Byproducts of Urban Infrastructure Interfaces: Evidence from Parking Compliance

Konstantinos Pelechrinis
School of Information Sciences
University of Pittsburgh
kpele@pitt.edu

ABSTRACT

The increased levels of urbanization have resulted in the demand for developing urban technologies that can realize the vision of smart cities, i.e., urban environments that are sustainable, livable and resilient. Electromechanical infrastructure is substituted by intelligent, cyber-physical infrastructure (e.g., coin-based ticket fare collectors are substituted by smart cards) in an effort to both reduce costs, increase efficiency as well as improve the user-friendliness of the system. Significant efforts and resources have been allocated in the area of public transportation, including the modernization of subway and bus networks. However, one of the most-discussed aspects of public transportation in our automobile-dominated cities is that of parking infrastructure. While research has concluded that appropriate pricing of metered parking zones is essential to allow local businesses to flourish and even reduce congestion, there is still a lot of hesitance on implementing the appropriate policies. Hence, parking zones are still significantly underpriced. The problem is further pronounced by poor enforcement. However, during the last years most of the coin-based parking meters are being substituted by "smart" meters that accept various types of payments (e.g., credit cards, mobile etc.). While these meters have been installed to mainly make parking payments more convenient to drivers, they appear to have important indirect benefits. In particular, in this study we use quasi-experimental techniques to analyze parking citation information from the city of Pittsburgh and we find that the installation of the new parking meters leads to increased compliance with parking rules. This can further have significant implication for the design of the urban infrastructure interfaces of the upcoming smart technology.

1. INTRODUCTION

The increased levels of urbanization [1] have created the need for increased sustainability, livability, resilience and efficiency. The advancements in computing technology has significantly helped towards these efforts, leading to the modernization of various aspects of urban life ranging from the ability to complete local government forms/applications online to having real-time public transportation information in your hand at any time. One of the areas that has significantly benefited is that of transportation. For instance, smart cards allow port authorities to collect detailed origin-destination trip information that can be further analyzed and design routes that are responsive to passengers' demand. Furthermore, localization technologies and near-field communication allow tracking of capacity and location of buses, which further improves the experience of public transit commuters. However, despite the improvements in public transportation made possible, mobility in cities is still dominated by automobiles and the increasing urban population leads to increased congestion and pollution among other side effects.

One of the major challenges associated with automobile mobility is the limited terminal capacity (i.e., curb parking space). The solution urban planners came up with in the mid 20th century for this problem, was to force every new development to create the required infrastructure for serving the increased parking demand from the new establishment. This further increased the urban sprawl that had already started appearing due to views supporting Le Corbusier's Radiant City model. Closing the vicious circle, this led to even more car trips and made parking one of the major transportation problems.

A lot of research has been done on studying the parking behavior of people dating back to the very first parking meter installed in 1935 in Oklahoma City. One recurring conclusion is that quick turnover of curb parking spots is beneficial on many levels, ranging from reduced congestion [19] to increased revenues for local businesses. For example, the latter rely on the arrival of new customers that can only be supported when the turnover of curb parking spots is fairly fast. There are two mechanisms that can be used to control this turnover of the curb spaces, namely, paid parking and time limits. Increased prices can lead to drivers parking for less time, while time limits forces them to leave after a pre-defined amount of time, thus, controlling the spot's turnover.

Despite the fact that curb parking spaces are significantly underpriced [20] many drivers still do not comply with the parking rules - as evidenced by the parking citations - and enforcement mechanisms are required. An interesting question that arises, is **what percentage of this *disobedience* can be attributed to reasons related with the complex *interface* of parking infrastructure ?** For example, parking restrictions signs can be particularly confusing (see

UrbComp'16, August 14, 2016, San Francisco, USA.
.

Figure 1: Physical "interfaces" of transportation infrastructure can be confusing, leading to unintended non-compliance. Source: http://la.curbed.com/2011/8/24/10446060/see-the-most-ridiculous-parking-restriction-sign-ever

Figure 1). Hence, even if a driver wants to comply with the (complex) parking rules, there is a possibility that he simply violated the rules by mistake. Another similar notorious example is that of the traditional parking coin-meters that accept only specific coin denominations !

