Vanderbilt University

College of Arts & Science

Understanding the Block-Level Price Elasticity of On-Street Parking Demand: A Case Study of San Francisco's SF*park* Project

Adam Shriver

Advised by Professor Malcolm Getz

shriver.adam.j@gmail.com

I greatly appreciate the guidance and thoughtful revisions of Professor Malcolm Getz through every stage of the development of this project. I also thank Professor Joel Rodrigue for numerous stimulating discussions and insightful comments.

1 May 2016

Abstract

This paper explores the the determinants of on-street parking demand in an effort to provide a robust empirical model for estimating block-level demand sensitivity. The SF*park* Project in San Francisco, which has implemented performance-based parking pricing since Aug. 2011, provides a unique setting in which to study these factors. First-order panel vector autoregression, i.e. pVAR(1), is implemented with controls for unobserved heterogeneity and serial correlation. Despite few structural specifications, no clear statistically significant relationships between parking occupancy and price emerge. The paper supports the conclusion that block-level price elasticities vary too widely to be reliably estimated using econometric modeling.

I. Introduction

Parking plays a pivotal role in a city, shaping its spatial layout, facilitating (or hindering) economic growth, and directly impacting the welfare of its residents (Shoup 1997, 2006). Despite its importance, many of the fundamentals of the market for urban parking are little understood by urban planners and transportation economists. In particular, what variables are most influential in determining the demand for on-street parking? These important factors may include location, time of day, day of the week, month of the year, availability of public transit and off-street parking nearby, and land use in the surrounding area. This paper explores which, if any, of these variables impact parking demand with the goal of providing urban planners more reliable information regarding how parking prices should be set to maximize social welfare.

Unlike traditional competitive markets, the market for on-street parking does not have an "invisible hand" to find the price that balances demand and supply. On-street parking is typically owned and controlled by the local government; however, there are instances where local governments contract parking management out to private businesses. In particular, Chicago controversially leased its parking meter system to a private group of investors in 2008 (Preston 2010). Some transportation economists discuss the plausibility of a monopolistically competitive market for on-street parking with many private owners (Anderson and de Palma 2004).

Despite these alternative parking market structures, local government ownership of parking systems prevails in the United States. Private monopoly, as in Chicago, leads to overpriced and under-utilized parking systems due to the profit motive of the owner, which diverges from policymakers' goal of maximizing social welfare. Another approach would be for the city to franchise parking to many small vendors, thus creating a monopolistically competitive market. However, monopolistic competition with many small vendors would be difficult to implement and sustain due to the spatial monopolies generated. If parking were franchised at too high of a level, such as the city's entire downtown financial district, there would be little competition in that particular area due to the resulting spatial monopoly. For competition to work, each parking district or neighborhood must have more parking vendors than would likely be able to collude, which may require an unreasonable number of vendors. Moreover, the small private vendors are unable to take advantage of the economies of scale associated with maintaining parking systems. These economies of scale arise from investment in new pricing software and technology, as well as employing parking patrol officers and keeping up with meter and other technological maintenance. Consequently, this paper specifically studies city-owned parking systems to better understand how urban planners should determine parking prices in order to achieve socially optimal outcomes.

Today, many cities use rigid hourly parking rates. These rates often do not vary by the time of the day, day of the week, or even the location of a given parking space. For example, as of December 2015 the City of Boston charges $1.25 per hour for any on-street parking space between 8 a.m. and 8 p.m., Monday through Saturday, with no charge on Sunday ("Parking Meters"). Frequently, meters stop operating after a specific time, often around 6 p.m., which is a tradition retained from when most stores closed near that time and when meters had one or two hour limits to promote turnover (Pierce and Shoup 2013). As a result, when there is a surge of shoppers on a Friday evening, for example, the increased demand causes parking to be underpriced since prices do not adjust to reflect the change in market conditions.

When the price of parking does not adequately reflect demand, externalities result that negatively affect societal welfare. First consider the case where the price of parking is below the market-clearing price (i.e. parking is underpriced). This may occur if there is a special event, such as a concert or sporting event, that involves many drivers seeking parking spaces within a narrow time interval within a relatively small area. Or, if the parking rate is flat all day, the mid-afternoon may see a surge in demand due to an influx of people going out for lunch or to shop on their lunch breaks. When most spaces are filled, drivers must "cruise" to find an open parking space. Shoup (2006) found that out of 16 studies in large cities, the average percent of traffic cruising for parking was 30 percent and the average time spent cruising was 8.1 minutes. However, the standard deviation of these studies' estimates is large, and Arnott and Inci (2006) point out that the methods used in the individual studies were rudimentary and that the studies are likely biased since researchers looked for cruising in locations where they expected to find cruising. Despite these potential flaws, the studies strongly suggest that cruising is a major consequence of underpriced parking.

Cruising wastes individual drivers' time and resources, adds to congestion, harms the environment, and puts pedestrian and driver safety at risk. In a dense urban area, the cost of cruising for each individual driver in terms of his or her time and money spent on fuel is relatively insignificant. But aggregating this cost over all drivers in that area can give surprising results. In a manner similar to Shoup (2006), consider a situation where the average time spent cruising is 4 minutes. If 10 cars park in a particular space in a given day and there are 10 spaces on a block, then that block generates 400 minutes of cruising per day. Now, if the average speed while cruising is 10 mph, then over a year that block generates about 2,433 *hours* of cruising and 24,333 additional Vehicle Miles Traveled (VMT). Arnott and Inci (2006) argue that a car cruising for parking adds more to congestion than a car in transit. A cruising car is typically moving slower than the throughput traffic, and the interaction between slow- and relatively fast-moving cars causes more congestion than if all cars were travelling at the same speed (Verhoef, Rouwendal, and Rietveld 1999). This affects not just automobile drivers, but also users of public transit who are negatively impacted by slow-moving traffic. Besides adding to congestion, the additional VMT increase emissions, resulting in lower air-quality and other potentially negative long-term environmental effects. Finally, a little studied topic is cruising's effect on pedestrian and driver safety. Drivers cruising for parking are often distracted while looking for an open space, which puts both pedestrians and drivers at risk of collisions. The SF*park* Pilot Project Evaluation (2014) cites safety improvement as one of the parking Project's secondary benefits, working through the mechanism that fewer distracted drivers cruising for parking reduces collisions.

The other case of mispriced parking to consider is overpriced parking. Overpriced parking occurs when rigid prices do not adjust downwards in order to compensate for periods of low demand. Pierce and Shoup (2013) argue that overpriced parking causes fewer drivers to demand parking services than the socially desirable quantity. The consequence is detrimental to the economic vitality of the surrounding area, since nearby businesses lose customers and hire fewer employees. Moreover, the local government loses the opportunity to collect revenue on those empty spaces.

An intriguing question arises regarding overpriced parking: would it ever be socially desirable for parking to be overpriced? Suppose that the local government has the primary objective of reducing traffic and that the implementation of roadway congestion pricing (i.e. tolls) is not feasible. Since raising the price of parking can discourage automobile travel, the local government may choose to raise parking prices beyond the equilibrium price in order to reduce the number of automobile trips and consequently the amount of traffic in a particular area. Compared with road pricing, parking pricing is comparatively easy to put into effect from an administrative standpoint (parking pricing has fewer privacy concerns and does not require additional legislation) and relatively inexpensive to implement since most cities have some type of parking pricing infrastructure in place (Litman 2015). However, raising parking prices beyond their equilibrium level to reduce traffic is still a second-best strategy for traffic control to road pricing because it affects individual behavior indirectly and therefore does not affect through traffic.

Given the presence of these externalities associated with the parking market, it is desirable for the urban planner to set the price of parking such that marginal social benefits are equated with marginal social costs (Arnott and Inci 2006). This is a basic problem from microeconomics, but the solution is difficult to find and implement in practice. While supply is easy to measure (simply count the number of spaces) in the short-run, what should the long-run socially optimal number of parking spaces be? This is a tricky question, since this quantity should, to some extent, depend on the level of societal demand for parking. However, parking demand is extremely difficult to measure directly for practical considerations (Shoup 1995). In particular, for a particular price the quantity demanded of on-street parking on a given block reflects how many vehicles would park on that block if there were sufficient supply (Ottosson et al. 2013). But the "sufficient supply" assumption is often not met in practice when blocks are full. This raises the following question: how can economists and urban planners determine if a particular driver cruising for parking would actually park on a fully occupied block if there were open spaces at the quoted price? This issue is layered on top of the practical difficulty of distinguishing between throughput traffic and cars cruising for parking. Given these difficulties, urban planners sometimes even make policy decisions (in particular, set minimum parking requirements) by assuming that all parking is free and then roughly estimating demand based on land use (Shoup 1995).

Another problem with estimating demand accurately is that the demand for on-street parking is stochastic. Unpredictable random variables, such as the weather, may impact how many drivers are seeking parking at a given point in time. On the other hand, urban planners can try to forecast how many people will attend a non-random special event, such as a concert or a sporting event, but their forecasts are simply estimates since the number of additional vehicles seeking parking cannot be known with certainty ahead of time. Moreover, in the case of a special event, there may be a "crowding out" effect on other vehicle trips; for example, if a person wants to go downtown to shop on a Sunday afternoon but there is an NFL football game that night taking place downtown, that person may choose to delay his or her trip to another day. This means that the effect of a special event on parking demand is not purely additive. The extent of this crowding out is uncertain and cannot be perfectly predicted beforehand. As a result, the randomness associated with parking demand adds another layer of difficulty for determining prices.

In light of this uncertainty in demand measurement, transportation economists have advocated "performance-based pricing," also known as "demand-responsive pricing" (Ottosson et al. 2013; Pierce and Shoup 2013). First define the occupancy rate of a block to be the percent of parking spaces on that block that are occupied at a given point in time. An average occupancy rate over an interval of time can easily be computed from a set of time-specific occupancy rates. Typically, the target occupancy range is around 60 to 80 percent, as is the case in San Francisco, so that there should always be one or two spaces open on every block. Other cities have also implemented performance-based pricing using cutoffs near this range, such as Seattle, Los Angeles, and Washington D.C. Some economists advocate a target occupancy of 85 percent (Weinberger, Kaehny, and Rufo 2010). Under the performance-based pricing régime, urban planners choose a target occupancy rate per block (typically a range of rates) and adjust prices up or down until the target occupancy is reached. Ottosson et al. (2013) justify the use of occupancy rates to adjust prices by arguing that occupancy rates are highly related to parking demand and that occupancy rates can be used as an estimator of parking demand. Performance-based pricing is appealing to planners because it renders directly estimating parking demand unnecessary since occupancy rates are used to adjust prices.

Determining the optimal target occupancy range is difficult. If the range were 90 to 95 percent occupancy, blocks would often be completely full during peak periods of demand. In addition, small stochastic fluctuations in demand would cause blocks to become entirely occupied, such as a small surge in demand caused by a large party having dinner at a restaurant.

On the other hand, the range should not be set too low (say at 30 to 50 percent) because nearby businesses would lose customer traffic and some cruising would simply be shifted to surrounding areas. The key to determining the appropriate target occupancy or occupancy range is to balance the aforementioned tradeoffs between having overcrowded blocks with high targets and under-occupied blocks with low targets. This is why the target range is typically near 60 to 80 percent, ensuring that one or two spaces are always open on each block.

