included since there is not a framework to estimate this effect.

**Roadmap:** In the following section we present our method. We then describe our experimental setup and results, while we further discuss the limitations of our study. Finally, we discuss relevant to our work studies and conclude our study.

## 2 Analytical Methods

Let us denote the total volume of revenue within area $\alpha$ at day $t$ with $\rho_{t,\alpha}$. Furthermore, $\mathcal{T}_\alpha$ is the set of days that a street fair took place within area $\alpha$. The trending of $\rho_{t,\alpha}$ by itself cannot reveal anything with respect to the contribution of the street fair at the revenue generated in area $\alpha$. Hence, in order to account for various confounding factors and other externalities we will need to get a "baseline" for comparison. When experimental design and implementation is possible this happens with random assignment of the treatment (in our case the street fair) to the experimental subjects. However, in our case this is not possible and hence, we rely on matching techniques and more specifically we use propensity score matching. Matching techniques provide us with the ability to analyze observational data in a way that mimics some of the particular characteristics of a randomized trial. In particular, we choose a matched, with area $\alpha$, neighborhood, say, $\alpha_m$, to analyze and compare the corresponding revenues generated.

Our analysis is inspired by the difference-in-differences method [5]. In brief, we compare the daily revenue differences between the area with the street fair and the corresponding matched area(s) both during the period of the street fairs as well as during the period without any street fair. The comparison with the matched area(s) - that are exposed to the same externalities - accounts for various confounding factors that can affect revenues, and hence, any observed difference can be attributed to the treatment, i.e., the street fairs in our case. In what follows, we describe in detail the building blocks of our analysis, i.e., propensity score matching and difference-in-differences.

### 2.1 Propensity score matching

Propensity score matching can be used to reduce (or even eliminate) the effect of confounding variables on the analysis of observational data. To reiterate propensity score matching allows an analysis in a way that mimics a randomized trial. In our own context, the **_treatment_** of interest is whether or not there is a street fair in neighborhood $i$. The propensity score of each (untreated) instance (i.e., every untreated

neighborhood) represents the probability of this instance to be treated, conditional on a set of confounding variables. In a real randomized experiment, the instances are randomly assigned to the treatment and control groups. This ensures (given sufficiently large number of instances) that on average the two groups will only differ with respect to the reception of the treatment. In the case of observational data, the treatment is not randomly assigned but usually the "treated" instances are chosen due to some specific characteristics (i.e., the confounding factors). Therefore, in order to identify an appropriate control group we need to calculate the probability of the untreated instances obtaining the treatment.

In order to calculate the propensity scores, i.e., the conditional probabilities of the instances receiving the treatment, we employ a logistic regression model similar to [7]. In particular, given a feature vector $\boldsymbol{Z}$ that is formed by a set of neighborhood characteristics (i.e., the confounding factors) we estimate the following conditional probability:

$$\Pr(b_i = 1|\boldsymbol{Z}_i) = \frac{\exp(\boldsymbol{a}_i^T \cdot \boldsymbol{Z}_i)}{1 + \exp(\boldsymbol{a}_i^T \cdot \boldsymbol{Z}_i)} \tag{1}$$

where $b_i$ is a binary indicator variable, which takes the value 1 if area $i$ is treated and 0 otherwise. In our case, $\boldsymbol{Z}_i$ includes **three types of features** for every type of establishment $T$ that exists in neighborhood $i$ that captures **(a)** the fraction of type $T$ venues in $i$, as well as, **(b)** the fraction of the revenue (check-ins in our case) within $\alpha$ that was generated by venues of type $T$. Finally, we for every business venue type, we use **(c)** the "stickiness" of the users in this type as an additional feature. The "stickiness" is defined as the ratio between the total number of check-ins in the corresponding category over the number of unique users that generated these check-ins.

After training the aforementioned logistic regression model, we estimate the probability from Equation (1) for all neighborhood instances $i \in \mathcal{N}$ (both treated and untreated), where $\mathcal{N}$ is the set of areas/neighborhoods. Then we match the treated neighborhood $\alpha$, with:

$$\alpha_m = \min_{i \in \mathcal{N} \setminus \{\alpha\}} |\Pr(b_i = 1|\boldsymbol{Z}_i) - \Pr(b_\alpha = 1|\boldsymbol{Z}_\alpha)| \tag{2}$$

Essentially, this means that area $\alpha_m$ is the one that has the closest probability of hosting a street fair to that of area $\alpha$, under the assumption that the only features that affect the decision are the ones captured by the observable confounding variable vector $\boldsymbol{Z}$.