In this study, using (i) evidence from the parking citations and meters in the city of Pittsburgh and (ii) quasi-experimental, econometric techniques, we quantify the impact of the "pay-by-plate" parking meters installed in the city during 2012 on the drivers' compliance with parking rules. Our main finding indicates that the installation of this new technology led to a reduction of the citations handed by the enforcement officers. In particular, a reduction of 505 citations on average per month was observed as compared to counterfactual if the new meters were not installed.

Related literature: Focusing on parking and transportation in general, a lot of work has appeared that aims into modeling and describing the behavior of commuters with respect to the various aspects of the transportation system. For instance, various dynamic pricing schemes have been suggested in order to achieve an "optimal" utilization by altering drivers behavior (e.g., [17, 5]), while other studies have identified factors that can impact the parking choice of drivers (e.g., [11, 2, 12]). In another direction, computing and information systems for identifying and predicting open curb spots have been developed (e.g., [15, 4, 21, 6, 7, 10, 8, 14]). However, the impact of the electro-mechanical interfaces on the dweller-urban infrastructure interactions have not gained a lot of attention in the existing literature and our work is the first to examine the impact of the parking meter interface on the parking bylaws enforcement.

Roadmap: The rest of the paper is organized as follows. Section 2 describes the data used in this study as well as our research hypothesis, while Section 3 introduces the difference-in-differences method that was used to quantify the impact of the "pay-by-plate" meters. Section 4 presents our analysis and results. Finally, Section 5 concludes our work.

2. DATASETS AND HYPOTHESIS DEVELOPMENT

In this section we will describe the datasets we used for our study as well as the curb parking metering infrastructure of the city of Pittsburgh.

2.1 Citation and Parking Data

In order to perform our analysis we need access to parking citation data. Under the Freedom of Information Act, we requested and obtained from the Pittsburgh Parking Authority the parking citations in 5 neighborhoods in the city of Pittsburgh, namely, Oakland, Shadyside, Squirrel Hill, Downtown and Brookline, for the period 01/01/2011-31/12/2013. Every citation correspond to one data point and the information associated with it include the following tuple: `<ticket #, day, time, street, meter ID, district, bylaw>`. We aggregate these data in a monthly granularity, i.e., the number of citations during month m in neighborhood n are given by $c_n[m]$. Hence, for every neighborhood we have a time-series with 36 data points.

We also obtained a parking meter dataset from the "pay-by-plate" meters installed in the city starting July 2012. In particular, the system logs every payment received by the system in the format: `<Purchase Date, Terminal ID, Payment Type, Ticket #, Payment Amount, Start Time, End Time>`. These parking data logs are especially useful in our study for analyzing the payment types as a function of the payment amount.

2.2 Pay-by-plate Meters

"Pay-by-plate" meters offer the convenience of various types of payment options for parking including credit card. These stations are powered via solar panels and include a cellular interface for communicating with the central database that stores all the payment information. The operator can also alter the parking rates remotely without the need of manually intervening with the meters. These stations started being installed in the city of Pittsburgh during July 2012. Their installation was gradual with the older coin-meters in the neighborhoods of Oakland, Shadyside, Squirrel Hill and Downtown being substituted during July 2012, while the infrastructure in Brookline was updated later, and in particular in the end of May 2013. This setting is crucial for examining the research hypothesis of our work.

Pay-by-plate stations are clearly more user-friendly to drivers since they accept different types of payments and hence, drivers do not have to be equipped with coins for paying. Our hypothesis is that this new interface between the citizens (i.e., drivers) and the transportation infrastructure will lead to increased compliance with parking rules. In other words, this implies that a large fraction of the parking violations are not due to the unwillingness of drivers to pay for parking but due to the "unfriendly" nature of the payment interface. Formally, the research hypothesis that we examine in this study is:

HYPOTHESIS 1. `[Impact of "Pay-by-Plate" stations on parking rules compliance]`: *The installation of "Pay-by-Plate" parking meters has lead to increased levels of compliance by the drivers with the parking rules.*