However, several additional issues regarding the implementation have yet to be considered. These include how often prices should be changed, as well as at what increments prices should be changed in order to reach the target occupancy range. Pierce and Shoup (2013) posit that narrower intervals improve efficiency and ensure that target occupancy rates are met more frequently. However, in order to answer the latter question regarding price adjustment magnitudes, it is essential to be able to predict how occupancy rates will change in response to a given change in price. For a given percentage change in price, what is the corresponding percentage change in average occupancy rates? In economics terminology, it is necessary to understand the price elasticity of on-street parking demand. If this measure is known, urban planners can use it to adjust prices more accurately and, consequently, find the "right" price of parking quicker (Ottosson et al. 2013).

There is strong empirical evidence that the price elasticity of parking demand varies widely at the block level (Ottosson et al. 2013; Pierce and Shoup 2013). Key potential determinants of the price elasticity of parking demand for a particular block include the time of the day, day of the week, month of the year, location relative to driver's destination, land use in surrounding neighborhood, and availability of public transit and off-street parking. In order to understand how the demand sensitivity to price varies by block, this paper presents and tests a model based on some of these variables using data from the SF*park* Project. The particular variable of interest in this paper is the location relative to the driver's destination. To my knowledge, this paper represents the first attempt to systematically model how distance between a block and drivers' destinations within a neighborhood affects the price elasticity of demand.

The SF*park* Project is a cutting-edge project that has utilized performance-based pricing to implement multiple block-level price changes since August 2011. The primary objective of the Project is to improve parking availability, making it easier for drivers to find parking spaces at all times of the day during each day of the week on any given block (SFMTA 2014). The San Francisco Municipal Transportation Agency also cites reduction of cruising-related externalities as secondary benefits of the Project (SFMTA 2014). Seven parking management districts (PMD), i.e. neighborhoods, with varying characteristics are included in the Project, as well as over a dozen city-managed parking garages. The Project utilizes new technologies to facilitate highly accurate data collection on occupancy rates, overcoming the need to impute occupancy from payment data. These data are made publically available online in an effort to promote transparency and stimulate research. This is an ideal setting for decomposing the price elasticity of parking demand by its principle determinants and, in particular, modeling the spatial variation of demand sensitivity to price.

II. Literature Review

As noted previously, the economics of urban parking literature is thin. There have been several empirical studies of how employer parking-pricing policies affect commuter mode choice in the United States (Shoup 1997; Washbrook, Haider, and Jaccard 2006; Hamre and Buehler 2014; Ommeren and Russo 2014). An increase in the price of parking is an increase in the cost of individual automobile travel. Essentially, these studies are therefore estimating how substitutable commuters find single-occupancy automobile vehicle travel and carpooling, public transit, biking, or walking to work. Another notable study by Donald Shoup (1999) studies minimum parking requirements and analyzes their effects on land use, congestion, and the environment.

Several theoretical modeling attempts have been made to understand various aspects of the parking market. Arnott and Inci (2006) model the interaction between cars cruising for parking and cars in transit. Arnott and Rowse (2009) present an expanded model that incorporates the availability of off-street parking, concluding that local governments should raise on-street parking rates in order to effect efficiency gains and raise revenue. There have also been modeling exercises that capture stochastic variation in parking demand (Arnott and Rowse 1999; Qian and Rajagopal 2014). Anderson and de Palma (2004) consider a long, narrow city with parallel access roads, and Anderson and de Palma (2007) use a monocentric model with discrete zones based on the distance from the center. Both of these attempts incorporate distance from the center of the city as a key variable in the analysis.

However, these theoretical studies often run into problems due to lack of reliable empirical information regarding key parameters in their models. For example, Arnott and Inci (2006) take for granted -0.2 as an estimate for the price elasticity of travel demand during peak periods. This parameter value is then used in the remainder of the numerical analysis without the authors justifying its use beyond positing that it is a “standard estimate” (without reference). Anderson and de Palma (2004) do not mention the price elasticity of demand for parking, nor its potential spatial variation. However, these authors’ assumptions regarding price elasticities of parking are not a sign of a lack of rigor; instead, they indicate the dire need for additional empirical evidence regarding the determinants of parking demand sensitivity.

There is a limited body of existing literature on price elasticities of parking demand, although relatively recent technological advancements have spurred transportation economists to begin to more closely study this topic. Even within the last decade, researchers have used meter payment data to estimate how many cars parked in each space during a day, but since not all drivers comply to payment requirements, this simply gives an estimate of occupancy. Now, in-ground parking sensors can be used to precisely determine how many hours of the day any given parking space is occupied. Moreover, "smart meters" facilitate payment and ease of use for drivers, consequently increasing payment compliance rates. When compliance rates are higher, pricing strategies are more effective. If a driver is going to park in a space but not pay regardless of the price, no pricing structure will influence that driver's behavior. These new technologies aid parking meter enforcement, reducing the cost of hiring Parking Control Officers (SFMTA 2014). Finally, the real-time data provided by these instruments is time and location specific, which enables urban planners and transportation economists to perform rigorous analysis and make data-driven decisions (SFMTA 2014).

Outside the United States, Kelly and Clinch (2009) study the block-level price elasticity of demand of on-street parking in the central district of Dublin, Ireland, and derive an average estimate of -0.29. The same study also suggests temporal variation in price elasticities, with demand being more sensitive in the morning and less sensitive in the afternoon and at night. However, a shortcoming of this study is that it does not report standard errors for its price elasticity of demand estimates, which are computed using the arc elasticity formula and adjusted for changes in average income over time. The lack of reported standard errors hides potentially large block-level variability in elasticities.

To my knowledge, the first block-level study of the price elasticity of demand of on-street parking in the United States was conducted by Ottosson et al. (2013) using data from a single, neighborhood-level price adjustment near Downtown Seattle. Each neighborhood contained multiple blocks. The price adjustment was determined using a combination of the neighborhoods' previous peak period occupancy and average occupancy rates, with target occupancy between 71 and 86 percent. Block elasticities varied based on the time of day, with demand being most elastic in the morning and most inelastic mid-day. The study finds also that as a block's distance increases from the core of the Central Business District (CBD), parking demand becomes more sensitive. In neighborhoods with price increases, public transit availability had a significant negative effect on occupancy rates. Surprisingly, the price of off-street parking did not have a significant effect on occupancy rates.

Several problems arise in the Ottosson et al. (2013) analysis. In particular, only one price change was studied. All analysis, as a result, was conditioned on whether a particular block had a price increase, decrease, or no change. This prevents direct, reliable comparison (without additional assumptions) between a block that had a price increase and another that had a decrease. Moreover, occupancy rates were estimated based on meter payment data. This carries the implicit assumption that payment compliance is uniform over space. If a certain block is more heavily patrolled by Payment Compliance Officers, then compliance should be higher on that block. This would cause the price change to have a larger effect on the high-compliance block than a similar price change on a low-compliance block. If payment compliance varies in some systematic way, especially if it is correlated with distance from the center of the CBD (which is a likely possibility), then the analysis cannot definitively attribute price elasticity variation to distance from the urban core.

In the same year as Ottosson et al. (2013), Pierce and Shoup utilized data from the SF*park* Project to study the block-level price elasticities of on-street parking demand. The SF*park* dataset has a distinct advantage over the Seattle data used in Ottosson et al.: SF*park* tracks many performance-based price changes, occurring approximately every two to three months, as well as the corresponding changes in block-level occupancy. Additionally, SF*park* uses in-ground sensors to directly observe occupancy rates, rather than imputing occupancy rates based on meter payment (SFMTA 2014). The study found that price elasticities varied by neighborhood, time of day, day of the week, initial price, magnitude of price change, and date of price change. As a result of this observation, Pierce and Shoup conclude that "planners will never be able to develop a robust theoretical model to predict the correct prices needed to achieve the target occupancy for every block" (2013). However, a considerable factor undermining this conclusion is the study's lack of analytical rigor. Pierce and Shoup simply observe that elasticities vary greatly and are related to many variables, and they jump to the conclusion that there is not a systematic way to model how these variables interact to determine price elasticities. Furthermore, they do not address the availability of public transit or off-street parking prices, as well as location of blocks relative to likely driver destinations (i.e. block distance from center of neighborhood).

The above commentary reveals a discrepancy in the existing literature. Ottosson et al. (2013) maintain that transportation economists can rigorously model parking demand sensitivity to predict future occupancy rates based on price changes, while Pierce and Shoup (2013) claim that this prediction process is not possible since block-level elasticities vary so greatly. Instead, Pierce and Shoup advocate a trial-and-error process for determining prices that achieve occupancy goals. This paper seeks to address this disagreement in the literature.

III. Methodology

Study Setting

The City of San Francisco has experienced considerable growth moving out of the Great Recession in terms of population and economic viability. From 2011 to 2013, San Francisco's citywide population grew an average of 1% per year with average growth in real metropolitan GDP of 2% per year (author's calculation) during the same time period (SFMTA 2014; BEA 2015). Growth in real metropolitan GDP skyrocketed in 2014, reaching 5.2% and more than doubling the national average of 2.3% (BEA 2015).

These demographic and economic changes in San Francisco have had interesting influences on transportation demand. As of 2014, San Francisco-Oakland is ranked by researchers from The Texas A&M Transport Institute and INRIX, Inc. to be the third most congested metropolitan area in the United States, most notably *ahead* of New York, Boston, and Chicago (Schrank et al. 2014). Vehicle miles traveled to and from San Francisco decreased by approximately 5% per year from 2011 to 2013; weekday BART ridership (the heavy-rail system servicing the San Francisco Bay Area) to and from San Francisco increased by 7% per year over the same time period (SFMTA 2014). These data suggest that commuters from outside San Francisco have shifted away from driving automobiles to work in favor of utilizing public transit. However, it is surprising that automobile ownership in San Francisco increased one-to-one with citywide population from 2011 to 2013, adding 16,000 cars to the city (SFMTA 2014).

A possible explanation of this phenomenon is that wealthier individuals, who are able to afford the extraordinarily expensive housing in the heart of San Francisco, moved into the city during this time period, bringing their cars with them. Then, the increased congestion and demand for parking within San Francisco may have discouraged workers from outside the city from commuting by car, thus shifting their demand to public transit systems. This becomes a plausible story when one considers the growth of the technology industry in Silicon Valley. Many of the recent employees at growing tech companies, such as Facebook, are young and mobile, choosing to live outside the city limits in order to avoid monumental housing costs within the city. These young, reasonably wealthy commuters often utilize public transit services, such as BART, rather than owning cars in order to avoid the headaches of traffic and parking in the city. High traffic and parking costs have induced some large tech companies to offer their workers subsidies to relocate closer to the city. Notably, Facebook has begun to offer a $10,000 or more subsidy to workers who relocate near the Facebook headquarters in South San Francisco (Lam 2015).

As discussed in the Introduction, cruising is a major problem associated with mispriced urban parking. When parking access is not readily available, drivers waste time and money searching for open spaces. This increases congestion, harms the environment, and poses safety risks to drivers and pedestrians. In response to worsening traffic conditions and decreasing parking availability in recent years, the San Francisco Municipal Transportation Agency (SFMTA) launched the SF*park* Project in August 2011. Under this unique project, the SFMTA implemented performance-based pricing for on-street parking in 7 Parking Management Districts (PMDs), i.e. neighborhoods, and for off-street parking in 14 garages (SFMTA 2014). The SFMTA operates San Francisco's public transit system, roads, on-street parking, and a significant amount of its off-street parking. The primary goal of this project was to improve parking availability, with secondary benefits being reduced cruising and its associated externalities (SFMTA 2014).

