New Algorithms for Parking Demand Management and a City Scale Deployment

Onno Zoeter Christopher Dance Stéphane Clinchant Jean-Marc Andreoli
Xerox Research Centre Europe
6 Chemin de Maupertuis
38240 Meylan
France

ABSTRACT

On-street parking, just as any publicly owned utility, is used inefficiently if access is free or priced very far from market rates. This paper introduces a novel demand management solution: using data from dedicated occupancy sensors an iteration scheme updates parking rates to better match demand. The new rates encourage parkers to avoid peak hours and peak locations and reduce congestion and underuse. The solution is deliberately simple so that it is easy to understand, easily seen to be fair and leads to parking policies that are easy to remember and act upon. We study the convergence properties of the iteration scheme and prove that it converges to a reasonable distribution for a very large class of models. The algorithm is in use to change parking rates in over 6000 spaces in downtown Los Angeles since June 2012 as part of the LA Express Park project. Initial results are encouraging with a reduction of congestion and underuse, while in more locations rates were decreased than increased.

Categories and Subject Descriptors

H.2.8 [**Database Management**]: Database Applications—*Data Mining*

General Terms

Algorithms

1. INTRODUCTION

Advances in sensor technologies and advances in information dissemination technologies, most notably the widespread use of smartphones, allow for a radical change in the way governments manage public infrastructure. The deployment of sensors allows this to be data driven, and the information channels allow this to be adaptive. It forms an interesting application domain for data mining, machine learning, and optimization techniques. And in line with this year's KDD special theme it is a good opportunity for data mining for social good since it allows scarce public utilities to be used more efficiently and helps to reduce unwanted externalities such as congestion and pollution. This paper introduces new algorithms that leverage these new sensor and dissemination opportunities to improve the way on-street parking is used.

KDD'14, August 24–27, 2014, New York, NY, USA.

ACM ACM 978-1-4503-2956-9/14/08 ...$15.00.
http://dx.doi.org/10.1145/2623330.2623359 .

On-street parking is in many city centers a scarce resource. If on-street parking is free, or priced significantly below market rates, it will be used inefficiently, because drivers are not properly incentivized to avoid peak hours and peak locations. William Vickrey [9] argued that users of on-street parking, and in fact users of any publicly owned utility, should be charged as close as possible to the externality they impose (i.e. the inconvenience they cause others).

Imagining a solution using 1950s technology, Vickrey proposed to connect 20 meters and to make them turn faster as the 20 spaces filled up. The meter could turn slowly if there were less than 17 cars parked (anybody that would want to park in the block could easily do so), would turn faster if the block filled up from there, and would turn fastest if all 20 spaces were full.

There are several problems with these ex-post (pay at the end) meters. The current infrastructure is not ready for it and the upfront uncertainty about what parking costs might not be acceptable to users. But perhaps more importantly, the ex-post meters put the burden of predicting demand for parking on the shoulders of every driver, while the overall parking system has all the data to do this best!

Several parking specialists (e.g. [8]) have suggested a more cautious approach: revise parking rates at regular intervals, say every month, based on data obtained from parking sensors. That way drivers can memorize the rates around their office, favorite restaurant, etc. and adjust their behavior accordingly.

1.1 Contributions

This paper studies the problem of iteratively improving on-street parking rates based on observed parking demand data to increase the efficiency with which these scarce resources are used and to reduce congestion and pollution. It makes several contributions.

- It describes why the often proposed straightforward method of basing rate updates on average occupancy data does not correspond to a reasonable utility model and can lead to incorrect rate changes (Sec 3.1).

- It introduces an iteration scheme that is based on a trade-off between congestion and underuse. This iteration scheme is simple to communicate, fair, not based on any model assumptions, and solves the problems with the naive scheme (Sec 3.2).

- The limiting behaviour of this iteration scheme is analyzed and proven to converge to a reasonable distribution under a very large class of models. This is particularly useful since there was no data available to build and select models before the first deployment of our methods (Sec 3.3).
- An extension of the algorithm automatically determines time-of-day windows for rates trading-off the closeness to the parking patterns and ease of communication (Sec 4).
- The paper describes the real-world deployment of these ideas. The developed methods have been used to adjust parking rates for 6300 on-street parking spaces since June 2012 in Los Angeles (Sec 5).

2. VCG PAYMENTS, POSTED PRICES AND FREE PARKING

Before we discuss iteration schemes in the next section, it is interesting to study the motivations for demand management in more detail. A simple stylized example to demonstrate the impact of parking rates on efficiency is presented in Figure 1.
This considers a single block with:

- capacity $C = 10$, a single time step
- $n \sim \text{Poiss}(\lambda = 15)$ parkers
- values for parking $v_i \overset{\text{iid}}{\sim} P(v) = \text{Gam}(\alpha = 5, \beta = 2)$ for $i = 1, \ldots, n$.

Free parking implies a first-come-first-served mechanism. A posted price mechanism will filter parkers with a very low value for parking such that spaces are kept for parkers with a high valuation. The figure shows the sum of values of parked cars (social welfare) as a function of parking rate (blue, solid line shows mean, dashed 25% and 75% quantiles). Whereas we see that a non-zero posted price can indeed improve social welfare, a too high price risks leaving spaces empty. The horizontal line shows the expected social welfare under an optimal allocation, e.g. if parking allocation could be done using a Vickrey-Clarke-Groves mechanism [10]. The green lines indicate revenue. As expected the revenue maximizing rate is higher than the welfare maximizing rate [6]. The second set of lines is added to emphasize the difference in objectives: demand management aims to increase efficiency, yield management to increase revenue.

A properly chosen take-it-or-leave-it price can provide a constant factor approximation to the optimal revenue ([1], where the same result is also claimed for social welfare). The result depends on the assumption that valuations are drawn independently. For parking several factors could break this assumption. For instance in front of a stadium all parkers will have a low-ish valuation if there is no match on, but all will have a high-ish one if there is. So if we can only set a single rate, parkers' valuations are dependent. However if we have a match-time rate and a non-match rate, for each of the two cases independence is a reasonable assumption. So for demand based pricing it is important to condition rates on important factors that significantly influences valuations.