3. METHODS AND EXPERIMENTAL SETUP

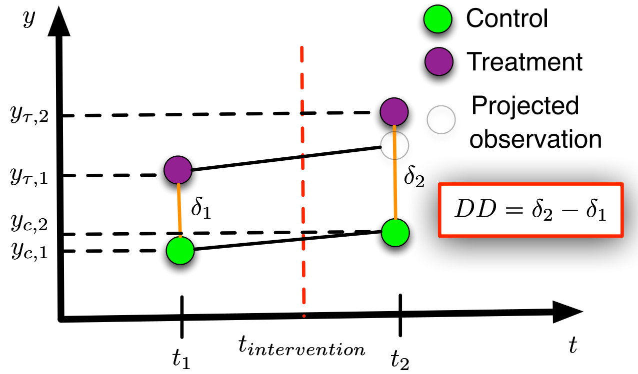

Figure 2: The difference-in-differences method.

In this section we will introduce the analytical methods utilized in our analysis. In particular, we will describe the difference-in-differences as well as statistical bootstrap.

3.1 Difference-in-differences

The difference in differences (DD) method [3] is a quasi-experimental technique that aims in identifying the effect of an intervention using observational data. DD requires observations of the metric of interest y, obtained in different points in time, e.g., t_1 and t_2 ($t_1 < t_2$), for both the control (e.g., $y_{c,1}$ and $y_{c,2}$) and the subjects that receive the treatments (e.g., $y_{\tau,1}$ and $y_{\tau,2}$). The treatment is applied at time t_τ with $t_1 < t_\tau < t_2$, and hence, the treatment subjects are exposed to the intervention only during t_2. The difference between $y_{\tau,2}$ and $y_{c,2}$ quantifies both the impact of the intervention as well as other "intrinsic" differences between the two groups. An estimate for the latter can be captured by the difference between the treatment and the control during time t_1, i.e., $y_{\tau,1} - y_{c,1}$, where none of them was exposed to the intervention. The DD estimate is then (see Figure 2):

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

If $\delta_{\tau,c} > 0$ ($\delta_{\tau,c} < 0$), then the treatment has a positive (negative) impact on y, while if $\delta_{\tau,c} = 0$ there is not any impact from the intervention. Equation (1) 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: $(y_{\tau,2} - y_{c,2}) = (y_{\tau,1} - y_{c,1})$, 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 many times overlooked when the method is applied. However, we formally show in Section 4 that the parallel trend assumption holds in our dataset, and hence, the results obtained are reliable.

Exactly the 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:

$$y_{int} = \alpha_n + \beta_t + \delta \cdot D_{nt} + \epsilon_{int} \tag{2}$$

where y_{int} is the dependent variable for instance i (at time t and neighborhood n), α_n and β_t are binary variables that capture the fixed effects of the neighborhood and time respectively, D_{nt} is a dummy variable that represents the treatment status (i.e., $D_{nt} = \alpha_n \cdot \beta_t$) and ϵ_{int} is the associated error term. The coefficient δ 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 (1). In particular, if $\overline{y}_{nt}$ is the sample mean of y_{int} and $\overline{\epsilon}_{nt}$ is the sample mean of ϵ_{int}, and using Equation (2) we have:

$$\begin{aligned}(\overline{y}_{11} - \overline{y}_{10}) - (\overline{y}_{01} - \overline{y}_{00}) = \quad & \delta(D_{11} - D_{10}) - \delta(D_{01} - D_{00}) \\ & + \overline{\epsilon}_{11} - \overline{\epsilon}_{10} + \overline{\epsilon}_{00} - \overline{\epsilon}_{01}\end{aligned} \tag{3}$$

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

$$\mathrm{E}[(\overline{y}_{11} - \overline{y}_{10}) - (\overline{y}_{01} - \overline{y}_{00})] = \delta(D_{11} - D_{10}) - \delta(D_{01} - D_{00}) \tag{4}$$

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

$$\hat{\delta} = (\overline{y}_{11} - \overline{y}_{10}) - (\overline{y}_{01} - \overline{y}_{00}) \tag{5}$$

which is essentially the same as Equation (1). Therefore, one can estimate the DD using either of the Equations (1) or (2). Figure 2 further visualized the estimation process.