As a part of the project, the SFMTA developed and tested cutting-edge technologies. StreetSmart Technologies, now known as Fybr, installed and maintained in-ground sensors starting in September 2010 (SFMTA 2014). These parking sensors provide real-time occupancy data on individual spaces by detecting the presence of a parked car regardless of whether or not the meter is paid. This overcomes one of the major difficulties in the Ottosson et al. (2013) study that relied on meter payment data to estimate occupancy. New "smart" parking meters were also installed for the Project from July 2011 to January 2012 in an effort to facilitate payment and increase payment compliance, as well as enable wireless communications and updates that allowed for programmed price changes according to occupancy rates (SFMTA 2014). These meters accept credit cards, as opposed to traditional coin- and cash-only meters. The wireless communication of these new technologies also laid the groundwork for the creation of an SF*park* mobile app. The app uses real-time data to inform drivers seeking parking where there are open spaces and what the hourly rates are at those spaces. It has become widely-used, with 73,500 downloads from iPhone and Android users as of July 2013 (SFMTA 2014). The data collected by these technologies is available on the SF*park*.org website for public use, as well as the source code for their software. This is part of an organized effort to promote data-driven parking management decisions and encourage novel research to inform policy.

The SF*park* Project is an ideal setting for evaluating performance-based pricing. Parking prices vary at the block level by time of day and weekday/weekend. Each day is segmented into three time bands: morning (open to 12 p.m., where opening time varies by PMD and is usually around 9 p.m.), afternoon (12 p.m. to 3 p.m.), and evening (3 p.m. to close, where closing time varies by PMD and is usually around 6 p.m.). Parking is free overnight, i.e. from closing to opening time. As of January 2016, there are two classes of days: weekdays (Monday through Friday) and weekend days (Saturday and Sunday).

Prices are updated approximately every two months in response to observed average occupancy rates during the previous two-month period. The target occupancy range is 60 to 80 percent. Average occupancy over 80 percent indicates a $0.25 hourly increase; between 60 and 80 percent, no change; between 30 and 60 percent, $0.25 hourly decrease; and less than 30 percent, $0.50 hourly decrease. During the day, rates cannot fall below $0.25 or exceed $6.00 per hour.

To illustrate this pricing system, consider Beach Street 500 in Fisherman's Wharf from 12:00 to 3:00 p.m. on the weekend. From October 2011 to December 2011, the price of parking on this street was $3.25 and the average occupancy was 84%. This indicates a $0.25 hourly price increase on this street from 12:00 to 3:00 p.m. on weekends, resulting in a new price of $3.50 per hour. During the next period, from December 2011 to February 2012, the occupancy declined slightly to 83%. This indicates another incremental increase of $0.25 per hour, giving a new price of $3.75 per hour. This trial-and-error process is designed so that prices adjust according to observed occupancy during the previous period until target occupancy goals are reached. Ideally, prices will be adjusted such that target occupancy will eventually be met every day, in every time band, and on every block.

To perform a qualitative review of the effectiveness of the performance-based pricing system, the SFMTA conducted a pilot project of SF*park* from late 2011 to mid-2013, making a total of ten price adjustments. Performance-based pricing as outlined above was instituted in seven “pilot” Parking Management Districts with varying characteristics and geographic locations. During this time period, two control Parking Management Districts with no price changes were tracked. These control PMDs were added in order to mitigate the influence of possible confounding factors, such as city growth, increased transportation demand, and other demographic shifts. Before undergoing the pilot project, payment compliance was measured on each block by comparing in-ground sensor data with meter payment data. Assuming payment compliance remains constant over time, these data can be used to analyze the relative efficacy of price changes in high payment versus low payment neighborhoods. A high payment block, or “HP block”, is one in which 85 percent or more of the hours parked on that block were paid for (SFMTA 2014).

As expected, qualitative analysis has shown that the performance-based pricing was highly effective over that time period. First let a “full” block be a block with, on average, 90 to 100 percent of spaces occupied. Pilot HP blocks were full 45 percent less often than before the project began, and other pilot blocks were full 16 percent less often. In contrast, control blocks were full 51 percent more often (SFMTA 2014). This statistic also suggests that the pricing régime was more effective in high payment blocks. The pilot blocks performed better than control blocks on a number of categories, such as average time spent cruising, greenhouse gas emissions, number of citations, congestion, and traffic volume and speed. Pilot HP blocks also experienced greater improvements than other pilot blocks in percent of time meeting target occupancy goals and percent of time at full occupancy (SFMTA 2014).

The above qualitative considerations strongly suggest that SF*park* was effective during the pilot project. This has led the SFMTA to continue the SF*park* Project, which continues to be in effect as of January 2016. Nevertheless, several aspects of SF*park* require improvement.

Pierce and Shoup (2013) argue that SF*park* should refine its time periods and take a proactive approach in predicting future occupancy rather than implementing price changes based solely on the previous period's occupancy. Narrower time bands, as well as later meter hours, would improve the amount of time that target occupancy is met. Furthermore, seasonal effects should be considered when adjusting prices. Just because occupancy rates are high in December due to the holiday and shopping season does not indicate that parking prices should be raised in January. In general, a more thorough understanding of the determinants of parking demand would allow planners to estimate future demand by means other than simply relying on past occupancy rates.

In order to accomplish the shift from reaction to prediction, a more developed investigation of the price elasticity of on-street parking demand is necessary. If planners do not understand how sensitive parking demand is to price, how can they possibly predict the hourly price required to reach the desired target occupancy?

Data Set

The data set is arranged as panel time-series (longitudinal) data. Each city block in SF*park* constitutes a panel, and each price change corresponds to a time period. There are N=256 blocks in the SF*park* data set, and there have been T=17 price changes as of January 2016. I refer to an "observation" as all data measures associated with block i at time t. SF*park* contributes 12 measurements per observation: price per hour during the period in the morning, afternoon, and evening, as well as average occupancy during the period in the morning, afternoon, and evening. These measures are supplied for both weekdays (Monday through Friday) and weekends (Saturday and Sunday). I use the term "morning" to designate opening time to noon, "afternoon" to indicate noon to 3 p.m., and "evening" to denote 3 p.m. to 6 p.m.

Some blocks have a fourth "night" time period that runs from 6 p.m. to close, but I exclude these night periods from the sample due to the relatively small number of measurements that fall into this category (254 out of 4,353). Observations begin in August 2011 and end in December 2015. Note that weekend observations before 2013 only represent Saturday price and occupancy data because San Francisco only began to charge parking fees for Sunday parking during the first weekend of January 2013 (Baker 2013).

While each time period on average lasts approximately three months, there is considerable variation in the length of time between price changes. The shortest period lasted only one month, while the longest period lasted almost a full year (10 months). In stark contrast to the transparent, rule-oriented pricing rules, the rules about when price changes are implemented are unclear. This creates a potential source of endogeneity for the econometrician: are prices changed when urban planners begin to account for potential future demand influences? For example, suppose BART needs major repairs, and construction will begin in six months and last for up to a year. Such an occurrence may reduce the accessibility of public transit, thus shifting mode choice toward automobile travel and increasing the demand for parking. Expecting this change in transportation demand, urban planners may decide to implement a price change just as construction begins and just after the increased surge in demand, say in 6 months and 7 months from today, respectively. The first price change "resets" the average occupancy measurement so that the increased occupancy during the 6 to 7-month time frame is not diluted by lower average occupancy rates before the construction. Therefore, the average occupancy rate over the 6 to 7-month period fully reflects the increased demand for parking induced by the construction. This example serves to illustrate that the choice of time periods is likely not random in the SF*park* Project, and as a result, each price change not only reflects the previous period's average occupancy but also planners' expectations of future parking demand. Despite this observation, the remainder of this paper will treat the time periods as having been imposed deterministically, without regard to expectations of future conditions. Without further information about how time periods are selected, this assumption is necessary for data analysis. The results will be interpreted with caution given this underlying assumption. It is appropriate to point out that the average time period length without the 10-month outlier is still 3 months.

In addition to the data provided by the SFMTA, I have gathered additional data pertaining to distance in order to perform analysis. I estimated the physical center of each of the 7 PMDs as they are laid out in the "SF*park*: Pilot Project Evaluation" but with two key exceptions: the bordering Downtown and South Embarcadero districts. I reallocated the blocks in the southern portion of the Downtown district to the South Embarcadero district to force land usage, distance, and parking district to better coincide. The center of the modified Downtown district is determined based on borders that more closely resemble the Financial District displayed on Google Maps, and the center of the modified South Embarcadero district is actually closer to many of the reallocated blocks than the center of the modified Downtown district is. The original and modified Downtown and South Embarcadero districts are depicted in Figure 1.

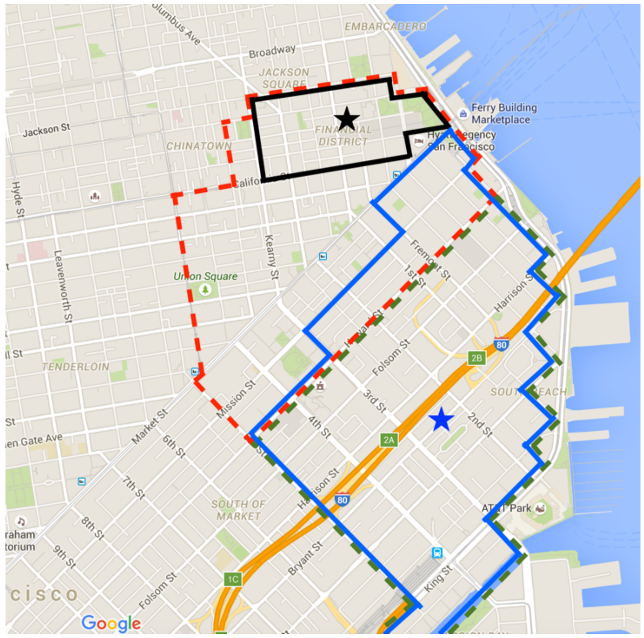

Figure 1. Modified Districts: Downtown and South Embarcadero
Notes: The red and green dashed lines encompass the unmodified financial district and South Embarcadero (SE) district, respectively. The black and blue lines encompass the modified financial district and SE district, respectively, with their estimated centers denoted by stars.

The "center" of each PMD was determined as the centroid of the physical area encompassed by the blocks in the neighborhood. Suppose each block is a point of mass (each with uniform density) in the two-dimensional plane. Averaging the horizontal and vertical measurements gives the coordinate of the centroid of the neighborhood. Determining the center of each PMD allows for the measurement of the network grid distance from the midpoint of each block to the center of the PMD. These distances were measured using Google Maps, and the distance was recorded as the shortest walking distance in kilometers. One distance measurement was taken for each of the 256 blocks. All distances are greater than zero.

Public transit is not included in this data set. Visual inspection of the PMDs using Google Maps suggests that public transit stops are (reasonably) uniformly spread over space within the districts. Many of the blocks have well over 20 public transit stops within a 500m radius. Due to the lack of variation in public transit accessibility points, this variable is assumed to have a uniform effect on all blocks. Additionally, the price of off-street parking is not included in the analysis since Ottosson et al. find it to have no effect on on-street occupancy rates. If these assumptions do not hold and the number of transit accessibility stops or the price or availability of off-street parking are correlated with distance from the center, then the results of the regression analysis will be biased due to omitted variables. Potential bias is further addressed in the Discussion.