In many scenarios (such as in our case study) we might only have one treated area $\alpha$, i.e., only one area has hosted a street fair. In this case, evaluating Equation (2) is *trivial*, since, the minimum is observed for the area $i$ for which the vector distance $d(\boldsymbol{Z}_i, \boldsymbol{Z}_\alpha)$ is minimized. Simply put, the matched area $\alpha_m$ is the one whose feature vector $\boldsymbol{Z}_{\alpha_m}$ is closer to that of the treated area $\boldsymbol{Z}_\alpha$. We would like to emphasize here that, there might be other, unobserved, factors that lead to the choice of an area for a street fair. This is a limitation of the quasi-experimental techniques in general and propensity score matching can only account for observable confounders $\boldsymbol{Z}$.

One way we propose to use in order to alleviate some of the potential problems associated with the aforementioned limitation is to initialize the matching process with expert knowledge. In particular, the matched area $\alpha_m$ can be chosen using expert knowledge (e.g., urban planners in our case). The benefit of this approach is that the domain expert is - implicitly or explicitly - considering various (potentially unobserved) confounders simultaneously. We can then use the expert matching as a "seed" for matching more than one neighborhoods to $\alpha$ using the propensity scores.

In particular, with $\pi_{m,e}$ being the propensity score of the (domain expert) matched area $\alpha_{m,e}$, we can pick the following set of matched areas:
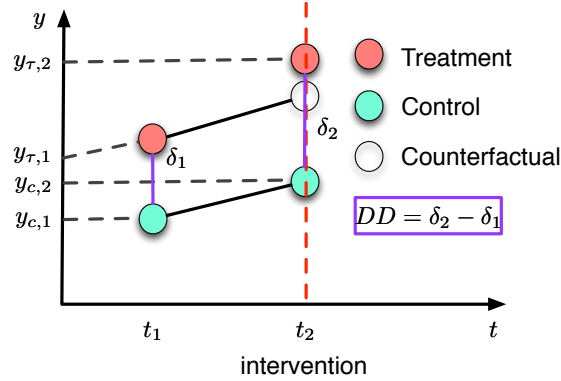
**Figure 1.** The difference in differences method.
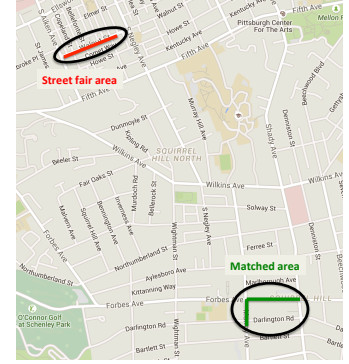


**Figure 2.** Treated and domain expert matched neighborhood.

$$\mathcal{A}_m = \{\alpha_{m_j} : \ |\pi_{m_j} - \pi_\alpha| < |\pi_{m,e} - \pi_\alpha| + \epsilon\} \tag{3}$$

Essentially, as per Equation (3), the set $\mathcal{A}_m$ includes neighborhoods that have propensity scores that are closer to the score of the treated area (within a tolerance factor $\epsilon$) as compared to the expert matched area. Once set $\mathcal{A}_m$ is obtained we can analyze the corresponding revenues generated using the difference-in-differences method described in what follows.

## 2.2 Difference-in-Differences

The difference in differences (DD) method [5] is a quasi-experimental technique that aims in identifying the effect of an intervention using observational data. DD requires observations obtained in different points in time, e.g., $t_1$ and $t_2$ ($t_1 < t_2$), for both the control (e.g., $y_{m,1}$ and $y_{m,2}$) and the treatment (e.g., $y_{\tau,1}$ and $y_{\tau,2}$) groups. The treatment group is exposed to the intervention only during $t_2$. The difference between $y_{\tau,2}$ and $y_{m,2}$ does not only include the effect of the intervention but it also includes other "intrinsic" differences between the two groups. The latter can be captured by their difference during time $t_1$, i.e., $y_{\tau,1} - y_{m,1}$, where the treatment group has not been exposed to the intervention. The DD estimate is then:

$$\delta_{\tau,m} = (y_{\tau,2} - y_{m,2}) - (y_{\tau,1} - y_{m,1}) \tag{4}$$

If $\delta_{\tau,m} > 0$ ($\delta_{\tau,m} < 0$), then the treatment has a positive (negative) impact on $y$, while if $\delta_{\tau,m} = 0$ there is not any impact from the intervention. Equation (4) captures the impact of the intervention assuming that both the treatment and control follow a **parallel trend**. In particular, in order for the conclusions drawn from a difference-in-differences analysis to be reliable, the parallel trend assumption needs to hold. This assumption essentially states that the average change in the control group represents the counterfactual change expected in the treatment group if there was no treatment. Simply put, if there was not any treatment applied, we would have: $(y_{\tau,2} - y_{m,2}) = (y_{\tau,1} - y_{m,1})$, that is, the two groups would have a stable difference. This assumption is crucial for the conclusions from a difference-in-differences analysis to hold and is many times overlooked when the method is applied.