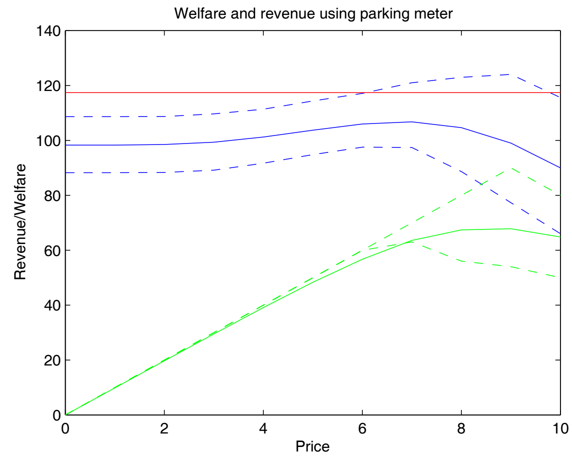

Figure 1: Welfare (dark, blue) and revenue (light, green) as function of price in a simple parking model. Solid lines represent means, dashed lines 25% and 75% quantiles. The horizontal line is the expected welfare under optimal assignment.

The work from [1] is studied in an on-line bipartite matching problem inspired by the parking problem in [5].

The stylized example of this section also makes clear that defining a utility model for a realistic parking scenario is not straightforward. In the example only the interests of the parkers have been taken into account. It could be argued that the congestion caused by drivers circling for a space should also be taken into account, or the interests of residents, shopkeepers, or others. Similarly any parking demand model will be non-trivial in a realistic setting and cannot be learned before several rate changes have been made. Even if many years worth of data is available before the first rate change, it will not distinguish between models that aim to predict how demand changes after a rate change.

3. DEMAND BASED PRICING

To get robust improvements that do not rely on a particular parking model, we use a simple iteration scheme that is inspired by stochastic approximation. Rates are on a discrete pricing ladder (in LA rates are in {\$0.5, \$1, \$1.5, \$2, \$3, …, \$6}). At the end of every review period (say every quarter, or every month) parking data is studied and for every block-face (one side of a street in between two side streets) it is determined if the rate should go up one step, down one step, or should stay the same. The review period cannot be too short, because drivers need to learn about the new rates and change their habits, before data about the new equilibrium becomes available. Intuitively if we have observed that demand has been too high, rates should go up, if demand has been too low, rates should go down. But as we will see in the next section, care has to be taken to implement that intuition.

3.1 The problem with updates based on average occupancy

It is interesting to observe that the simplest possible approach, used e.g. in a project running in parallel in San Francisco (`sfpark.org`), basing rate changes on the average occupancy in a review period, is not true to most reasonable utility models. Most utility models would say that underuse is bad (possible shoppers or other users have been discouraged and spaces remain unused that would be used at a lower rate) and congestion is bad (parkers with a high valuation are blocked, drivers circling for a space congest and pollute). If rates are changed based on average occupancy, a period of underuse followed by congestion can lead to a perfect "neither too empty, nor too full" (Goldilocks) situation — on average, and hence mask the problems that have actually occurred. This is of particular importance if time windows are determined algorithmically based on data: a quiet morning followed by a congested afternoon risk to be grouped together, since jointly they will have a perfect average occupancy. The data from the LA deployment indicates that this is a common case in practice (See Figures 4 and 9).

3.2 Rate change rules based on congestion and underuse patterns

A simple but effective alternative is to study the number of minutes each block-face was congested and how many minutes the block-face was underused. Several practitioners (e.g. [8]) have considered 85% occupancy per block-face as an ideal occupancy target. The city of LA has decided to treat occupancies below 70% as underused, and above 90% as congested. Such a discrete classification is a simplification, but forms a reasonable and easy to communicate basis for policy change.

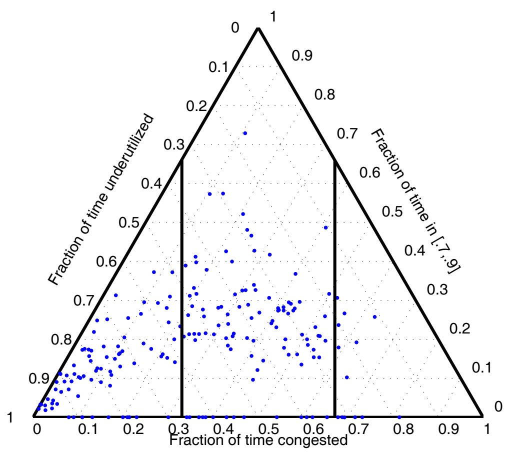

Figure 2: A ternary plot demonstrating the rate change rules. At the end of a review period each block is represented by three numbers: the fraction of time underused, the fraction of time congested, and the fraction of time just right. Since these numbers add to one, a block can be represented in a ternary plot. Blocks (dots) in the right region have a fraction of time congested minus fraction of time underused greater than 1/3 and see a rate increase. Blocks in the left region see a rate decrease. Data from April-May 2012.

At the end of every month, per block, the fractions of time spent in each of the three categories (underused, just right, congested) adds to one, and can be represented in a ternary plot. Figure 2 shows the 817 blocks in the LA project area based on data from April and May 2012 before the first rate change.

The blocks that fall close to the bottom right corner are predominantly in congested state, and a rate increase would be in order. Blocks close to the bottom left corner are predominantly underused, and rates should be reduced. Points at the top are predominantly just-right and could keep the same rate. Points in the middle close to the bottom demonstrate both non-negligible problems of congestion and of underuse. A change to a single rate can only try to improve one of the two problems at the risk of worsening the other. Studying the patterns of demand might reveal that the underuse and congestion patterns consistently appear at different parts of the day. In such cases charging different rates at different parts of the day can target the two problems appropriately. We will introduce algorithms to optimally determine such windows based on data in Section 4.

A rate change rule based on the difference between the fraction of time congested and the fraction of time underused, is a simple rule that agrees with the above sketched ideas. Rates are increased only if congestion is the predominant problem, reduced only if underuse is the predominant problem, and conservatively kept the same if both congestion and underuse form a problem. The rate changes in LA are based on thresholds 1/3 and −1/3: the regions demarked by the vertical lines in Figure 2. Algorithm 1 summarizes the method and introduces notation.

Algorithm 1 The default rate change iteration

For all blocks b

- Compute the *congestion index* $I_c^{(b)}$ as the fraction of operating hours in the review period that b is congested (occupancy $>$ %90).
- Similarly compute the *underuse index* $I_u^{(b)}$ (underuse defined as occupancy $<$ %70).
- Define the *congestion-underuse balance* as

$$B_{cu}^{(b)} = I_c^{(b)} - I_u^{(b)} .$$

For all b with $I_{cu}^{(b)} > 1/3$ (congestion dominant problem)

Increase the rate by one step in the ladder.