Model

With many panels and relatively few time periods, the data are most readily analyzed with micro-panel techniques ("Panel time-series modeling"). Specifically, I utilize a first-order panel vector auto-regression model, i.e. pVAR(1), implemented with Stata programs created by Abrigo and Love (2015). The first-order model was selected based on model-fit criteria proposed by Andrews and Lu (2001). Appendix B.1 describes the selection process of the lag-order of the pVAR model in detail. The model can be represented as

$$\vec{y}_{i,t} = A\vec{y}_{i,t-1} + B\,X_{i,t} + \vec{u}_i + \vec{e}_{i,t} \tag{1}$$

where $\vec{y}_{i,t}$ is a 6x1 vector of the natural logarithm of average occupancy rates (expressed as percentages, all of which are greater than 0 and less than or equal to 100) at different times of the day on weekdays and weekends on block *i* during time period *t*. A is a 6x6 matrix of time-invariant parameters that are estimated, and $\vec{y}_{i,t-1}$ is a 6x1 vector of the natural logarithm of lagged average occupancy rates; $\vec{u}_i$ is a 6x1 vector of dependent variable-specific fixed-effects errors; and $\vec{e}_{i,t}$ is a 6x1 vector of idiosyncratic error terms.

$X_{i,t}$ is a time-variant column vector of exogenous variables that can be expanded as follows

$$\begin{bmatrix} t \\ \vec{s}_t \\ \vec{D}(\Delta x_{i,t}) \\ dist_i \, \vec{x}_{i,t} \\ D_{1,i}(pmd) \, \vec{x}_{i,t} \\ \vdots \\ D_{5,i}(pmd) \, \vec{x}_{i,t} \end{bmatrix} \quad (2)$$

where t is a linear time trend and $\vec{s}_t$ is a 4x1 vector representing the "weighted-average season." The linear time trend controls for linear changes in on-street parking demand over time that impact all blocks. Since the time periods are irregular, controlling for seasonal effects is not a straight-forward task. To deal with this issue, I let March through May be spring months; June through August be summer months; September through November be fall months; and December through February be winter months. If a time period spans multiple seasons, then I count the number of months of the period that fall in each season to allocate the time period across the season (by proportion). The sum of the seasonal terms must therefore equal 1.

Next, $\vec{x}_{i,t}$ is a 6x1 vector of the natural logarithm of the hourly price of parking at different times of the day on weekdays and weekends on block *i* during time period *t*. $\vec{D}(\Delta x_{i,t})$ is a 6x1 vector of dummy variables defined as follows

$$D_j(\Delta x_{i,t}) = \begin{cases} 1, & x_{i,t} - x_{i,t-1} \neq 0 \\ 0, & x_{i,t} - x_{i,t-1} = 0 \end{cases} \qquad j = 1,\ldots,6 \quad (3)$$

In plain words, $D_j(\Delta x_{i,t})$ represents whether or not there was a change in price from *t-1* to *t* on block *i* and is calculated for each time zone (i.e. time of day, day of week). Following Ottosson et al. (2013), the inclusion of this term establishes a control group—a group of blocks with constant price between over the time interval—against which blocks with price changes can be compared.

$dist_i\ \vec{x}_{i,t}$ is simply the distance of block *i* from the center of its associated PMD (in km) multiplied by the price vector, $\vec{x}_{i,t}$. The estimated coefficient on $dist_i\ \vec{x}_{i,t}$ represents the change in price elasticity associated with a 1 kilometer increase in distance from the center of the PMD.

$D_{k,i}(pmd)$ equals 1 if block *i* belongs to PMD *k* and is 0 otherwise, where *k*=1,…,5. $D_{k,i}(pmd)\ \vec{x}_{i,t}$ represents the interaction between PMD and the price vector. The estimated coefficient on $D_{k,i}(pmd)\ \vec{x}_{i,t}$ represents the price elasticity of on-street parking demand for PMD k if $dist_i$=0.

Here I note that *k*=1,…,5 since the preferred model excludes the Marina and Mission PMDs due to numerical considerations. The omission of these two PMDs leaves 5 remaining PMDs for estimation, which have a total of 209 blocks. The decision to exclude these neighborhoods does not have a significant impact on any of the main results, but the elasticity estimates break down for these particular neighborhoods due to large estimates (in magnitude) and large standard deviations on their $D_{k,i}(pmd)\ \vec{x}_{i,t}$ slope coefficients. This issue and potential causes will be further addressed in the Model Results subsection.

An assumption in the analysis is that $dist_i$ is a time-invariant scalar. A case in which this assumption may not hold is if new commercial developments are implemented asymmetrically on one edge of a PMD over time during the SF*park* Project. If the SFMTA installs additional meters near the development, then the regression analysis could be performed before and after the natural experiment with an updated center utilized after the change. This paper does not take this approach since to my knowledge, no such information regarding meters installed due to commercial development is available. If asymmetric development occurs but the SFMTA does not install additional meters, then the "real" center of the neighborhood may shift despite the observed center (based on SF*park* block inclusion data) not moving. However, this fundamental alteration in the system would not be reflected in the data. This would lead to biased parameter estimates and would obscure the role of distance in the system.

It is important to notice that the model does not include interactions between distance and PMD. In other words, I assume that the effect of distance is constant across neighborhoods. This assumption is imposed for practical purposes. Initial regressions were performed including this interaction, but the number of parameters in the model was so large that reliable parameter estimates were not attainable. As more data from SF*park* become available, this may be a fruitful relationship to further explore.

In the absence of readily available natural experiments, such as some supply shock or some other externally induced change to the environment, including the lagged dependent variable as a regressor is often the best approach to dealing with the possible joint determination of price and quantity that often presents a barrier to this type of regression analysis. Under the assumptions that price changes are entirely determined by the previous period's occupancy and that their timing is independent of expectations of future market conditions, it is valid to include the price vector $\vec{x}_{i,t}$ as exogenous regressors after "conditioning" on lagged occupancies $\vec{y}_{i,t-1}$. Ottosson et al. (2013) approaches this issue in a similar way, using a first-order autoregressive, i.e. AR(1), model to account for possible endogeneity. They posit that the inclusion of lagged occupancy as a regressor ensures that "any measured influence of the interest … price elasticity in our case … is conditioned on this history and represents the effect of the new information (price change)" (Ottosson et al. 2013).

While the inclusion of lagged dependent terms on the right-hand side of System 1 reduces concerns of joint-determination of occupancy and price, it presents the need for sophisticated econometric tools to mitigate bias in estimating system parameters. It has been shown that significant bias still remains when using an equation-by-equation OLS approach to estimating a system of equations with lagged dependent variables used as regressors, even with large N (Nickell 1981) or T as large as 30 (Judson and Owen 1999). Due to this consideration, I jointly estimate System (1) using the pVAR package of programs in Stata in order to avoid bias in my results (Abrigo and Love 2015).

Dependent variable-specific fixed effects errors, $\vec{u}_i$, are removed using forward orthogonal deviations (FOD). Under this method, the mean of all future observations in the block is subtracted from the observation at time *t* (Baum 2013). I utilize this method in order to preserve as large of a sample size as possible with the goal of improving estimation reliability. The data set is unbalanced (contains gaps), so a method such as taking first differences (FD) magnifies the gaps in the data and further reduces the sample size. Utilizing forward orthogonal deviations dramatically reduces this loss of data while effectively controlling for unobserved heterogeneity at the block level and eliminating the omitted-variable bias associated with $\vec{u}_i$. Sensitivity analysis shows that parameter estimates are sensitive to choice of FOD or FD. The preferred model utilizes FOD for the reasons cited above.

The first four lags of the dependent variables are used as moment conditions in estimating the pVAR(1) model (see Appendix B.1 for details on the selection of moment conditions). Since the panels are unbalanced, I replace missing observations with zero, as recommended by Holtz-Eakin, Newey, and Rosen (1988). This substitution is valid under the standard assumption that the instruments are uncorrelated with the error terms (Abrigo and Love 2015). The econometric gain from the substitution is an increase in the number of observations and panels that can be utilized for model estimation.

Under panel vector auto-regression, idiosyncratic error terms are assumed to be serially uncorrelated with mean zero, and the contemporaneous covariance matrix is assumed to be positive semi-definite (Abrigo and Love 2015). Model estimation is robust to heteroskedasticity. In the analysis, I relax the assumption that idiosyncratic error terms are serially uncorrelated within blocks by clustering $\vec{e}_{i,t}$ at the block level. That is, errors may be dependent within blocks but not across blocks. Kézdi (2004) shows that N>50 is sufficient to use cluster-robust standard error (CRSE) estimation at little to no cost, even in the absence of error clustering. Since the number of blocks under study is much greater than 50, I confidently use CRSE estimation.

Moreover, blocks are assumed to be independent. This is a strong assumption for this study, since it is likely that there is spatial correlation between blocks. However, the distance term interacted with price will control for this spatial relationship between the blocks. An alternative approach would be to substitute the distance-price term with a weighted matrix that controls for the distance between each block (Ottosson et al. 2013; "Panel time-series modeling").

Descriptive Statistics

Table 1 contains descriptive statistics of the data set. A graphical representation of the occupancy statistics for each time period, not just the first and most recent periods, can be found in Appendix A.1 for each combination of time of day and day of week.

Table 1. Descriptive Statistics of Occupancy and Price

	Weekdays						Weekends					
	Morning		*Afternoon*		*Evening*		*Morning*		*Afternoon*		*Evening*	
Variable:	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.
Avg. Occupancy Aug. 2011 (%)	61.52	19.39	72.59	16.60	69.38	16.67	51.58	23.48	66.63	22.51	67.70	20.19
Avg. Occupancy Dec. 2015 (%)	63.68	12.84	67.96	12.43	67.62	11.71	58.20	13.89	68.41	15.02	68.39	14.03
Avg. Occupancy Change (%)	2.16	-	-4.63	-	-1.76	-	6.62	-	1.78	-	0.69	-
Price Aug. 2011 ($/hr)	2.87	0.61	2.87	0.61	2.87	0.61	2.87	0.61	2.87	0.61	2.87	0.61
Price Dec. 2015 ($/hr)	1.90	1.46	3.23	1.43	2.82	1.54	1.16	1.28	2.78	1.88	2.70	1.95
Price Change ($/hr)	-0.97	-	0.36	-	-0.05	-	-1.71	-	-0.09	-	-0.17	-

Notes: Based on the prices and average occupancy rates by time of day and day of week of 205 blocks.

There are several striking features of Table 1. First, average occupancy rates vary widely by time of day and day of the week. In Aug. 2011, occupancy was lowest on weekend mornings and highest on weekday afternoons. After 17 price changes, by Dec. 2015 occupancy was still lowest on weekend mornings but had increased by over 6 percentage points, moving it closer to the target range of 60 to 80 percent. Also, weekday afternoon occupancy had decreased by more than 4 percentage points, narrowing the gap between its average occupancy and the other time periods' average occupancies.

Even more striking is the decrease in the standard deviation of average occupancy rates across blocks from Aug. 2011 to Dec. 2015 for all times and days. The price changes appear to have squeezed more neighborhoods into the target occupancy range, reducing the number of outliers (i.e. blocks with very high or very low occupancy rates) and, more generally, the number of blocks not meeting occupancy goals. This finding does *not* mean that the occupancy rate on any particular block is more stable; it means that overall variability in average occupancy rates across blocks decreases through performance-based price changes. Ottosson et al. (2013) also report decreased standard deviations after the performance-based price change in Seattle, but since their study was conducted over the course of just one year and based on a single price change, the change in standard deviation is small enough (less than 2% in all cases) that the researchers did not mention the reduced standard deviation in their analysis. This is a novel result from this paper: performance-based pricing appears to decrease the variability in average occupancy rates across blocks through time and subsequent price changes.