The control and treatment subjects are defined based on the spatial dimension. In particular, we are focusing on neighbors within the city of Pittsburgh and hence, every neighborhood is a single subject. With this setting, the neighborhoods of Oakland, Squirrel Hill, Shadyside and Downtown received their treatment (i.e., pay-by-plate meters) during July 2012, while the neighborhood of Brookline did so during May 2013. One of the experimental setup decisions that we have to make is what exactly are the two time-points that we will examine. Naturally, the pre-treatment period is the period between January 2011 and June 2012, while the post-treatment spans the period between August 2012 and May 2013 (the meters in Brookline were installed during the end of May). Our control neighborhood is Brookline, which did not receive the treatment at any point between January 2011 and May 2013.

Our metric of interest is the number of parking citations handed on a monthly basis. Simply put, the components of Equation (1) are given by:

$$y_{n,T} = \frac{\sum_{t \in T} c_n[t]}{|T|} \tag{6}$$

As alluded to above one of the important things when applying the difference-in-differences method is to verify that the parallel trend assumption holds. Hence, we will use the first 8 months of our dataset to examine this assumption (see Section 4). Consequently, we are using the 9-month period between September 2011 and July 2012 as our pre-treatment period, while the 9-month period between August 2012 and

May 2013 is our treatment period. In order to estimate DD we will calculate the average monthly number of citations handed out in each neighborhood during each period. In order to obtain a robust estimate for the average we will rely on statistical bootstrap, a resampling method that we describe in what follows.

3.2 Bootstrap

Statistical bootstrap [9] is a robust method for estimating the unknown distribution of a population's statistic when a sample of the population is known. The basic idea of the bootstrapping method is that in the absence of any other information about the population, the observed sample contains all the available information for the underlying distribution. Thus, resampling with replacement is the best guide to what can be expected from the population distribution had the latter been available. Generating a large number of such resamples allows us to get a very accurate estimate of the required distribution. Furthermore, for time-series data, block resampling retains any dependencies between consecutive data points [13].

In our study, we will use bootstrap to estimate the distribution of the difference-in-differences estimator. In particular, we will resample with replacement the monthly citations during the period of interest and hence, create the empirical distribution of δ, $f(\delta)$. This will allow us to further estimate the statistical significance of the estimator.

4. RESULTS

In this section we will present the results from our analysis.

Parallel trend: In order to verify the existence of a parallel trend a typical approach that is being followed in the literature [18, 16] is to estimate the difference-in-differences coefficient during a period where there is no treatment (i.e., a pseudo-treatment time will be assigned randomly). In our case, we used the 8 first months of our dataset for this purpose. In particular, we use the period January-April 2011 as the pre-treatment period, while the period May-August 2011 as the (pseudo) treatment period. We then calculate the difference-in-differences between the control neighborhood of Brookline and the rest of the Pittsburgh neighborhoods in our dataset. Our results indicate that the overall null difference-in-differences is 242.75. However, the corresponding p-value is 0.6, which means that one cannot reject the hypothesis that this coefficient is zero[1]. The estimated coefficients for every neighborhood are presented in Figure 3. As we can see there is not any clear positive or negative trend, so **overall there is not significant evidence against the parallel trend assumption**. Also note that a positive value of the null coefficient essentially means that **if** there is a trending component for the citations, these increase faster in the neighborhoods of Shadyside, Squirrel Hill and Oakland as compared to the control neighborhood of Brookline.

Treatment impact: Having provided evidence that validate the presence of the parallel trend assumption between the treatment and control subjects prior to the application of the treatment, we can now examine the impact of the treatment. As alluded to above, we use as our pre-treatment period the time between September 2011 until July 2012, while the treatment period is between August 2012 and May 2013. During each of these periods and for each one of our subjects (i.e., neighborhoods), we perform bootstrap 100 times for estimating the average monthly citations. Consequently we obtain to corresponding difference-in-differences estimator and its p-value. Table 1 presents our results. As we can see the difference-in-differences is negative for all the neighborhoods (and statistically significant at the 0.01 level), which means that **the installation of new parking meters has lead to an increase in the driver's compliance with respect to the parking bylaws**. Note here that, our analysis essentially controls for other potential confounding factors such as reduced enforcement, pay rate changes etc. since both the treated and control subjects are exposed to the same externalities; the only difference is the presence of a pay-by-plate station.