For all b with $I_{cu}^{(b)} < -1/3$ (underuse dominant problem)

Decrease the rate by one step in the ladder.

3.3 Studying the limiting distribution

The iteration scheme has a resemblance to stochastic approximation [7]. A big difference is that the updates here are restricted to a discrete set of easy-to-remember rates. Whereas stochastic approximation relies on a specific cooling schedule for the step sizes to guarantee convergence, the discrete nature of the rates prevents such convergence to a point.

In practice we observe significant mid- and long-term fluctuations in demand. In LA we see seasonal effects, changes apparently coming from the opening of new restaurants, shops, etc., and from trends that could come form general macroeconomic shifts. So one could argue that convergence to a single "perfect" rate is not desired. However it is interesting to understand the iteration scheme. To do so we study it in the context of stationary demand.

Let us denote the demand distribution $P(z|r)$, with z the sequence of observed occupancies during a review period for a particular block and $r \in \{r_1 \leq r_2 \leq \ldots \leq r_L\}$ its hourly rate. Together with the rate change rules represented in Figure 2 this defines $P_{i,i+1}$, the probability of a rate increase from i to $i+1$, and $P_{i,i-1}$, the probability of a decrease. If the demand is stationary the iteration defines a stationary Markov chain on the rates.

To be able to motivate rate changes to the general public it is very reasonable, if not required, to consider only one step up or down in the ladder. That is, the transition matrix is tri-diagonal. Any non-pathological demand distribution will have a non-zero (but possibly very small) probability of increasing or decreasing any rate and will define an aperiodic Markov chain. Hence the Markov chain will be ergodic and will have a unique stationary distribution.

This stationary distribution can be characterized under very general conditions. The only mild assumption we will make is that the combination of demand distribution and rate change rules are such that a suggestion to increase rates becomes more likely as rates go down, i.e. $P_{i+1,i+2} \leq P_{i,i+1}$. For instance for the same location we expect to see a rate increase sooner if parking costs 50 cents per hour than if it is \$1 per hour. Similarly we assume that rates are more likely to go down if the current rate is higher: $P_{i+1,i} \leq P_{i+2,i+1}$. These monotonicities are desired for any rate change rule, and will for our choice hold under many models, but cannot be formally guaranteed without specific knowledge of the underlying demand distribution.

Under these assumptions the stationary distribution is uni-modal. To show this we first observe that the stationary distribution for tri-diagonal Markov chains can be easily characterized using the transition probabilities.

LEMMA 1. *For every tri-diagonal transition matrix P there exists a vector s such that $s_i P_{i,j} = s_j P_{j,i}$ for all i and j.*

PROOF. For all diagonal and for all off-tridiagonal elements the equality holds trivially. So the only non-trivial equations that need to hold are $s_i P_{i,i+1} = s_{i+1} P_{i+1,i}$ for $i = 1, \ldots, L-1$. These can be made to hold, and we can get the stationary distribution (up to a normalizing constant) by setting $s_1 = 1$ and $s_{i+1} = s_i P_{i,i+1}/P_{i+1,i}$ for $i = 1, \ldots, L-1$. □

This allows us to demonstrate uni-modality of the stationary distribution and characterize its mode.

THEOREM 1. *If the demand distribution is stationary and the rate change rules are such that $P_{i+1,i+2} \leq P_{i,i+1}$ and $P_{i+1,i} \leq P_{i+2,i+1}$ for all i, the stationary distribution s_i over rates is uni-modal with a mode at the smallest i with $\frac{s_{i+1}}{s_i} = \frac{P_{i,i+1}}{P_{i+1,i}} < 1$, or L if there is no such i.*

PROOF. To show unimodality of the stationary distribution we look at the likelihood ratio of subsequent states. Since $s_i P_{i,i+1} = s_{i+1} P_{i+1,i}$ we have $\frac{s_{i+1}}{s_i} = \frac{P_{i,i+1}}{P_{i+1,i}}$. To check monotonicity we would like to have that as soon as this ratio becomes smaller than 1, the ratio stays below 1. So $P_{i,i+1} < P_{i+1,i} \rightarrow P_{i+1,i+2} < P_{i+2,i+1}$. This holds since $P_{i+1,i+2} \leq P_{i,i+1}$ and $P_{i+1,i} \leq P_{i+2,i+1}$ from the monotonicity of our pricing engine rules.

The ratio of the subsequent states and the unimodality leads directly to the characterization of the (not necessarily unique) mode which is given by the smallest i with $\frac{s_{i+1}}{s_i} = \frac{P_{i,i+1}}{P_{i+1,i}} < 1$, or L if there is no such i. □

With one year of data and only one or two changes up or down allowed by city ordinance limits, we have not enough data to study these properties in practice. Studies using artificially generated data demonstrate what we would expect: the stationary distribution is sharply peaked around a value that avoids both congestion and underuse (See Figure 3). A longer period between rate changes (more observations) leads, in a stationary model, to less variances in the parking patterns observed, and therefore to a more peaked stationary distribution.

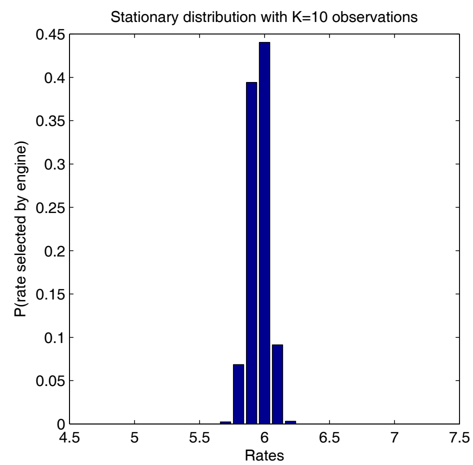

Figure 3: An empirical study of the convergence of the iteration. The plot shows the stationary distribution. The analysis is based on a truncated Poisson demand model with arrival rate function $\lambda(p) = 3(1 - 1./(1 + \mathbf{exp}(-(p-5)))) + 0.05$, parking price per-hour $p \in \{5, 5.1, \ldots 7\}$ and a block-face capacity of 20.

4. DEMAND DRIVEN TIME-OF-DAY WINDOWS

Parking demand clearly varies over different times of day and between weekdays and weekends. If the differences are significant, a single rate cannot be appropriate for all times. In Figure 4, the evolution of I_c (an aggregate over blocks) is shown over the week. Midday on a weekday is on average the most congested whereas parking is overall underused during nights.