The average hourly price of parking was $2.87 for all times and days in Aug. 2011 before the first price change. The reason for this is that the price on each block did not change by time of day; in other words, prices were completely rigid at the block level. Variation did occur across blocks, resulting in non-zero standard deviation in each time zone. To achieve the average weekend morning occupancy increase of 6.62 percentage points, prices were lowered an average of $1.71, which is about a 60 percent decrease from their original level. Similarly, the price was raised $0.36 on average to reduce average weekday afternoon occupancy rates by 4.63 percentage points.

Unlike the decrease in standard deviation over time of occupancy rates, the standard deviation of prices rose with time. Blocks started out at a few predetermined, rigid rates, and subsequent incremental adjustments caused their distribution to spread out. In Dec. 2015, the standard deviation in price ranges from $1.28 during the weekend mornings to $1.95 weekend evenings. Price distribution histograms by time of day and day of week can be found in Appendix A.2.

Both occupancy and price show more variability on the weekends than on the weekdays both in Aug. 2011 and Dec. 2015. A possible explanation for this phenomenon is that some shops are closed on the weekends, especially weekend mornings. This may lead to reduced parking demand on blocks near such closed shops, but nearby blocks with fewer or no closings do not experience this drop in demand. On the other hand, shops are likely more consistently open during the weekdays, reducing this effect on the weekdays. As a result, occupancy rates are more likely to show greater variability on weekends than on weekdays. Since price adjustments are based on observed occupancy, it is not surprising then that the larger variability in occupancy rates across blocks during the weekends forces the prices to spread out more over time relative to weekday prices.

Finally, if the SF*park* Project is successfully moving individual blocks towards the "right" price through incremental price changes, then one would expect more and more blocks to fall in the target occupancy range over time. An alternative way to look at this is that fewer blocks would be expected to required price changes during each subsequent time period. Indeed, this is generally what has happened in SF*park*, as shown in Figure 2 below.

Figure 2. Number of Price Changes Over Time

After the first 6 price changes, the general trend in number of price changes appears to be downwards, which supports the efficacy of SF*park*'s pricing strategy.

Model Results

Table 2. Parameter Estimation of Auto-Regressive Component (Matrix A in Sys. 1)

			Natural Logarithm of Lagged Occupancy at Time *t* - 1					
			Weekdays			Weekends		
			Morning	*Afternoon*	*Evening*	*Morning*	*Afternoon*	*Evening*
Natural Logarithm of Occupancy at Time *t*	Weekdays	*Morning*	0.519*** (0.114)	0.272* (0.144)	-0.062 (0.112)	-0.013 (0.045)	-0.023 (0.053)	0.003 (0.060)
		Afternoon	0.092 (0.085)	0.542*** (0.133)	-0.019 (0.090)	-0.010 (0.040)	0.028 (0.040)	0.028 (0.048)
		Evening	-0.020 (0.109)	0.227 (0.171)	0.428*** (0.115)	0.023 (0.039)	-0.017 (0.053)	0.025 (0.061)
	Weekends	*Morning*	-0.028 (0.252)	-0.008 (0.375)	0.244 (0.274)	0.145* (0.088)	-0.226*** (0.084)	0.115 (0.119)
		Afternoon	0.043 (0.264)	0.447 (0.400)	-0.289 (0.279)	0.072 (0.088)	-0.238** (0.099)	0.143 (0.113)
		Evening	0.054 (0.247)	0.319 (0.423)	-0.165 (0.353)	0.105 (0.080)	-0.355*** (0.103)	0.285** (0.114)

Notes: Based on 2,514 observations, 205 panels, average number of T=12.263; the exogenous variables, as well as 4 lags of the dependent variable, were used as instruments for the pVAR(1) model; Hansen's J-statistic for the overall model was 87.20 with 108 degrees of freedom giving a p-value of 0.929; FOD used to control for unobserved heterogeneity; *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 2 reports the estimated coefficients of Matrix A in System 1. Recall that forward orthogonal deviations are utilized to control for unobserved heterogeneity rather than first differences. This significantly impacts the interpretation of these results. The parameters in Table 2 capture the extent to which lagged occupancy during time *t*-1 is linearly correlated with occupancy at time *t* during different times of the day and days of the week after removing dependent variable-specific fixed effects at the block level. Using FOD, a positive, statistically significant coefficient along the main diagonal indicates that $y_{i,t-1}$ is positively linearly correlated with $y_{i,t}$ during the particular time of day and day of the week corresponding to the row and column label.

In general, for autoregressive models to be stable, the estimated coefficients on lagged terms should be less than one in absolute value. This precludes the possibility of selfperpetuating effects in the system that amplify in magnitude over time rather than diminish. For example, if there is a large music festival that increases average occupancy rates at time *t*-1, occupancy rates after the festival at time *t* may remain slightly elevated from their "normal" levels, all else equal, but one would not expect occupancy to increase every period thereafter at *t*, *t*+1, *t*+2, … without any other changes to the system. Eventually, average occupancy rates should settle back down to their pre-festival rates unless there is some other change to parking demand or supply. Glancing over Table 2, it is clear that all statistically significant coefficients are less than one in absolute value. This supports the proposition that the model is stable. A technical test for parameter estimation stability is presented in Appendix B.3, and the model satisfies the stability condition.

Except for afternoon weekends, all occupancy rates are positively correlated with their own lagged values after controlling for price, seasonality, distance, neighborhood fixed effects, and block-level unobserved heterogeneity. Since occupancy during some time of the day and day of the week at time *t*-1 is positively correlated with occupancy at *t*, this simple result bolster's SF*park*'s practice of using the previous period's average occupancy during a particular time zone to adjust that time zone's hourly price for the next period.

On the other hand, the lagged weekend afternoon occupancy is actually negatively correlated with current weekend morning, afternoon, and evening occupancy. This is an unusual result and requires further attention. A possible explanation for this is that individuals are pay close attention to occupancy rates on the weekend afternoons when planning future automobile trips. For example, suppose an individual goes shopping one weekend afternoon and finds that all of the stores are overcrowded, resulting in an unpleasant shopping experience. He likely experiences difficulty with vehicle and pedestrian traffic, as well as finding parking, in such a situation. So when subsequent weekends roll come around, the hypothetical individual is less likely to make a trip out on the weekend, perhaps assuming that the weekends are too busy and crowded for an enjoyable outing.

Table 3 presents the parameter estimation of the distance-price interaction in System (1). Recall that the interpretation of these coefficients is the change in the price elasticity of demand associated with a one kilometer move away from the center of the PMD. It is immediately apparent that none of the coefficients are statistically significant.

Table 3. Parameter Estimation of Distance-Price Interaction in Sys. 1

			Distance*Natural Logarithm of Price at Time *t*					
			Weekdays			Weekends		
			Morning	*Afternoon*	*Evening*	*Morning*	*Afternoon*	*Evening*
Natural Logarithm of Occupancy at Time *t*	Weekdays	*Morning*	-0.635 (0.554)	0.603 (0.722)	-0.736 (0.543)	0.708 (0.416)	-0.185 (0.719)	-0.042 (0.647)
		Afternoon	-0.002 (0.442)	0.454 (0.496)	-0.258 (0.345)	0.153 (0.316)	-0.476 (0.473)	0.365 (0.409)
		Evening	0.204 (0.658)	0.138 (0.812)	-0.448 (0.436)	0.010 (0.506)	-0.444 (0.728)	0.492 (0.608)
	Weekends	*Morning*	-0.525 (1.456)	-1.191 (1.423)	0.299 (1.216)	0.258 (0.998)	0.900 (1.389)	-0.868 (1.138)
		Afternoon	1.701 (1.378)	-0.957 (1.361)	-0.583 (1.195)	-0.765 (0.937)	0.962 (1.364)	-0.589 (1.186)
		Evening	1.737 (1.399)	-1.269 (1.551)	-0.102 (1.366)	-0.957 (1.043)	0.018 (1.265)	0.455 (1.306)

Notes: *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 4 presents (partial) parameter estimation results of the PMD-price interaction in System 1. These coefficients can be interpreted as follows: they represent the percentage change in occupancy during a time of the day and day of the week associated with a one percent change in price during that *same* time zone *in a particular PMD when* $dist_i$*=0.* A full table of output would include cross-elasticities that capture the percent change in occupancy corresponding to a percent change in price during *different* time zones (i.e. a different day of the week and/or time of day). Note that none of the coefficients in Table 3 are statistically significant. The magnitude of the coefficients and associated standards is large for Fisherman's Wharf during weekend afternoons and evenings. Large fluctuations in price elasticities at the block level within the PMD during these time zones could potentially cause this phenomenon. It is also possible that there remains some PMD-level omitted-variable bias that disproportionately impacts Fisherman's Wharf.

Table 4. Parameter Estimation of PMD-Price Interaction in Sys. 1

			Parking Management District*Natural Logarithm of Price at Time *t*				
			Civic Center	*Downtown*	*Fillmore*	*Fisherman's Wharf*	*South Embarcadero.*
Natural Logarithm of Occupancy at Time *t*	Weekdays	*Morning*	-0.111 (0.308)	-0.893 (1.370)	0.449 (0.524)	0.125 (0.276)	0.315 (0.317)
		Afternoon	-0.225 (0.425)	0.710 (1.189)	-0.079 (0.581)	-0.936 (0.618)	-0.247 (0.283)
		Evening	0.208 (0.267)	-1.038 (2.006)	0.114 (0.329)	-0.330 (0.406)	0.567 (0.413)
	Weekends	*Morning*	0.077 (0.331)	-0.675 (0.641)	-0.738 (1.121)	-0.475 (0.646)	-0.158 (0.488)
		Afternoon	-1.211 (0.941)	-0.921 (1.759)	-1.431 (1.258)	5.933 (6.110)	-0.767 (0.986)
		Evening	0.829 (1.094)	-0.635 (1.359)	0.388 (1.449)	-4.808 (5.142)	-0.526 (0.773)

Notes: *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

The coefficient estimates from Table 3 and Table 4 can be used to predict the price elasticity of demand on any given block i. Let β be the coefficient on the distance-price interaction from Table 3, and let α_k be the coefficient on the PMD-price interaction from Table 4 for block k. Then the price elasticity of demand during a particular day of the week and time of day on block i, e_i, can be estimated as

$$\widehat{e_i} = \beta dist_i + \sum_{k=1}^{5} \alpha_k D_{k,i}(pmd) \tag{4}$$

Note that e_i is assumed to be time-invariant. The estimated $\widehat{e_i}$ are presented below in Table 5.

Time Zone	Mean $\widehat{e_\iota}$	S.D.	Min.	Max.
Weekday Morning	-0.233	0.389	-1.243	0.417
Weekday Afternoon	0.006	0.466	-0.913	0.959
Weekday Evening	-0.176	0.457	-1.284	0.563
Weekend Morning	-0.228	0.334	-0.725	0.307
Weekend Afternoon	0.625	2.663	-1.382	6.702
Weekend Evening	-0.603	1.955	-4.785	1.101

Table 5. $\widehat{e_\iota}$ Summary Statistics
Notes: Based on 209 blocks

The results from Table 5 should be interpreted with caution since none of the coefficients in Table 2 or Table 3 are statistically significant. Moreover, the large coefficient estimate on Fisherman's Wharf during the weekend afternoon and evening introduces considerable variability to these two overall price elasticity estimates.