The exactly same estimate for the DD can be formally derived through a linear regression that models the dependent variable $y$. In particular, we have the following

model:

$$y_{ilt} = \gamma_0 + \gamma_1 \cdot \alpha_l + \gamma_2 \cdot \beta_t + \delta \cdot D_{lt} + \epsilon_{ilt} \tag{5}$$

where $y_{ilt}$ is the dependent variable for instance $i$ (at time $t$ and location $l$), $\alpha_l$ and $\beta_t$ are binary variables that capture the fixed effects of location and time respectively, $D_{lt}$ is a dummy variable that represents the treatment status (i.e., $D_{lt} = \alpha_l \cdot \beta_t$) and $\epsilon_{ilt}$ is the associated error term. The coefficient $\delta$ captures the effect of the intervention on the dependent variable $y$. It is then straightforward to show that the DD estimate $\hat{\delta}$ is exactly Equation (4). In particular, if $\overline{y}_{lt}$ is the sample mean of $y_{ilt}$ and $\overline{\epsilon}_{lt}$ is the sample mean of $\epsilon_{ilt}$, and using Equation (5) we have:

$$(\overline{y}_{11} - \overline{y}_{01}) - (\overline{y}_{10} - \overline{y}_{00}) = \quad \delta(D_{11} - D_{01}) - \delta(D_{10} - D_{00}) \quad + \overline{\epsilon}_{11} - \overline{\epsilon}_{01} + \overline{\epsilon}_{00} - \overline{\epsilon}_{10}$$

Taking expectations and considering the i.i.d. assumptions for the errors for the ordinary least squares we further get:

$$\mathrm{E}[(\overline{y}_{11} - \overline{y}_{01}) - (\overline{y}_{10} - \overline{y}_{00})] = \delta(D_{11} - D_{01}) - \delta(D_{10} - D_{00}) \tag{6}$$

Given that the dummy variable $D$ is equal to 1 only when $l = 1$ and $t = 1$ (i.e., for the treatment group after the intervention), we finally get for the DD estimator:

$$\hat{\delta} = (\overline{y}_{11} - \overline{y}_{01}) - (\overline{y}_{10} - \overline{y}_{00}) \tag{7}$$

which is essentially the same as Equation (4). Therefore, one can estimate the DD using either of the Equations (4) or (5). Figure 1 further visualizes the estimation process. The control and treatment subjects in our setting are urban neighborhoods. Treated subjects includes neighborhoods that host street fairs.

## 2.3 Hypothesis development

Having introduced our basic methodology we are ready to formally state the research hypotheses that are the focus of our study. In particular, we will examine the following two hypotheses.

**Hypothesis 1** [`Street fairs impact on local businesses`]: *Street fair events lead to an increase in customer visitations for nearby business venues.*

**Hypothesis 2** [`Spatial impact of street fairs`]: *The impact of street fairs on the customer visitations is geographically contained in a very small area.*

In order to support or reject Hypotheses 1 and 2 we will rely on data we collected from Foursquare described in the next section, utilizing the difference-in-differences method described in Section 2.2. We will further examine contextual dependencies, i.e., whether specific types of business venues benefit more than others.

# 3 Experimental Setup and Results

In this section we will present the dataset we collected, as well as, the setup for our analysis. We will then present our results and finally, we will discuss the implications and the limitations of our analysis.

## 3.1 Dataset

For the purposes of our study we collected time-series data using Foursquare's venue public API. We queried daily all Foursquare venues in Pittsburgh for the three-month period between **06/01/2015 - 08/30/2015**. This period includes six street fairs/events[2] that took place at a specific neighborhood in the city of Pittsburgh (see the street marked with red in Figure 2).