[1]Of course, when the null hypothesis cannot be rejected, this can be due to an under-powered test. A small sample typically does not provide enough statistical power to detect small but significant differences. While this is possible in our case, we can confidently reject the hypothesis that there is a "large" trending component.

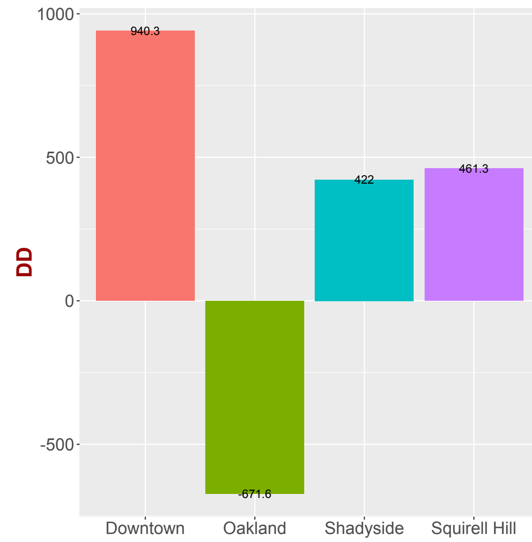

Figure 3: The parallel trend assumption that needs to be satisfied for the difference-in-differences method to provide robust results and conclusions appears to hold.

Neighborhood	δ	p-value
Squirrel Hill	-399.776	2.374^{-8}
Shadyside	-109.204	0.008235
Oakland	-1252.3	$<2.2e^{-16}$
Downtown	-262.668	6.532^{-6}

Table 1: The installation of pay-by-plate parking meter stations has lead to a decrease in the number of parking violations cited by the Pittsburgh Parking Authority enforcement officials.

The underlying reasons for this phenomenon cannot of course be revealed by our analysis, but it does not seem very plausible that the new stations suddenly made the drivers in Pittsburgh more obedient. Our hypothesis is that the friendly interface of the new parking meter stations, and in particular their ability to accept credit card payments, gave the opportunity to drivers to pay for parking in cases where they would not have otherwise (e.g., because they did not have enough coins). To explore this further, we analyzed the method payments logged in the new system, and more specifically the number of tickets paid through credit cards and the number of tickets paid with coins. In particular, we are interested in examining the cost of parking tickets paid by card and by coins. An important thing to note here is that when paying with credit card there is a minimum charge of $1, regardless of whether the real ticket costs less or not.

The parking log data indicate that the median ticket price paid with coins is 75c while the mean is 0.97c, both less than minimum credit card charge of $1. Hence, drivers appear to prefer to pay for their tickets with credit cards unless this incurs over-payment. In particular, from all the tickets that cost more than $1, 76% of them were paid with credit card! To reiterate, these are just observational results and as such we cannot be sure on whether these credit card payments are due to the convenience of paying by card, or due to the "power of habit" of using a card or simply because the driver did not have any quarters. Nevertheless, they point to important evidence that having the option to pay with credit card makes it easier for drivers to pay for parking and hence, obey the bylaws. We believe that many of the tickets that were paid by credit card - especially the ones with high cost - would not have been able to be paid if coins was the only option for payment.

5. CONCLUSIONS

In this study we examined the byproducts of user-friendly urban interface design. In particular, we focused on the case of parking meters and on their impact on drivers' compliance with parking bylaws. We used quasi-experimental techniques, and in particular, the difference-in-differences method, and found that the number of citations recorded have significantly reduced after the installation of the new infrastructure. In the future, we plan on examining further improvements in the urban infrastructure and their impact on other aspects. For example, the parking authority in the city of Pittsburgh has recently added a capability for mobile payments, which has also eliminated the minimum charge of $1. We believe that our work will trigger more research on the design of effective urban infrastructure interfaces as we move towards the implementation of the smart cities vision.