I would expect these coefficients to be less than 0; however, the weekend afternoon estimate is much larger than 0. This is an odd result and is attributable to the influence of Fisherman's Wharf on this estimate. In all time zones, the standard deviation across block price elasticities is large enough such that some block-level elasticities are positive and others are negative.

This is not inconsistent with rough calculations using the midpoint formula to calculate price elasticities. Table 6 gives the results obtained from calculating observed elasticities using the midpoint formula and aggregating the results according to time zone.

Time Zone	**Obs.**	**Mean**	**S.D.**	**Min.**	**Max.**
Weekday Morning	1,578	-0.535	2.129	-15.526	9.025
Weekday Afternoon	1,551	-0.895	2.355	-25.900	15.658
Weekday Evening	1,390	-0.908	2.589	-30.977	23.459
Weekend Morning	1,531	-0.504	2.685	-17.475	19.697
Weekend Afternoon	1,653	-0.961	3.107	-23.438	20.444
Weekend Evening	1,702	-0.780	2.831	-23.400	20.700

Table 6. Midpoint Elasticity Summary Statistics

The mean of the observed midpoint elasticities is negative in all cases, as expected. However, the standard deviation is larger than in the predicted model, and the range of values is much greater. A possible reason for these results is that, unlike the pVAR(1) model proposed in this paper, using the observed occupancy and price data to compute midpoint elasticities does not require the assumption that the price elasticity of demand be time-invariant. The implications of this will be further addressed in the Discussion.

IV. Discussion

The model proposed in this paper makes very few structural assumptions, controls for serially correlated error terms within blocks, and accounts for block-level unobserved heterogeneity. Despite these advantages, the model fails to show statistically significant relationships between occupancy and price, distance, and PMD.

In plain words, the implication of this finding is that block-level price elasticities are not easy to predict. As the qualitative analysis by Pierce and Shoup (2013) argues, it appears to be the case that there is simply too much variability in price sensitivity at the block level in order to accurately predict elasticities, at least using the econometric model presented in this paper. The novel contribution of this paper to the literature is that price elasticity prediction is not difficult because of the variability in price elasticity but rather the lack of structure of this variability. In this sense, this paper goes one step further than Pierce and Shoup (2013), who simply observe wide price elasticity variation and jump to the conclusion that economists cannot develop robust econometric models to predict block-level price elasticities.

A logical question that arises is the following: why is there no apparent structure in the variability of on-street parking demand sensitivity to price? One possible explanation is that the price of on-street parking is relatively unimportant in determining occupancy rates. When a person considers where to park, he takes many variables into account, both consciously and subconsciously. His trip has a purpose, and this purpose is unobservable to economists, as well as other factors such as his preferences regarding physical exercise, ability to parallel park, and a number of location- and time-specific potential influences. The goal of this study was to try to understand aggregate behavior: how average occupancy rates change in response to price changes. Nevertheless, realizing that ultimately individual choices and preferences influence parking decisions unlocks the door to a number of possible confounding variables.

To elaborate on one example, suppose that in 2014 a groundbreaking study was published that definitively established a causal relationship between physical exercise and cancer mortality rates. In response, public health campaigns may begin to encourage people to walk a certain number of steps per day in order to reduce the risk of cancer. All of a sudden, individuals do not try to park immediately near their destinations anymore; they seek to park further away so that they can walk more and maintain a healthy lifestyle. Now, for many individuals the relationship between distance between desired destination and parking spot is not primarily based on the opportunity cost of time spent walking but rather represents a more complex tradeoff that invokes the value of health and life. Therefore, if the pre-2014 model is used to estimate post-2014 data, the parameter estimates will suffer from omitted variable bias.

While the above scenario is clearly fabricated, the idea is generalizable. There are so many changing conditions that may impact fundamental relationships in any proposed econometric model of parking that reliable prediction of parameters is not feasible.

A significant limitation in my study is that I assume that the block-level price elasticity in each time zone is constant over time. In other words, the model assumes that price elasticity is a function of PMD and distance from the center of the PMD only, and this price elasticity is estimated for all time zones (i.e. times of day and days of the week). In the model's estimation, there is no allowance for the possibility for structural changes to this relationship as time progresses. If some other influential variables were to arise in, say, 2013, then the model's parameter estimates would be biased.

The calculation of observed midpoint elasticities provided an interesting exercise against which I could compare the model results. These point estimates of the midpoint elasticities were all negative on average but showed huge variability across individual estimates. Pierce and Shoup (2013) conduct a similar exercise and utilize this observation as key evidence for their claim that elasticities vary too widely to be estimated reliably. Note that the mean price elasticity estimates are reasonably comparable between Table 5 and Table 6, but the standard deviations are not directly comparable. I say the means are "reasonably" comparable because the model provides exactly one estimate per block per time zone. In contrast, the midpoint method may provide many estimates for some blocks and none for others; in other words, the midpoint method does not ensure that the blocks are equally represented in calculating the average elasticity during a particular time zone.

My model also omits off-street parking based on the findings of Ottosson et al. (2013) and public transportation based on lack of variation in number of accessibility points per block. If either one of these variables did in fact turn out to be important determinants of price elasticity, then my study would also suffer from omitted variable bias.

Ottosson et al. (2013) have statistically significant price elasticity estimates and find that distance is an important determinant of this demand sensitivity. I suggest that their findings are due to the single price change used in the study. Their study takes place over a single year, which reduces the chance that time-dependent structural changes in the price-elasticity of demand determinants occur between the beginning and end of their study. This starkly contrasts the approximately 4-year period under study in this paper, where there is greater likelihood that omitted variable bias becomes an issue sometime during the course of the study and thus disproportionately affects a portion of the data.

Future research on this topic should seek to more carefully control for potential structural changes over time in the relationship among the determinants of the price elasticity of demand. In addition, future research may consider adding off-street parking and public transit data in order to explore the effects of these variables further.

V. Conclusion

This study has presented a sophisticated econometric model that seeks to relate occupancy rates and prices after controlling for endogeneity and confounding factors. The preferred model does not show a statistically significant effect of key variables of interest, such as distance, PMD, and price, on occupancy rates. However, there remain promising avenues for future research in the parking literature, particularly in the context of the SF*park* Project.

Works Cited

Abrigo, Michael R.M. and Inessa Love. "Estimation of Panel Vector Autoregression in Stata: a Package of Programs." University of Hawaii, Department of Economics. Feb. 2015. Web. 10 March 2016.

Andrews, Donald W.K. and Biao Lu. "Consistent model and moment selection procedures for GMM estimation with application to dynamic panel models." *Journal of Econometrics* 101 (2001): 123-164. Web. 13 March 2016.

Anderson, Simon P. and André de Palma. "The economics of pricing parking." *Journal of Urban Economics* 55 (2004): 1-20. *ScienceDirect*. Web. 26 May 2015.

---. "Parking in the city." *Papers in Regional Science* 86.4 (2007): 621-632. Web. 31 July 2015.

Arnott, Richard and Eren Inci. "An integrated model of downtown parking and traffic congestion." *Journal of Urban Economics* 60 (2006): 418-442. *ScienceDirect*. Web. 1 Jan. 2016.

Arnott, Richard and John Rowse. "Downton parking in auto city." *Regional Science and Urban Economics* 39 (2009): 1-14. *ScienceDirect*. Web. 1 Sept. 2015.

---. "Modeling Parking." *Journal of Urban Economics* 45 (1999): 97-124. Web. 31 Dec. 2015.

Baker, David R. "S.F. begins charging for Sunday parking." SFGate. 6 Jan. 2013. Web. 12 Apr. 2016.

Baum, Christopher F. "Dynamic Panel Data estimators." Boston College. Spring 2013. Web. 1 Apr. 2016.

Eberhardt, Markus. "Panel time-series modeling: New tools for analyzing xt data." University of Nottingham. 2011 UK Stata Users Group Meeting, Cass Business School, London. 16 Sept. 2011. Conference Presentation.

"Economic Growth Widespread Across Metropolitan Areas in 2014." *Bureau of Economic Analysis* 23 Sept. 2015. Web. 1 Jan. 2015.

Hamre, Andrea and Ralph Buehler. "Commuter Mode Choice and Free Car Parking, Public Transportation Benefits, Showers/Lockers, and Bike Parking at Work: Evidence from the Washington, DC Region 2014." *Journal of Public Transportation* 17.2 (2014): 67-91. Web. 31 Dec. 2015.

Holtz-Eakin, Douglas, Whitney Newey, and Harvey S. Rosen. "Estimating Autoregressions with Panel Data." *Econometrica* 56.6 (1988): 1371-1395. *JStor*. Web. 21 March 2016.

Judson, Ruth A. and Ann L. Owen. "Estimating dynamic panel data models: a guide for macroeconomists." *Economics Letters* 65 (1999): 9-15. *ScienceDirect*. Web. 14 March 2016.

Kelly, J. Andrew and J. Peter Clinch. "Temporal variance of revealed preference on-street parking price elasticity." *Transport Policy* 16.4 (2009): 193-199. *ScienceDirect*. Web. 31 Dec. 2015.

Kézdi, Gábor. "Robust Standard Error Estimation in Fixed-Effects Panel Models." Spec. issue of *Hungarian Statistical Review* 9 (2004): 95-116. Web. 7 Apr. 2016.

Lam, Stephen. "Facebook offers workers $10K to cut the commute." *Al Jazeera America*. Al Jazeera America, 18 Dec. 2015. Web. 1 Jan. 2015.

Litman, Todd. "Parking Pricing Implementation Guidelines: How More Efficient Parking Pricing Can Help Solve Parking And Traffic Problems, Increase Revenue, And Achieve Other Planning Objectives." Victoria Transport Policy Institute. 24 Aug. 2015. Web. 11 Apr. 2016.

Nickell, Stephen. "Biases in dynamic models with fixed effects." *Econometrica* 49.6 (1981): 1417-1426. *JSTOR*. Web. 14 March 2016.

Ommeren, Jos van and Giovanni Russo. "Time-varying parking prices." *Economics of Transportation* 3 (2014): 166-174. *ScienceDirect*. Web. 7 July 2015.

Ottosson, Dadi Baldur, Cynthia Chen, Tingting Wang, and Haiyun Lin. "The sensitivity of on-street parking demand in response to price changes: A case study in Seattle, WA." *Transport Policy* 25 (2013): 222-232. *ScienceDirect*. Web. 7 Oct. 2015.

"Parking Meters." *City of Boston*. CityofBoston.gov, n.d. Web. 20 Dec. 2015.

Pierce, Gregory and Donald Shoup. "Getting the Prices Right." *Journal of the American Planning Association* 79.1 (2013): 67-81. Web. 12 Oct. 2015.

Preston, Darrell. "Morgan Stanley Group's $11 Billion Makes Chicago Taxpayers Cry." *BloombergBusiness*. 2010 Aug. 8. Web. 20 Dec. 2015.

Qian, Sean and Ram Rajagopal. "Optimal occupancy-driven parking pricing under demand uncertainties and traverler heterogeneity: A stochastic control approach." *Transportation Research Part B: Methodological* 67 (2014): 144-165. *ScienceDirect*. Web. 31 Dec. 2015.