Our time-series data include information with respect to the number of check-ins $c_v[t]$ that have been generated in venue $v$ during day $t$. To reiterate, given the fact that we do not have actual revenue data for the businesses in Pittsburgh we rely on the check-in information as a proxy for the corresponding revenue of venue $v$, $\rho_v[t]$. This information will allow us to build the aggregate volume daily check-ins $c_\alpha$ within area $\alpha$, i.e., $c_\alpha[t] = \sum_{v \in \alpha} c_v[t]$. Every area is defined as a circle of radius $r$ centered at the centroid of the neighborhood under consideration. In our experiments, we examine various values for $r$ in order to explore the spatial distribution of the impact.

We have also collected meta-data information. In particular, Foursquare associates each venue $v$ with a type/category $T$ (e.g., restaurant, school etc.). This classification is hierarchical and at the top level of the hierarchy there were 9 categories at the time of data collection. In order to obtain the feature vector $\mathbf{Z}$, we use the top-level categories and hence $\mathbf{Z}$ includes 21 features (2 for each category and 3 for the stickiness of each type of business venue). Our final dataset includes 27,263 venues in the city of Pittsburgh, where 21.53% (5,869) are business venues (i.e., *Nightlife Spots*, *Food* and *Shops & Services*). There are in total 32,501 check-ins in our dataset, among which 44.46% were generated in business venues.

## 3.2 Experimental Setup

In our study we consider a single area $\alpha$ that has hosted street fairs during our data collection period. This area is a small business center, with a number of restaurants, cafes, retail stores (e.g., clothing stores, galleries etc.) and services (e.g., bank branches). The *treated* area is also accessible through public transportation, Pittsburgh's shared bike system as well as through private vehicle with parking facilities nearby. We (initially) perform the matching process based on the *expertise*[3] of local urban planners. Based on their recommendations we choose another small business area, with a similar urban form and accessibility patterns not very far from the treated area (approximately 2 miles away - green area in Figure 2). We have further used Equation (3) to build a set of matched areas. More specifically, we first pick 2,000 random points in the city of Pittsburgh and create a neighborhood of radius 0.3 miles around this point. We further eliminate areas with less than 60 venues. We consequently obtain the matched area set $\mathcal{A}_m$ using Equation (3) with $\epsilon = 0$ and we filter out overlapping matched neighborhoods, in order to remove possible dependencies in our datasets originating from the overlapping regions. In particular, when $k$ matched areas overlap we only keep the final matched set the area with a propensity score matching closest to the treated area. We would like to emphasize here that we have examined different values for the radius of the control neighborhood area selection and the tolerance factor $\epsilon$ and the results obtained were very similar.

## 3.3 Results

The metric of interest for our analysis is the mean number of daily check-ins in area $\alpha$, denoted with $y_\alpha$. For every area $\alpha$ we compute the average number of daily check-ins

---

[2] http://thinkshadyside.com/events/
[3] We have consulted with urban planners familiar with the city of Pittsburgh.
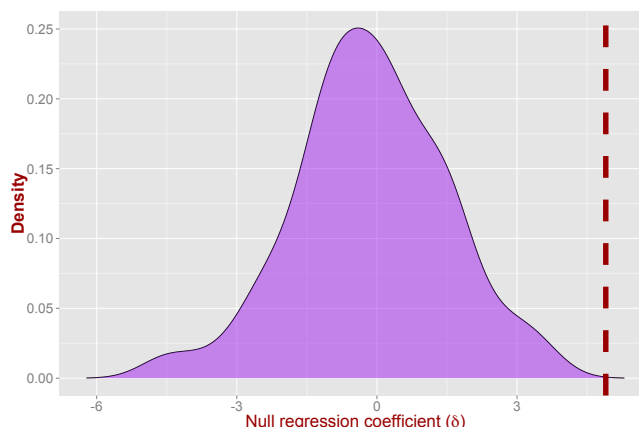
**Figure 3.** The null difference-in-differences coefficient is practically equal to 0, hence, allowing us to apply the model with high confidence.

during the treatment period, $y_{\alpha,\mathcal{T}_\alpha}$, as well as, during the days with no street fair, $y_{\alpha,\mathcal{T}_\alpha^c}$, where $\mathcal{T}_\alpha^c$, represents the complement of $\mathcal{T}_\alpha$, i.e., the set of days in our dataset where no street fair took place in $\alpha$. With this setting the difference-in-differences coefficient is equal to **4.95 (p-value ¡ 0.001)**. Simply put, there are 5 more check-ins every day with a fair in area $\alpha$ on average. This corresponds to an almost 100% increase in the check-ins in the area, since the average daily check-ins for the days with no event is 5.3.