If a block-face could have different rates at different times, it could potentially lead to increased effectiveness of the parking system. Following the externality based pricing principles outlined in Section 2 we would require parkers to pay a "blended rate", i.e. the integral under the rate function and not the rate at arrival time for the entire duration. In this section we introduce a suitable logic to choose nearly optimal, yet simple time-of-day windows.

Suppose we would ignore the fact that rates need to be memorized by drivers and therefore ignore the need for the

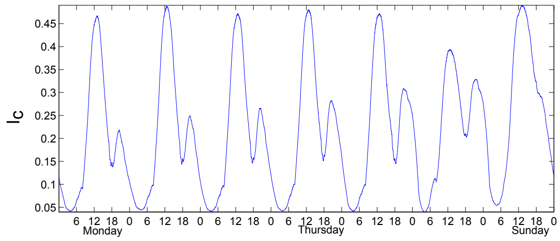

Figure 4: I_c aggregate over the blocks

rate structures to be simple. In particular suppose we could price each half-hour period differently. This is illustrated in Figure 5 on a deployment in over 800 blocks (See next section). The figure represents block-faces as rows. Every pixel represents a vote for a half hour in the week, starting with the first half hour of Monday on the left. The vote is color coded: black means a rate reduction, grey an unchanged rate or unpriced and white a rate increase. The blocks are sorted by the number of half hours that vote for an increase. An important comment is that even if B_{cu} can be low, some faces reached \$0.5 (the lowest rate allowed) so they do not get a price decrease. Secondly, one can observe that a period of low occupancy before 10 AM is common throughout the week. Finally, it is interesting to see that weekdays are rather similar, even if Mondays are slightly less busy than Fridays.

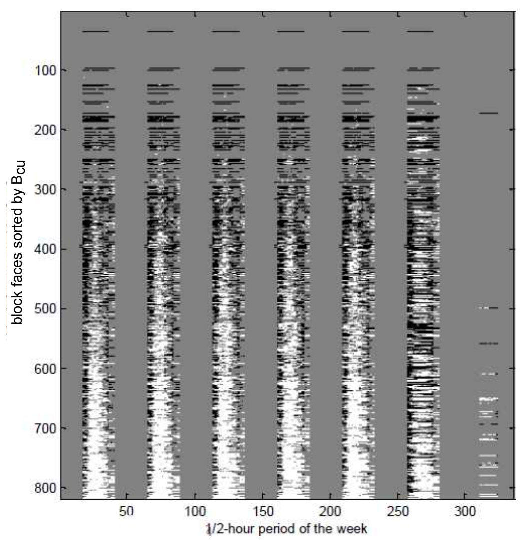

Figure 5: A matrix representing a block-faces (rows) and their vote for a rate change every half hour period in a week. Black means a suggested reduction, grey staying the same or being unpriced, white a rate increase.

Since the information about rates needs to be memorized by drivers before they can act on them, using the suggestions from Figure 5 directly is not acceptable. Encouraged by the similarity among blocks it is interesting to explore a segmentation in time-of-day windows that are *identical all over the city* while minimizing the number of half-hour periods that get the wrong price change. The fact that the windows are the same all over the city makes for a huge simplification and makes the system a natural extension of the often observed practice to have different rates for weekdays and weekends.

The optimal number of windows found is 5 and is illustrated in Figure 6 where it is compared to half-hour votes of 100 E 14th Street. For increased simplicity we add the constraint that we do not want too short segments (< two hours). The optimal partition has 5 segments, but that with the no-short-segments constraint has only 3 segments.

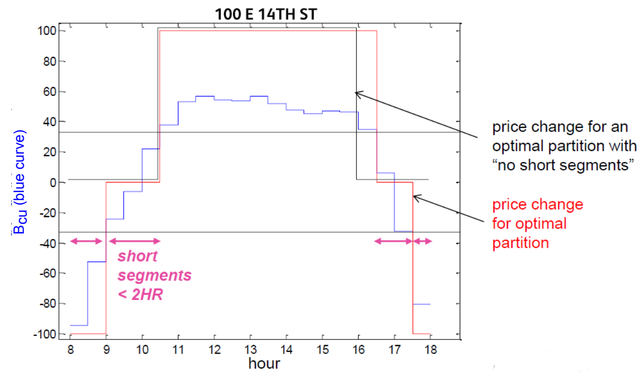

Figure 6: An example of suggested rate changes if they would be done on each half hour of the day independently (blue), and the optimal over the pilot area (red), and with an additional constraint that they cannot contain short (<2 hrs) sections.

Our full analysis on the Los Angeles pilot area, revealed that a simple partition into 3 parts for weekdays is optimal for over 90% of all stalls: (open-11:00,11:00-16:00,16:00-close) We also compared a solution with the 2 best partitions over the pilot areas but found little benefit to it. Furthermore, a single partition is much easier to communicate and to remember. We also analyzed the sensitivity of this partition to this B_{cu} threshold defined by the city of Los Angeles. If B_{cu} were concentrated near the threshold (e.g. 33%), then we would expect that the choice of partition might be rather sensitive to the choice of threshold. However, data shows the contrary. This is a pleasant property because it allows for adjustments to the thresholds on when rates increase or decrease without the data driven time-of-day windows becoming inapproporiate.

In Figure 7, we show a map of block-faces where the fraction of mispriced hours with flat pricing and Time-of-Day Pricing is color-coded. Flat pricing averages the B_{cu} criterion over a large period of time and can select rates that are inappropriate to deal with congestion and underutilization (see for example blocks colored in yellow on the map). This is often the case when operating hours are not well adjusted to the typical usage of a block-face. Time-of-Day pricing greatly reduces this phenomenon as shown on the map.