Acknowledgements: We would like to thank the Director of Meter Services of the Pittsburgh Parking Authority, Wesley Pollard Jr., for his help with accessing the data and his comments on our work.

6. REFERENCES

[1] United nations-world urbanization prospects: The 2011 revision - highlights: http://esa.un.org/unup.

[2] Modeling parking behavior under uncertainty : a static game theoretic versus a sequential neo-additive capacity modeling approach, 2013.

[3] O. Ashenfelter and D. Card. Using the longitudinal structure of earnings to estimate the effect of training programs. *The Review of Economics and Statistics*, 67(4):648–60, 1985.

[4] D. Ayala, O. Wolfson, B. Xu, B. DasGupta, and J. Lin. Parking in competitive settings: A gravitational approach. In *Mobile Data Management (MDM), 2012 IEEE 13th International Conference on*, pages 27–32. IEEE, 2012.

[5] D. Ayala, O. Wolfson, B. Xu, B. DasGupta, and J. Lin. Pricing of parking for congestion reduction. In *Proceedings of the 20th International Conference on Advances in Geographic Information Systems*, SIGSPATIAL '12, pages 43–51, New York, NY, USA, 2012. ACM.

[6] F. Caicedo, C. Blazquez, and P. Miranda. Prediction of parking space availability in real-time. In *Elsevier Expert Systems with Applications, Vol. 39, No. 8*, June 2012.

[7] M. Caliskan, A. Barthels, B. Scheuermann, and M. Mauve. Predicting parking lot occupancy in vehicular ad-hoc networks. In *IEEE VTC*, 2007.

[8] T. Delot, S. Illari, S. Lecomte, and N. Cenerario. Sharing with caution: Managing parking spaces in vehicular networks. *Mobile Information Systems*, 9:69 – 98, 2013.

[9] B. Efron and R. Tibishirani. *An Introduction to the Bootstrap*. Chapman and Hall/CRC, 1993.

[10] Y. Geng and C. Cassandras. Dynamic resource allocation in urban settings: A "smart parking" approach. In *IEEE CACSD*, 2011.

[11] S. Hess and J. W. Polak. An analysis of parking behaviour using discrete choice models calibrated on sp datasets.

[12] S. Hess and J. W. Polak. Mixed logit estimation of parking type choice.

[13] H. Künsch. The jackknife and the bootstrap for general stationary observations. *Annals of Statistics*, 17(3):1217–1241, 1989.

[14] N. Moini, D. Hill, and R. Shabihkhani. Impact assessments of on-street parking guidance system on mobility and environment. In *Transportation Research Board 92nd Annual Meeting*, 2013.

[15] S. Nawaz, C. Efstratiou, and C. Mascolo. Parksense: A smartphone based sensing system for on-street parking. In *Proceedings of the 19th annual international conference on Mobile computing & networking*, pages 75–86. ACM, 2013.

[16] K. Pelechrinis, M. Kokkodis, and T. Lappas. On the value of shared bike systems in urban environments: Evidence from the real estate market. *Available at SSRN*, 2015.

[17] G. Pierce and D. Shoup. Getting the prices right: an evaluation of pricing parking by demand in san francisco. *Journal of the American Planning Association*, 79(1):67–81, 2013.

[18] W. R. Shadish, T. D. Cook, and D. T. Campbell. *Experimental and Quasi-Experimental Designs for Generalized Causal Inference*. Cengage Learning, 2001.

[19] D. C. Shoup. Cruising for parking. *Transport Policy*, 13(6):479–486, 2006.

[20] D. C. Shoup et al. *The high cost of free parking*, volume 7.

[21] B. Xu, O. Wolfson, J. Yang, L. Stenneth, P. S. Yu, and P. C. Nelson. Real-time street parking availability

estimation. In *Mobile Data Management (MDM), 2013 IEEE 14th International Conference on*, volume 1, pages 16–25. IEEE, 2013.