As mentioned in Section 2.2 one of the crucial assumptions for the difference-in-differences to provide robust results is the parallel trend assumption. Typically the way that has been followed in the literature for verifying this assumption is to calculate the difference-in-differences coefficient for periods that the treatment has not been applied [8,9]. Hence, for the days that in reality no street fair occurred we randomly assign pseudo-treatments in order to calculate a null coefficient $\delta$. Figure 3 depicts the distribution of the corresponding coefficients obtained from 100 randomizations. As we can see the mass of the distribution is concentrated around $\delta = 0$, while the 95% confidence interval is $[-0.42, 0.37]$. Hence, we cannot reject the hypothesis that the null coefficient $\delta$ is actually 0, hence, verifying the parallel trend assumption needed for the difference-in-differences method.

We also want to examine the spatial extent of this impact, i.e., how the impact decays with space. For this, we compute the difference-in-differences coefficient for zones of different radius around the treated area making sure that there is not any overlap with control areas. In particular, we examine zones of $[0, 0.1]$, $[0.1, 0.3]$, $[0.3, 0.6]$ miles. Our results are depicted in Figure 4 where as we can see there is a clear decreasing trend of the impact. In fact, the coefficient for the range $[0.1, 0.3]$ miles is much smaller, and equal to 0.89 (p-value ¡ 0.1), while going further away from the area of the event (i.e., $[0.3, 0.6]$ miles) the effect is practically eliminated ($\delta_{[0.3,0.6]} = 0.33$, p-value = 0.61). These results indicate - as one might have expected - that the impact of a street fair event is highly localized within a very small area around the epicenter of the event.

We further examine the impact of each event individually, i.e., we consider a single day treatment. Table 1 presents our results. As we can see every event contributes to the overall local business sector a positive increase to the check-ins, which can further be translated to increase foot traffic and revenue. The only exception is the Vintage GP Car show. Compared to the other events, this attracts a very specific part of the population - i.e., car-lovers - and this might have affected its overall impact.

Our analysis until now has considered all of the business venues together regardless of their type. This essentially captures the aggregate impact of the street fair in the
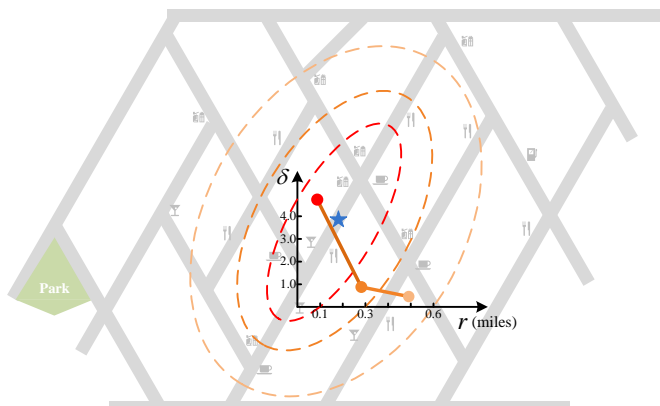
**Figure 4.** The impact of street fairs on local businesses rapidly decays with the spatial distance from the event.

| Event | Difference-in-differences coefficient $\delta$ |
|---|---|
| Jam On Walnut 1 | 9.7*** |
| Vintage GP Car Show | -2.01*** |
| Jam on Walnut 2 | 5.45*** |
| Jam on Walnut 3 | 6.64*** |
| Arts Festival on Walnut 1 | 4.45*** |
| Arts Festival on Walnut 2 | 5.53*** |

**Table 1.** All events - except the Vintage GP Car Show - exhibit a statistically significant and positive coefficient $\delta$. The reason why the car show does not impact the nearby businesses could potentially be attributed to the fact that compared to other events, it attracts a specific part of the population only. Significance codes: 0 '***' 0.01 '**' 0.05 '*' 0.1 '.' 1 ' '.

neighborhood. However, we would like to decompose this effect in order to understand better what type of establishments benefit from the fairs. In particular, we compute the difference-in-differences regression coefficient for the three different types of business venues our dataset contains. Figure 5 depicts our results, where the 95% confidence interval of the estimated coefficients is also presented. As we can see shopping venues are the ones that benefit the most from the street fairs, while nightlife and food establishment exhibit a much (but significant and positive) lower coefficient $\delta$. However, one crucial point here is that the coefficient provides the cumulative - additional to the counterfactual - check-ins recorded in all venues of the specific type. Hence, if a specific venue type is overrepresented in the area the estimated DD coefficient might be *inflated*[4]. In order to avoid similar issues, we can normalize the obtained coefficients from the regression model by the number of venues for every establishment type. In particular, the number of shop, nightlife and food venues in the treated area are 60, 13 and 25 respectively. Therefore, the normalized coefficients for the shop and nightlife are practically equal (0.066 and 0.061 respectively). However, the food venues still have a much smaller normalized coefficient, that is, 0.014.