Schrank, David, Bill Eisele, Tim Lomax, and Jim Bak. "2015 Urban Mobility Scorecard." The Texas A&M Transportation Institute and INRIX, Inc. Aug. 2015. Web. 1 Jan. 2015.

"SF*park*: Pilot Project Evaluation." *SFMTA*. SFpark.org., June 2014. Web. 26 Oct. 2015.

Shoup, Donald. "Cruising for parking." *Transport Policy* 13 (2006): 479-486. Web. 24 Feb. 2015.

---. "Evaluating the effects of cashing out employer-paid parking: Eight case studies." *Transport Policy* 4.4 (1997): 201-216. Web. 31 Dec. 2015.

---. “The High Cost of Free Parking.” *Journal of Planning Education and Research* 17 (1997): 3-20. Web. 30 May 2015.

---. “The trouble with minimum parking requirements.” *Transportation Research Part A* 33 (1999): 549-574. Web. 31 Dec. 2015.

Washbrook, Kevin, Wolfgang Haider, and Mark Jaccard. “Estimating commuter mode choice: A discrete choice analysis of the impact of road pricing and parking charges.” *Transportation* 33.6 (2006): 621-639. Web. 31 Dec. 2015.

Weinberger, Rachel, John Kaehny, and Matthew Rufo. “U.S. Parking Policies: An Overview of Management Strategies.” *Institute for Transportation & Development Policy*. ITDP, 23 Dec. 2010. Web. 31 Dec. 2015. Web.

Verhoef, Erik T., Jan Rouwendal, and Piet Rietveld. “Congestion Caused by Speed Differences.” *Journal of Urban Economics* 45.3 (1999): 533-556. Web. 31 Dec. 2015.

Appendix A

A.1 Time-Series Occupancy Summary

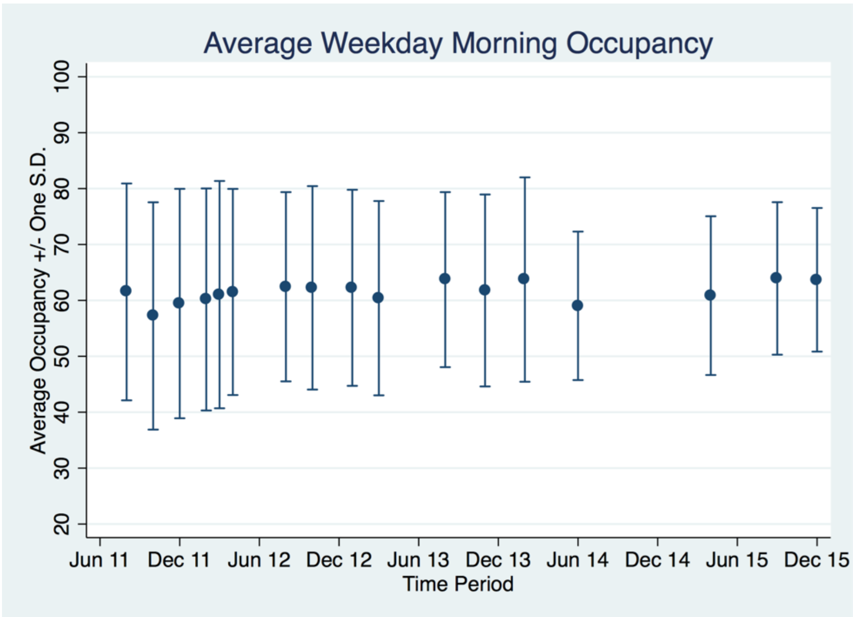

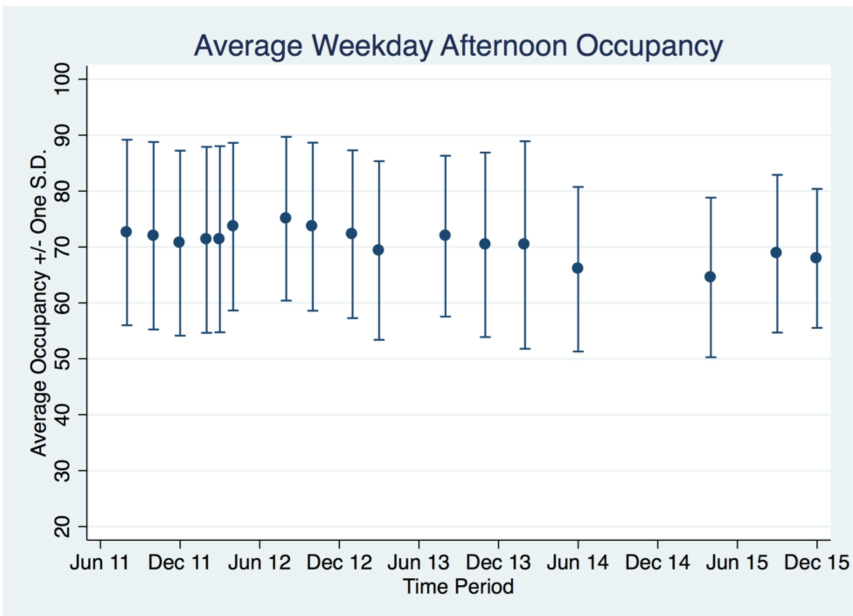

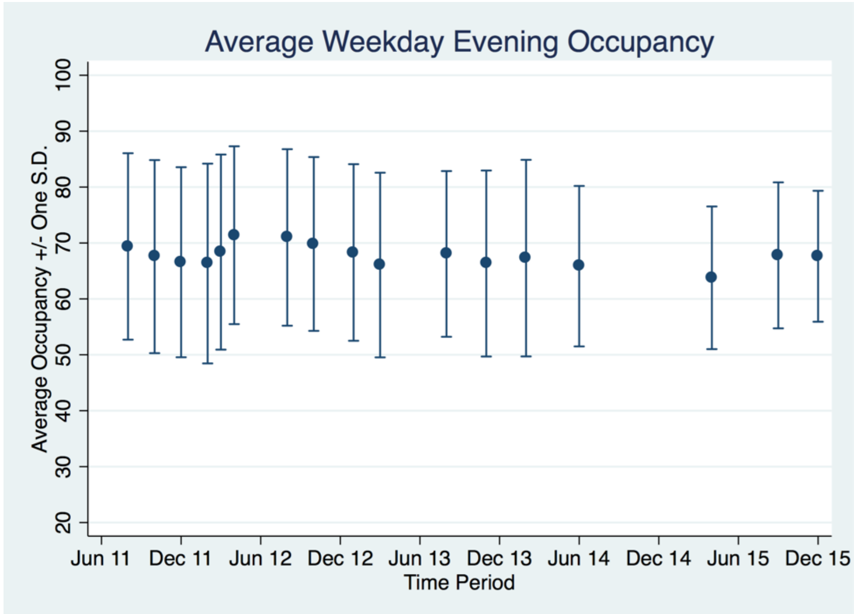
Average Weekday Evening Occupancy
Average Occupancy +/- One S.D.
100
90
80
70
60
50
40
30
20
Jun 11
Dec 11
Jun 12
Dec 12
Jun 13
Dec 13
Jun 14
Dec 14
Jun 15
Dec 15
Time Period

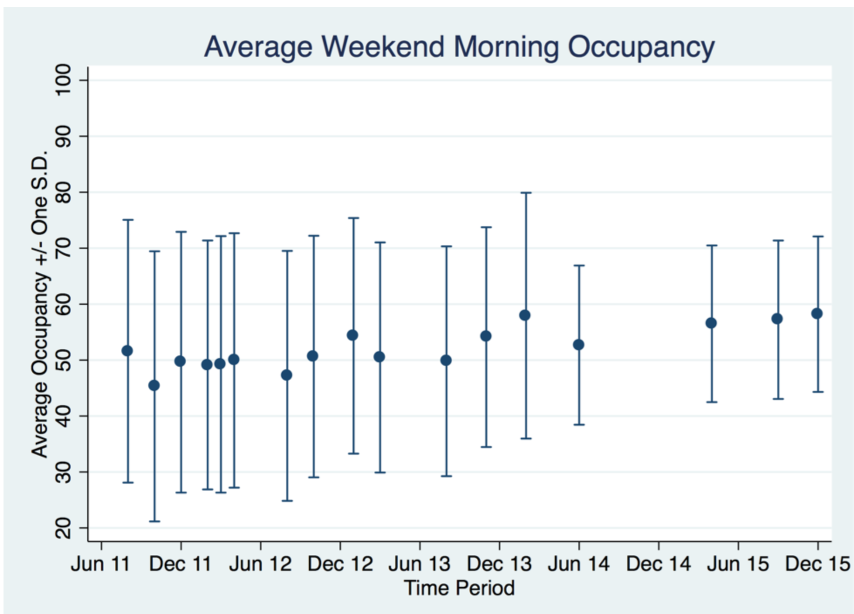
Average Weekend Morning Occupancy
Average Occupancy +/- One S.D.
100
90
80
70
60
50
40
30
20
Jun 11
Dec 11
Jun 12
Dec 12
Jun 13
Dec 13
Jun 14
Dec 14
Jun 15
Dec 15
Time Period

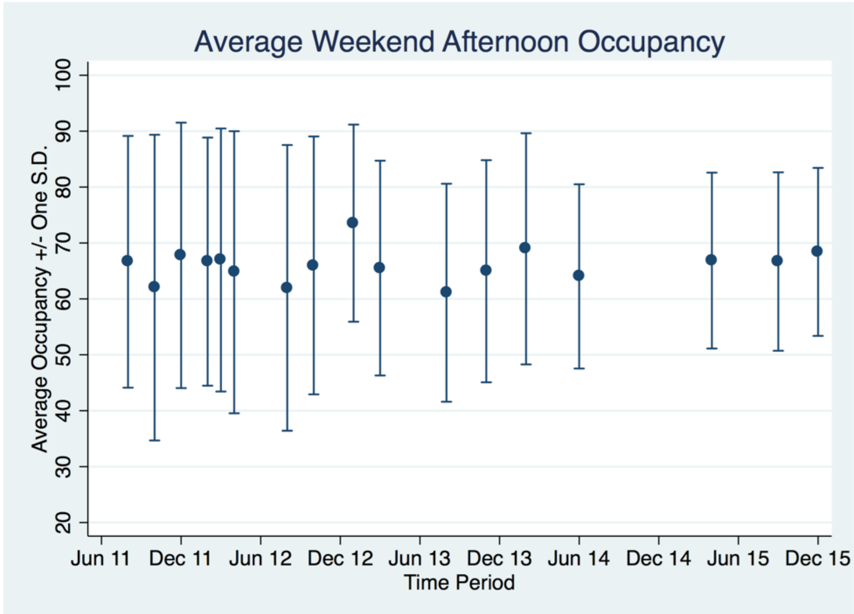
Average Weekend Afternoon Occupancy
Average Occupancy +/- One S.D.
100
90
80
70
60
50
40
30
20
Jun 11
Dec 11
Jun 12
Dec 12
Jun 13
Dec 13
Jun 14
Dec 14
Jun 15
Dec 15
Time Period