So, to our surprise, we found that over 800 block-faces could be priced using only 3 time-of-day segments for weekdays and a single one for weekends while being close to what data would suggest on individual half-hours. This simplification makes it easier for drivers to remember rates, and to pre-plan parking destinations accordingly. The early morning rates (open-11am) typically can be interpreted as a “smoothed-start” to paid hours, e.g. before shops open. As a

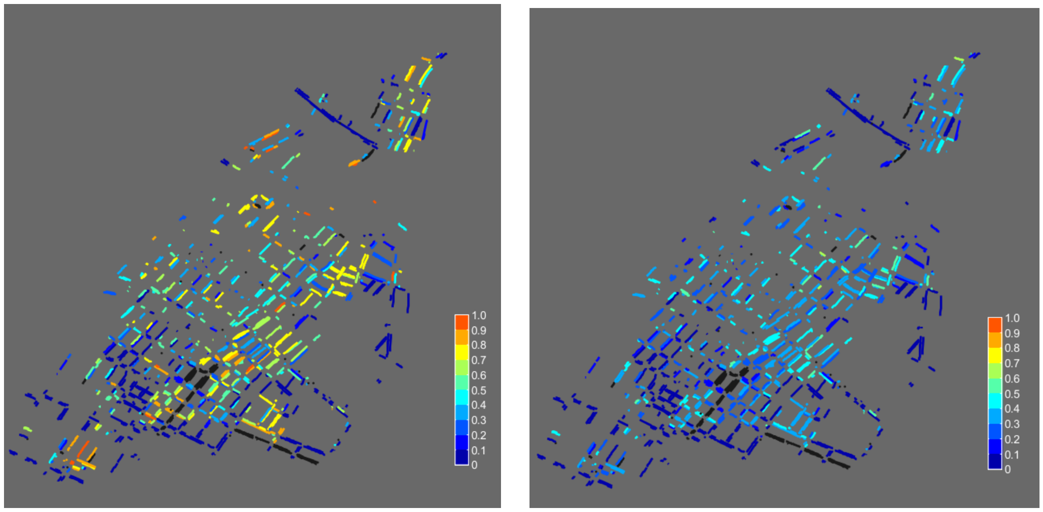

Figure 7: A map of the fraction of mispriced hours over 6 weeks of data with flat pricing (left plot) and time-of-day pricing (right plot) with 3 rates for weekdays and 1 flat rate for Saturdays. The map represents the project area as presented in Figure 8, for clarity no cartography data is shown.

last remark it is interesting to compare to a non data-driven choice of 9am-noon, noon-3pm, and 3pm-6pm as is done in e.g. San Francisco. Judging Figure 4 such a choice would lead to windows that have both underused and congestion in the same window, something that the introduction of time windows exactly tries to reduce!

5. A LARGE SCALE DEPLOYMENT: THE LA EXPRESS PARK PROJECT

The methods described in this paper are used in a large scale deployment in downtown Los Angeles in the LA Express Park project [2]. The goals of this $18.5 million federally funded project are: 1) to increase the availability of on-street parking; 2) to reduce traffic congestion and pollution and 3) to encourage a shift in travel choices. By applying the principles of demand-based parking pricing, the Los Angeles Department of Transportation (LADOT) sought to improve the distribution of parking so that ten to thirty percent of the on-street spaces on each block would be available most of the time.

Prior to LA Express Park, many block clusters had no available on-street parking spaces while others remained practically empty. Increasing the parking prices in high demand areas and lowering the prices in low demand areas shifted demand to yield a better distribution of parking within small geographic areas.

5.1 A Dedicated Sensor Infrastructure

The project depends upon the integration of new wirelessly communicating parking meters, real-time parking guidance systems (in the form of smartphone apps) and in particular upon on-street vehicle sensors that have been installed specifically for the project.

Vehicle sensors were installed in the paving for each of the 6,300 on-street parking spaces in the project area. The spaces are represented as dots in Figure 8. Measurements are based on a magnetometer. The sensors are battery operated and communicate through a wireless mesh network. A sensor provides the occupancy status of a space in real time. The real time data stream provides entries of the form $\langle$Space ID, Timestamp, [Arrival, Departure, Unknown]$\rangle$. The "Unknown" label is emitted if the sensor auto-detects that it cannot reliably determine the occupancy status. This can be due to several causes, among them interference and communication problems. It is particularly important to correct for these sensor inaccuracies, since ignoring data, or simply treating gaps to be missing at random might lead to biased estimates [11].

The sensors in the current generation have a lifetime of around three years. The sensor technology is progressing very rapidly with many manufactures supplying systems. To benefit from these rapid improvements it is important for a demand management system that it can work with any underlying sensor methodology and can automatically learn its specific failure characteristics from data. By learning the demand and noise model jointly this can be done with reasonable accuracy without detailed knowledge of the local interference sources, or the physics of the sensor [11].

The parking meters operate wirelessly as well and provide payment data. There is not a one-to-one mapping between parking and payment events due to multiple payments in one parking session ("topping up the meter"), and motorists that make use of money that is left in the meter when they arrive, and parkers that do not pay. We define *non-paying customers* to be those motorists that do not make a single payment during their stay. This concept is important because in California motorists with a reduced mobility that have a special permit can park at any space for free for an unlimited duration. In practice we see that in the very busy parts of downtown LA it is common to see over 90% of the spaces to be taken by non-paying customers. See also Section 6.2.

After correcting for sensor failures and combining with the payment data, the datasets can be processed off-line to yield entries of the form $\langle$Space ID, Arrival Time, Duration, [Paying, Non-Paying]$\rangle$.