Overall, we can say that our results support the two research hypotheses put forth in Section 2.3. In particular, street fairs have a positive impact on nearby businesses as captured by the check-ins on Foursquare and the difference-in-differences method. Furthermore, this impact is highly concentrated in the areas around the street fair (i.e.,

---

[4]Note here that, this is not an issue when we applied the difference-in-differences at the level of a neighborhood. In that case, we were interested in the total additional check-ins in the neighborhood as compared to the counterfactual. Hence, if a control area had a different number of venues this would not impact the results.
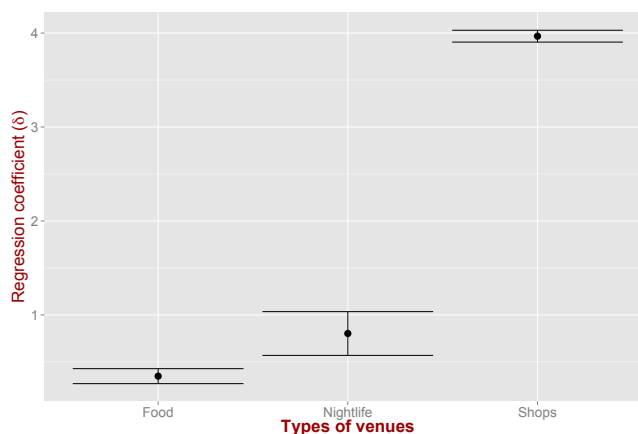
**Figure 5.** The shopping businesses appear to have the largest benefit from the street fairs among the local establishments around the area.

0.1, 0.2 miles) and drops extremely fast as we move further away.

## 3.4 Discussion and Limitations

One of the main critics that studies relying on social media get is that of the potential demographic biases that the data include. This is certainly true and is one of our study's limitation as well. Nevertheless, location-based social media is a very good, and accessible, proxy for the economic activities in urban areas. Certainly there will be noise in the obtained signal, but this information is valuable for providing supporting (or not) evidence in a variety of research hypotheses similar to ours. For example, similar datasets have been used to study urban gentrification, deprivation, emotions in a city [10–12] etc.

In our difference-in-differences regression model we included fixed time and location effects. One might argue that we should also control for the day of the week. However, this is not necessary since the null regression model essentially shows us that the different days of the week will exhibit the same "trending" on average (of course the absolute values of the check-ins will be different). To verify this we run the regression model by adding an independent variable that captures the day of the week. Our results for the various zones around the treated neighborhood are presented in Table 2.

| Radius $r$ | Difference-in-differences coefficient $\delta$ |
|---|---|
| [0, 0.1] | [3.71, 4.1] |
| [0.1, 0.3] | [0.17, 0.85] |
| [0.3, 0.6] | [0.51, 1.61] |

**Table 2.** Even when controlling for the day of the week, the impact of the street fair remains. The 95% confidence intervals of the difference-in-differences coefficient for the ranges [0.1, 0.3] and [0.3, 0.6] overlap; hence, we cannot support with confidence the presence of a larger impact in the further zone. Significance codes: 0 '***' 0.01 '**' 0.05 '*' 0.1 ' ' 1

As we can see even when controlling for the day of the week the impact is strong and significant. In fact, when controlling for the day of the week the impact appears to be significant even for distances beyond the 0.1 miles. Nevertheless, the impact itself is weak (i.e., the coefficient is small). Furthermore, even though it appears that the further zone has a stronger effect, the 95% confidence intervals for the two coefficients overlap, and hence, we cannot confidently support the presence of a trend.

# 4    Related Work

In this section we briefly discuss related methodological literature as well as literature relevant to the specific application domain.

**Quasi-experimental methodologies:**   The gold standard for evaluating the impact of a policy is a field experiment. However, when it comes to public policy many times this is not possible for a variety of reasons. In this case we need to rely on quasi-experimental techniques [2] in order to quantify the potential impact. Quasi-experimental designs allows to control the assignment to the treatment condition, but using some criterion different than random assignment as in field experiments.