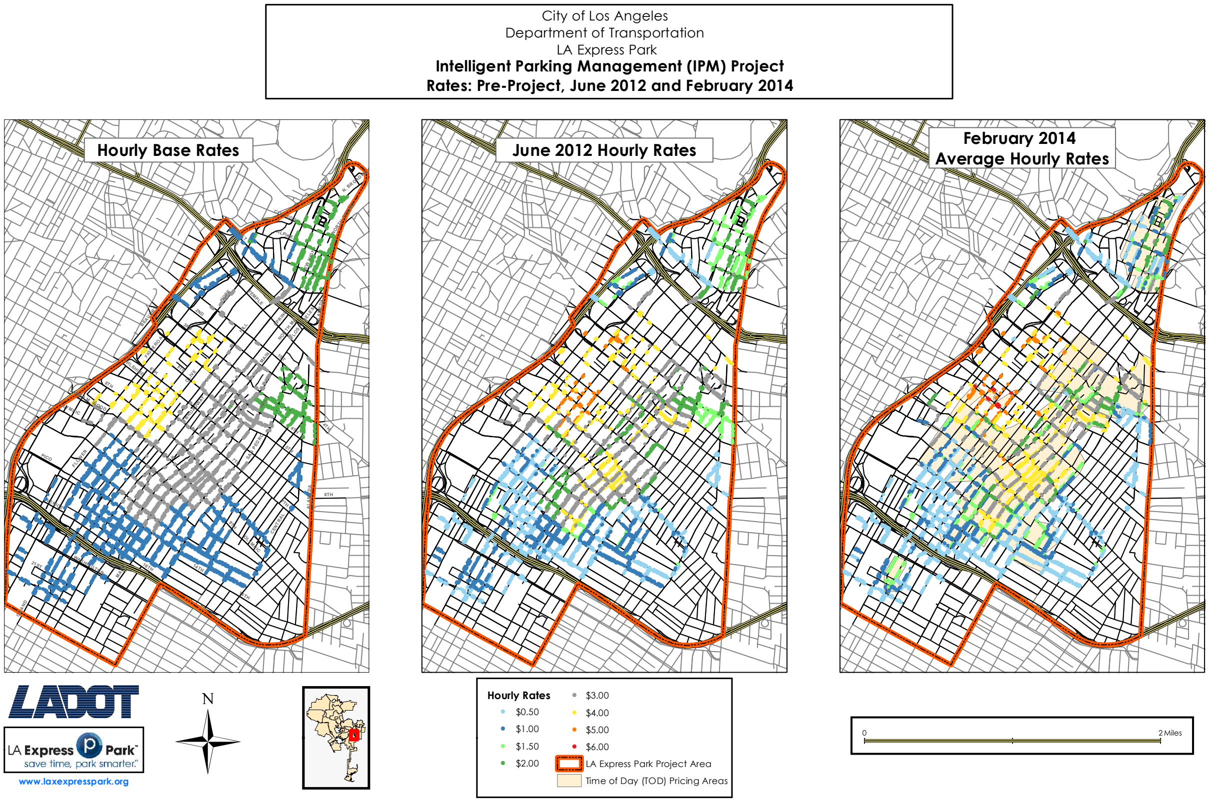

Figure 8: LA Express Park project area with pre-project rates (left) rates after the first change (middle), and rates at time of writing.

There is a large variability both in space and time, but on average a space sees close to 24 parking events per day. So with 6300 sensors and over 600 days of data our dataset consists of nearly 100 million parking events.

6. A CLOSER LOOK AT PARKING DATA

In the previous sections, we have discussed the foundations of the pricing engine: the congestion-underuse balance and its extension to time of day pricing. Parking data can be examined with different approaches:

- with simplistic raw occupancy fraction statistics;
- using arrival duration patterns for a given block-face;
- with temporal methods, assessing effects from the time of day, season, and (regular) events.

In the following sections, patterns of occupancy and arrival durations will be illustrated with block-faces from the pilot area. We then will discuss a first analysis of rate changes as reported by LADOT in [2]. We refer the interested reader to the forthcoming second report by LADOT [3] for a further analysis.

6.1 Occupancy

Figure 9 shows a year's worth of occupancy data for the 12 spaces at 701 South Olive Street in a colormap, sometimes also referred to as a "Jerry Garcia plot". Each row represents a weekday, so that from top to bottom Monday follows Tuesday and so on, with Saturday and Sunday omitted to ensure that the often different demand in the weekend does not hinder the study of the regular weekday patterns. The x-axis represents the minute of the day starting from midnight on the left, to midday in the middle, to midnight at the right. Every pixel color codes the fraction of spaces that are occupied. Blue tones represent occupancies less than 70%, green tones 70%- 90%, and red tones more than 90%. The horizontal gray lines represent policy changes. Plots like these provide a wealth of insight, we focus here on the impact of rate changes. If the change in rate is sufficiently high the change in behavior can be drastic. After the extension of operating hours from 6 p.m. to 8 p.m. in June 2012, we see a drastic change in occupancy patterns from red (congested) to blue and green (low and perfect utilization). This is a clear example that confirms the hypothesis that behavior can be changed by adjusting parking rates. We see the biggest changes in examples like these, where flat rates are maintained and operating hours are extended. The majority of rate changes are $1 up or down which leads to more subtle changes as presented in Section 6.3.

6.2 Arrival-Duration Patterns

The occupancy plots are interesting to study because they demonstrate when and where there is parking congestion

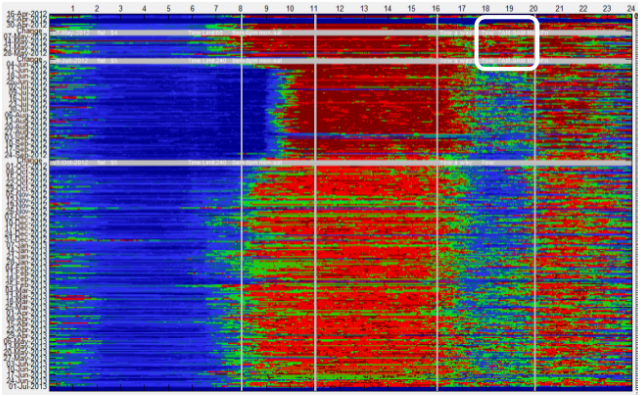

Figure 9: Occupancy data for the 12 spaces at 701 South Olive Street over a year. See text for details.

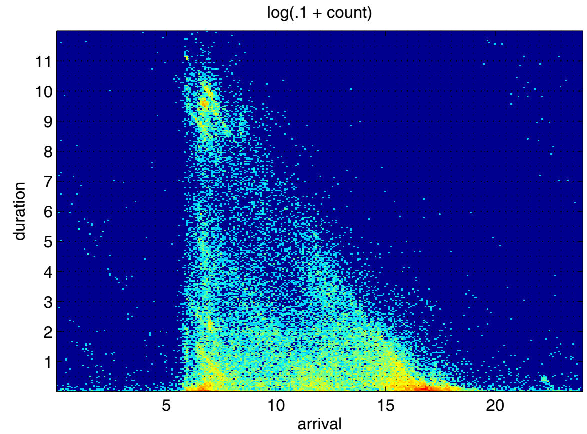

Figure 10: A 2D histogram for arrival duration patterns for 201 North Freemont Avenue. Counts are shown in logs for ease of interpretation. Note the cluster of parkers that arrive around 7 a.m. and stay for around 9 hours.

and motorists are potentially forced to circle to look for parking.

An important complementary view is one that studies arrival and duration patterns. This gives insight in how a block is used and how that use changes after a rate change. Even if occupancy stays the same, the use of the scarce resources can be improved. If for instance in front of a shop long-staying office worker is replaced by short-staying shoppers this increases the efficiency of use: one long stayer walks more, while many short stayers walk less and provide extra custom for local shops. Such a change is not directly observable in occupancy statistics.

Figure 10 shows a 2D histogram of arrival duration data for weekdays during the first year of the project on 201 North Freemont Avenue. Counts are shown in log-space (with a small offset to avoid zeros) for ease of interpretation. This example shows the significant group that arrives around 7 a.m. and parks for around 9 hours. This is longer than the 2 hour time-limit for this block. A study of the payment data confirms that this is a group of non-paying customers as defined in Section 5.1. This is a group that will not change behaviour after a change in rates.

Figure 11 shows for the same block-face and data the average occupancy, and what fractions are contributed by people that park less than 15 minutes, between 15-30 minutes, and so on. This plot shows that the long staying parkers form a significant contribution of the occupancy all throughout the day.

6.3 Analyzing Rate Changes

At the start of the project, enforcement hours (the hours during which parking needs to be paid for) were extended from 6 p.m. to 8 p.m. where there was sufficient demand. Time limits were extended from one hour to two or four hours based on the predominant use of the block.

Based on the algorithm described in Section 3.2 a major rate change was implemented on the first Monday in June 2012 impacting all blocks in the pilot area for which the sensors were successfully installed. Subsequently, rates were updated on every first Monday of the month. City officers used data of special permit use and detailed understanding of the local areas to phase the updates. For every iteration the city officials determined for which areas the recommendations of the pricing engine were put into effect, and which areas "skipped a beat" to wait for more data.

For instance, Time-of-Day (ToD) pricing was introduced in the beginning of August 2012 in two areas of pilot zones (Chinatown, and Fashion District). Then, ToD was progressively applied to other parts of the project area. As of January 2014, there have been a total of nine changes.

Initial results of the first rate change have been reported in [2]. In more cases did demand patterns suggest the rates to go down than to go up (39% of blocks saw a rate decrease, 14% a rate increase), the biggest effect of rate increases in congested areas was around 7pm with a reduction in occupancy of 15%, and the biggest effect of a decrease in rate at 1pm with an average increase of occupancy of 10%.

The absolute numbers are shown for all times of day in Figure 12. We see that a rate reduction seems to effectively draw back parkers in underused areas throughout the day (8am-6pm). Increasing rates leads to a reduction in demand throughout the period 8am-6pm, but the change is subtle. The big effects are in the period from 6pm-8pm. This can be explained from the fact that in several blocks the operating hours were extended from 6pm to 8pm. The two additional hours followed the rate already in use during the day. In many places this meant an increase from 0 to 3 sometimes 4$/hr. This is a strong positive answer to the question "Can rates influence behavior?". An example of this effect was presented in Section 6.1.

6.4 Further Analysis

In the previous sections we have shown examples of the data visualization tools that can be used, and highlighted some of the insights they can give. The first analysis of the impact towards the project goals have been reported by LADOT in [2]. We refer the interested reader to the second report in [3].

In this section we would like to highlight some aspects of this forthcoming work and discuss some of the challenges and lessons learned in the analysis.

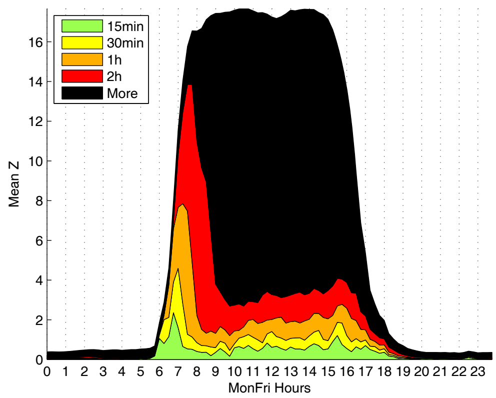

Figure 11: Average occupancy (Z) and contributions of different length parking events as a function of the time of the day. This plot shows 201 North Freemont Avenue and shows that the early arriving and long staying non-paying parkers seen in Figure 10 block a significant portion of the available spaces.

Figure 13 shows the weekly average of B_{cu} for all blocks of the project area and weeks. This is shown as a matrix where rows are block index and columns weeks. The weekly average of B_{cu} is then color coded. The bottom plot of Figure 13 displays a discretized view of the same data, where blue indicates underused (B_{cu} < -1/3), green means "just-right" or as much congested as underused ($-1/3 \leq B_{cu} <$ 1/3), red means congested most of the times ($B_{cu} > 1/3$).

Overall, these plots suggests a decrease in underutilization (blue is less frequent on the right part of the matrix) with a relatively less marked increase in congestion. In addition, it is interesting to mention the seasonality effects. Near weeks 30 and 80, there are visible decrease due to Christmas. Similarly, just before these Christmas events, a less marked effect correspond to Thanksgiving.

There are several lessons we have learned while doing the analysis:

- It is important to report on weekdays and weekends independently. Parkers inherently change their behavior during the week.

- Sensors are susceptible to noise. Both in known and unknown ways. It is important that this is carefully treated to avoid biases.

- Special permit use makes that demand based pricing can only have little impact in significant parts of the project, because the main source of congestion comes from parkers that do not need to pay.

- Events such as the opening or closing of schools, seasonal effects from holidays and overall macro-economic shifts have big impacts. Empirically of at least the same order of a 1$/hr rate change. This means that when a before/after change is observed a further analysis needs to ensure that this is not due to an exogenous effect.

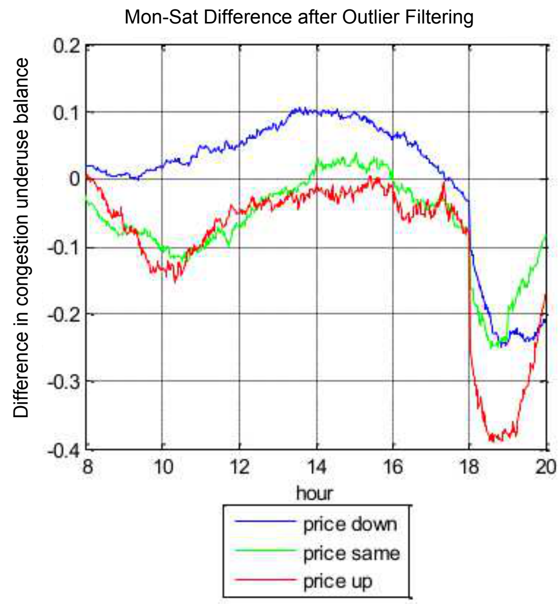

Figure 12: Difference of average B_{cu} before the first rate change (April 1 – June 4, 2012) minus average B_{cu} after rate change (June 4 – Aug 1, 2014).