There are various techniques that can be used depending on the type of observational data one has. For example, the difference-in-differences method [5] compares the average change over time in the outcome variable for the treatment group to the average change over time for the control group. One of the major problems when applying this method is the parallel trend assumption, that is, that the two groups exhibit the same temporal trend on their averages without the treatment. Regression discontinuity [13] is another technique that can be used to quantify the effects of treatments that are assigned by a threshold. The key idea is that observations lying very closely on either side of the threshold while differing in the reception of the treatment, they are *equal* for all practical purposes. Hence, their treatment assignment mimics that of a randomized control trial. It should be clear that not all quasi-experimental designs are applicable in all scenarios (for example regression discontinuity cannot be applied in our setting), while there can be settings were no method is applicable. A nice survey of various quasi-experimental techniques can be found in [14].

**Local businesses and urban economy:**   Small shops and businesses are the backbone of local economy and quantifying the effect of external events and policies on their prosperity is of utmost importance. Given the absence of large scale data, most of the existing studies have been based on survey data. For instance, a survey research conducted by Lee *et al.* [15] during the 2002 World Cup identified that the event-related tourists yielded much higher expenditure as compared to *regular* tourists, indicating that such mega-events could have a positive economic impact for local businesses. As another example, a report from a Toronto-based think tank has identified the positive impact that bike lanes have on the revenue of local businesses despite the fact that business owners systematically underestimate it [16]. In a similar direction, based on merchant and pedestrian surveys in Toronto's Annex Neighborhood, the "Clean Air Partnership" [17] recommended reallocating a curb parking lane to bike lanes, since this is likely to increase commercial activity. A recent study further showed that the installation of shared bike system can lead to an increase of the housing property values [9]. Moreover, in a briefing paper DeShazo *et al.* [18] using a survey conducted over a small sample of businesses quantified the effect of CicLAvia on local businesses. CicLAvia[5] is a car-free event that happens once every year in various areas in Los Angeles. Furthermore, anecdotal hard evidence from Seattle [19] show that increasing the price of curb parking can be beneficial to restaurants and local businesses mainly due to the increased turnover of each parking spot [1].

During the last years, and driven by the proliferation and availability of geo-tagged social media data, there has been a surge of studies on business analytics. For instance, Qu and Zhang [20] proposed a framework that extends traditional trade area analysis and incorporates location data of mobile users. Their framework can answer crucial questions in retail management such as "where are the customers of a business coming from?". As another example, Karamshuk *et al.* [21] proposed a machine learning framework to predict the optimal placement for retail stores, where they extracted two

---

[5]http://www.ciclavia.org

types of features from a Foursquare check-in dataset. Furthermore, these platforms can serve as mobile "yellow pages" with business reviews that can influence customer choices and business revenue. For example, Luca [22] has identified a causal impact of Yelp ratings on restaurant demand using the regression discontinuity framework. Closer to our study, Georgiev *et al.* [23] using data collected from Foursquare study the impact of the 2012 Olympic Games on the businesses in London, while Zhang *et al.* [3] quantify the effectiveness of special deals offered through location-based services as an affordable advertisement for local businesses.

To the best of our knowledge no one has examined the impact of street fairs on the adjacent businesses, even though local authorities expect this policy to have a positive outcome for businesses[6]. Studies that examine the economic effects of special events/festivals exist (e.g., [24]) but their focus is slightly different, focusing on the participating entities/kiosks themselves. On the contrary, our study is focused on the "network" effects a street fair can have for the nearby businesses.

## 5    Conclusions and Future Work

In this study we have used social media data and quasi-experimental techniques to evaluate the effect of street fairs on the local business sector. In particular, we have adopted quasi-experimental techniques, i.e., difference-in-differences, and synthesized them with domain expert knowledge. We consequently applied our method on street fairs and outdoors arts festivals that took place on a specific neighborhood in the city of Pittsburgh as a case study. Our results indicate that similar street fairs can boost local businesses and stimulate and contribute to a healthy local economy. Similar approaches can be used to evaluate the impact of different interventions (e.g., installation of new transportation modes, alterations on the street network etc.).