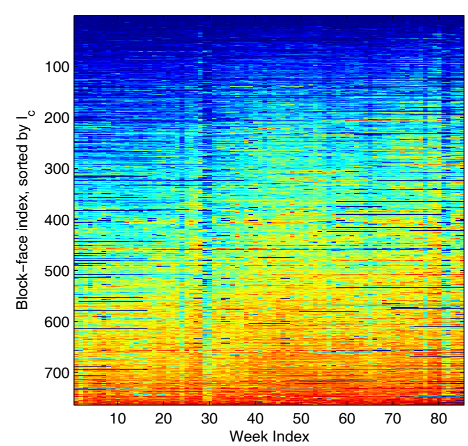

Figure 13: Weekly average of B_{cu} normalized by capacity. Week 1 starts June 4th, 2014.

- Not only the change in occupancy is of importance, but also the change in arrival duration patterns and the fraction of paying customers.

- It is important to understand overflow patterns to off-street and cheaper neighbouring alternatives.

7. DISCUSSION

New sensor technologies and learning algorithms make many new systems possible that can improve city mobility and reduce pollution. The LA Express Park project is a large scale (over 6000 spaces in over 800 blocks), multi-year project that provides a large scale implementation that tests the demand management principles set forth by Vickrey [9] and uses the iteration scheme presented in this paper.

Projects such as these do not only face technological challenges. In fact the political, and organizational challenges are at least as big, and pose important constraints on the type of methods that can be deployed. For instance the methods need to be simple to understand, easily seen to be fair, and lead to pricing policies that are easy for drivers to remember and to act upon. We have experienced that it is hard to elicit these constraints explicitly and any method needs to stand up to many different reasonable definitions of these constraints.

The demand based pricing solution deployed in LA keeps rates at simple to understand discrete values: ($0.5, $1, $1.5, $2, $3, ..., $6). Where demand patterns suggest it is best to charge different rates at different parts of the day, these periods are the same throughout downtown LA. Rates are increased only if congestion is a serious problem, and underuse is not.

Our iteration scheme is based on the fraction of time a block was congested and the fraction of the time it was underused. This "average utility" approach overcomes the basic problem of the "utility of the average" approach proposed earlier. In this method changes are based on average occupancies and can lead to obvious problems: an underused morning and a congested afternoon can lead to an average that is "just-right" even though at no point in the day the parking situation was considered to be so. We see in particular a difference between the two techniques when they are used to select time-of-day windows in a data-driven way. The average occupancy method leads to periods in the day that often exhibit both underuse *and* congestion, something that the idea of charging different rates for different times of the day expressly tries to avoid.

We have not yet seen a reason to change the simple pricing iteration. However, it can be argued that the iteration has not been tested to the fullest due to limits on the rates set by the city council. For the initial year the rates were allowed to increase or decrease by 50% from their starting values. With the pricing ladder used that meant one or at most two opportunities to inrease or decrease rates. When a larger ordinance allows more steps it might be necessary to extend the method to explicitly avoid oscillations between two rates that are nearly equally good. An obvious fix could be to keep longer periods between rate changes if such problems are detected.

Several worthwhile refinements to the basic iteration exist. An important one is to weight the occupancy status by traffic density when computing the congestion and underuse indexes. This requires accurate traffic measurements.

There are ample directions for future work. From a practical point of view, information provision can be improved. Surveys and field studies confirm our hypothesis that an understanding of drivers is a bottleneck [4]. For people to change their behavior they need to care and *know* about discounts. The aim of the techniques presented here is to be in the near future fully integrated with turn-by-turn navigation and reservation systems. Once the construction of a shortlist of options (e.g. closest, cheapest, best matched to preferences) is done by a device, the pricing systems will be more effective, and can also use more complicated rules.

To increase the impact it will also be beneficial to change the way special parking permits work. Currently motorists with a reduced mobility are allowed unlimited free parking in all on-street parking spaces. Empirically this leads to all-day use by a single car for a very high fraction of congested downtown areas. Hence for an important part of the project area smart pricing has a significantly reduced effect.

A careful and complete analysis of parking changes is a complex and challenging task as outlined in Section 6.4. The LA Department of Transportation has reported on the initial impact of the project in [2] and will report on the latest outcomes in [3].

Constructing models that describe parking behavior is an interesting challenge. In contrast to traffic flows on highways, the modeling of parking has received little attention. This is partly because data has never been available, but also because many more environment variables need to be understood and measured. Examples of such variables include the availability and price of off-street alternatives, the number of office spaces and shops in the neighborhood, and the presence of special events. A full understanding of parking behavior is useful in many applications. In the pricing problem studied here it can for instance help to anticipate the effect of changing the rates in neighboring blocks.

8. REFERENCES

[1] S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan. Multi-parameter mechanism design and sequential posted pricing. In *STOC '10*, pages 311–320, 2010.

[2] P. Ghent, D. Mitchell, and A. Sedadi. LA Express Park - curbing downtown congestion through intelligent parking management. In *Proceedings of ITS World Congress*, 2012.

[3] P. Ghent, A. Pudlin, E. Cardenas, S. Clinchant, C. Dance, and O. Zoeter. LA Express Park - Curbing downtown congestion through intelligent parking management. In *Proceedings of ITS World Congress*, 2014.

[4] J. Glasnapp, H. Du, C. Dance, S. Clinchant, A. Pudlin, D. Mitchell, and O. Zoeter. Understanding dynamic pricing for parking in Los Angeles: Survey and ethnographic results. In *HCI International*, 2014.

[5] R. Meir, Y. Chen, and M. Feldman. Efficient parking allocation as online bipartite matching with posted prices. In *Proceedings of AAMAS*, 2013.

[6] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, pages 58–73, 1981.

[7] H. Robbins and S. Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, 22(3):400–407, 1951.

[8] D. Shoup. *The High Cost of Free Parking*. APA Planners Press, 2005.

[9] W. Vickrey. The economizing of curb parking space. In *Traffic Engineering Magazine*, 1954. Reprinted in Journal of Urban Economics 36, (1994), 56–65.

[10] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.

[11] O. Zoeter, C. Dance, M. Grbovic, S. Guo, and G. Bouchard. A general noise resolution model for parking occupancy sensors. In *Proceedings of ITS World Congress*, 2012.