Of course, our specific case study exhibits limitations with respect to the available data as we elaborated earlier. In particular, while check-in information is intuitively a good proxy for the underlying revenues, demographic biases can provide us with a skewed view of the exact magnitude of the impact. Nevertheless, similar analysis can provide advocate citizens' organizations with a case for further scrutiny of any public policy in place. Social media data are "readily" available and accessible (at least most of the times) and can provide the basis for grassroots innovation in the space of policy evaluation. In the future we plan in examining other potential sources (e.g., sales tax data) and analyze information from other cities as well, in order to obtain a cross-city comparison with respect to street fairs and their impact on local economy. Furthermore, even though we have verified the parallel trend assumption, the increase in the check-ins (revenues) in the treated area might be partly attributed to a decrease in the rest of the areas. This interaction between neighborhoods is extremely interesting and can potentially be captured and analyzed through a network between the urban areas. Finally, the long-term effect of these events is also important. In particular, even though the street fair can potentially increase the revenues during its lifetime, does it have the ability to create new clientele for the area? We will further explore these points in our future work.

---

[6]E.g., `http://tinyurl.com/zdved39`

# References

1. Shoup D (2011) The High Cost of Free Parking. American Planning Association.

2. Shadish WR, Cook TD, Campbell DT (2001) Experimental and Quasi-Experimental Designs for Generalized Causal Inference. Cengage Learning.

3. Zhang K, Pelechrinis K, Lappas T (2015) Analyzing and modeling special offer campaigns in location-based social networks. In: AAAI ICWSM.

4. Rosenbaum PR, Rubin DB (1983) The central role of the propensity score in observational studies for causal effects. Biometrika 70: 41–55.

5. Ashenfelter O, Card D (1985) Using the longitudinal structure of earnings to estimate the effect of training programs. The Review of Economics and Statistics 67: 648-60.

6. Goetz CJ, Scott RE (1977) Liquidated damages, penalties and the just compensation principle: Some notes on an enforcement model and a theory of efficient breach. Columbia Law Review 77: pp. 554-594.

7. Aral S, Muchnik L, Sundararajan A (2009) Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. Proceedings of the National Academy of Sciences 106: 21544–21549.

8. Mora R, Reggio I (2012) Treatment effect identification using alternative parallel assumptions .

9. Pelechrinis K, Kokkodis M, Lappas T (2015) On the value of shared bike systems in urban environments: Evidence from the real estate market. Available at SSRN .

10. Hristova D, Williams M, Musolesi M, Panzarasa P, Mascolo C (2016) Measuring urban social diversity using interconnected geo-social networks. In: ACM WWW.

11. Venerandi A, Quattrone G, Capra L, Quercia D, ez Trumper DS (2015) Measuring urban deprivation from user generated content. In: CSCW.

12. Gallegos L, Lerman K, Huang A, Garcia D (2016) Geography of emotion: Where in a city are people happier? In: WWW.

13. Imbens G, Lemieux T (2007) Regression discontinuity designs: A guide to practice. Working Paper 13039, National Bureau of Economic Research. doi:10.3386/w13039.

14. Harris AD, McGregor JC, Perencevich EN, Furuno JP, Zhu J, et al. (2006) The use and interpretation of quasi-experimental studies in medical informatics. Journal of the American Medical Informatics Association : JAMIA 13: 16–23.

15. Lee CK, Taylor T (2005) Critical reflections on the economic impact assessment of a mega-event: the case of 2002 fifa world cup. Tourism management 26: 595–603.

16. Arancibia D (2012) Cycling economies: Economic impact of bike lanes. Report Toronto Cycling; Think and Do Tank .

17. CleanAir-Partnership (2009) Bike lanes, on-street parking and business. Report .

18. DeShazo J, Callahan C, Brozen M, Heimsath B (2013) Economic impacts of ciclavia: Study finds gains to local businesses. In: Briefing Paper - UCLA Luscin School of Public Fairs.

19. de Place E (2012). Are parking meters boosting business?: http://daily.sightline.org/2012/03/28/is-metered-parking-boosting-business/. URL `http://daily.sightline.org/2012/03/28/is-metered-parking-boosting-business/`.

20. Qu Y, Zhang J (2013) Trade area analysis using user generated mobile location data. In: ACM WWW.

21. Karamshuk D, Noulas A, Scellato S, Nicosia V, Mascolo C (2013) Geo-spotting: Mining online location-based services for optimal retail store placement. In: ACM SIGKDD.

22. Luca M (2011) Reviews, reputation, and revenue: The case of yelp. com. Technical report, Harvard Business School.

23. Georgiev P, Noulas A, Mascolo C (2014) Where businesses thrive: Predicting the impact of the olympic games on local retailers through location-based services data. In: AAAI ICWSM.

24. Carter RD, Zieren JW (May-June 2012) Festivals that say cha-ching!measuring the economic impact of festivalst. In: Main Street Now.