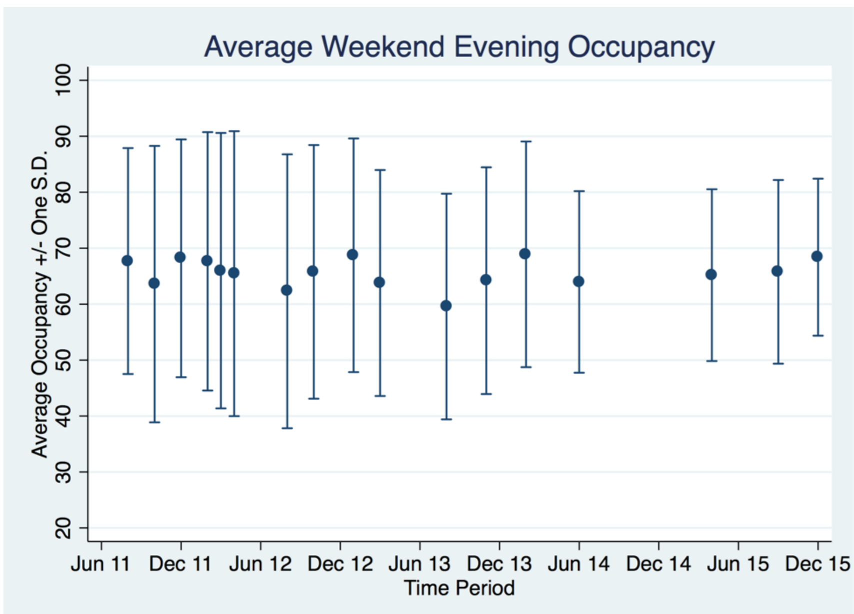
Average Weekend Evening Occupancy
Average Occupancy +/- One S.D.
100
90
80
70
60
50
40
30
20
Jun 11
Dec 11
Jun 12
Dec 12
Jun 13
Dec 13
Jun 14
Dec 14
Jun 15
Dec 15
Time Period

A.2 Price Distribution Histograms

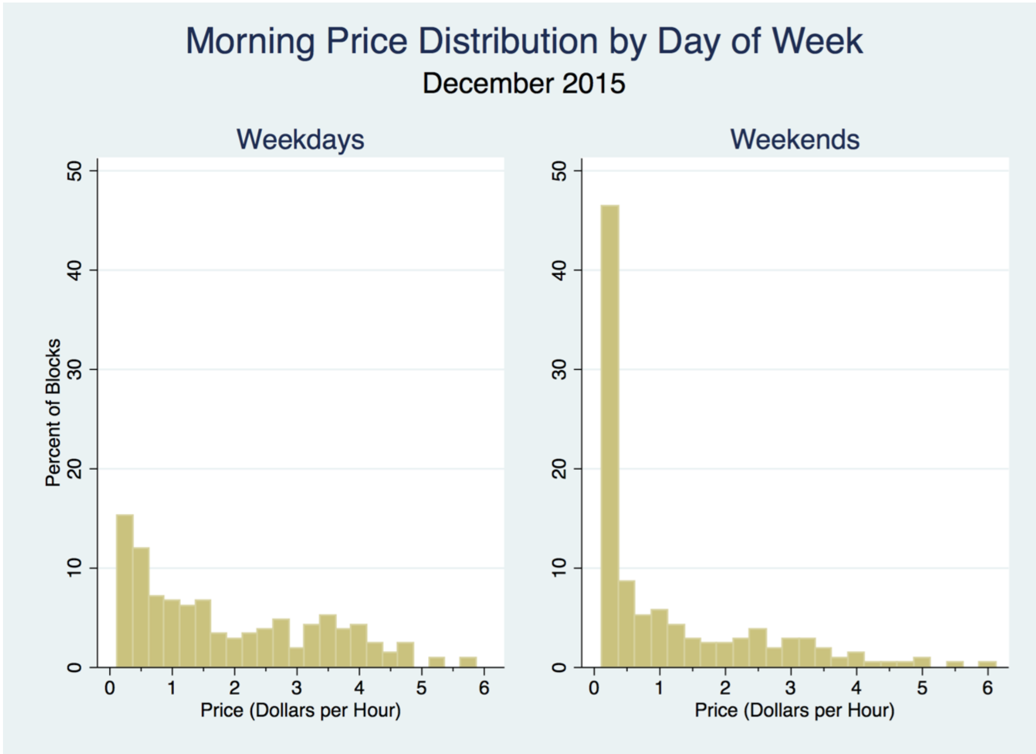

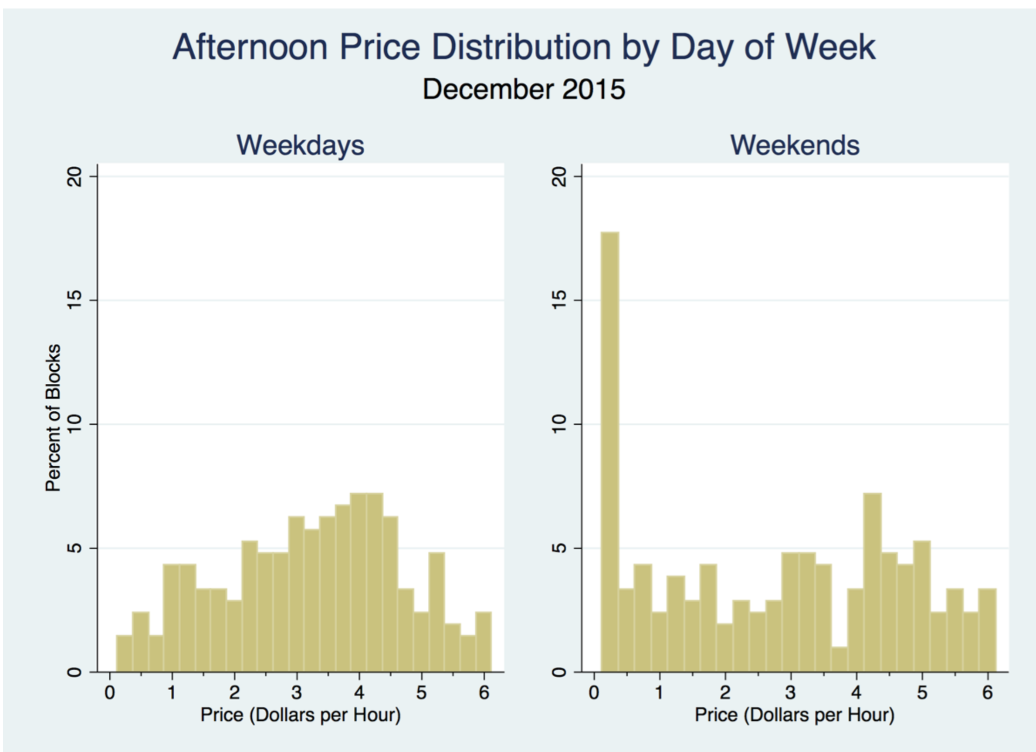

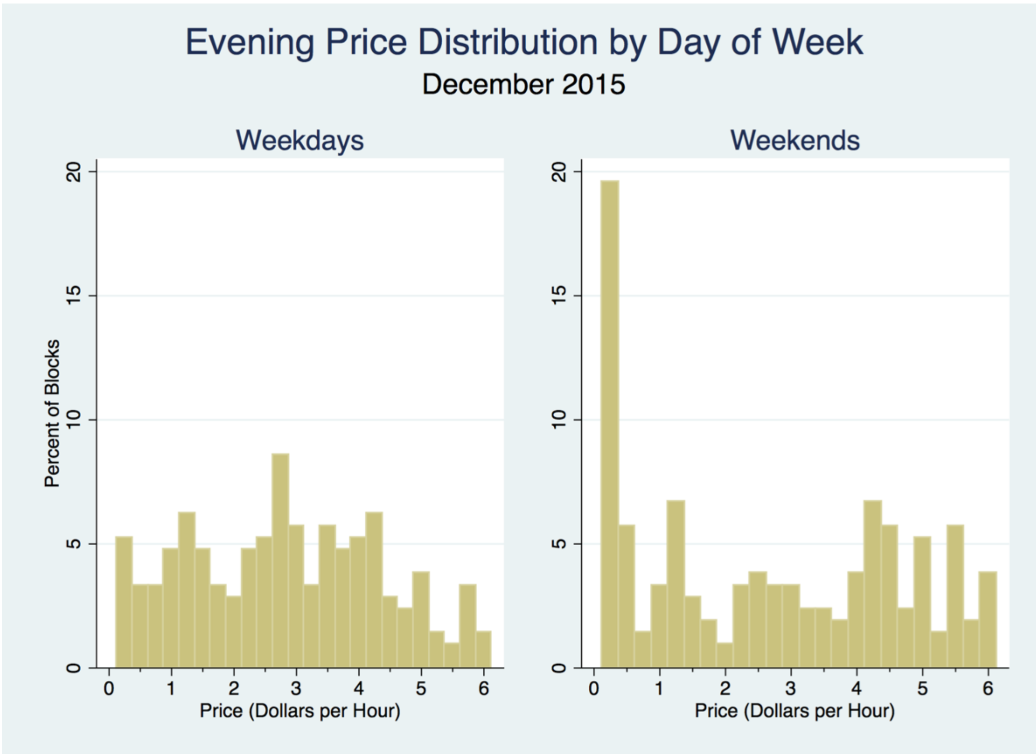
Evening Price Distribution by Day of Week
December 2015
Weekdays
Weekends
Percent of Blocks
0
5
10
15
20
0
1
2
3
4
5
6
Price (Dollars per Hour)
Price (Dollars per Hour)

Appendix B

B.1 Lag-Order Selection

The pvarsoc program, created by Abrigo and Love (2015), was utilized to determine to the appropriate lag-order for the panel auto-regression analysis. The dependent variables can be represented as $\vec{y}_{i,t}$, which is a vector of the natural logarithm of average occupancy rates (expressed as percentages) at different times of the day on weekdays and weekends on block *i* during time period *t*. The goal is to find the lag-order $p \in \{1, 2, \dots, T-1\}$ such that the model

$$\vec{y}_{i,t} = \sum_{j=1}^{T-1} A_j \vec{y}_{t-j} + B\, X_{i,t} + \vec{u}_i + \vec{e}_{i,t} \tag{B.1}$$

provides the best fit for the data. The measures used to determine the "best fit" of the data are the Andrews and Lu (2001) moment model selection criteria (MMSC), which are based on Hansen's J statistic of over-identifying restrictions, as well as the number of moment conditions, the number of endogenous variables, and the sample size (Abrigo and Love 2015). These include the MMSC-Bayesian Information Criterion (MBIC), MMSC-Akaike Information Criterion (MAIC), and the MMSC-Hannan-Quinn Information Criterion (MQIC). The best-fitting model will minimize the MMSC criteria. Additionally, pvarsoc reports the coefficient of determination (CD) of a particular lag-order as a representation of the proportion of variation captured by the model (Abrigo and Love 2015).

Since the time intervals are approximately three months long and therefore loosely resemble quarterly data, I follow Ryan-Collins (2016) in using 4 lags of the dependent variables as moment conditions. Specifying before testing that 4 lags be used as instruments prevents the temptation to "over-fit" the data. The lagged instruments are defined to be the first to fourth lags from the highest lag used in the pVAR model. For example, when testing a pVAR model of order 2, the second to fifth lags are used as instruments.

As a final consideration before testing, since the number of moment conditions must be greater than the number of lags of the endogenous variables in order to calculate the J statistic, the maximum lag order I test is 3. The lag-order selection test results are shown in Table B.1.

Table B.1. Lab-Order Selection Criteria*

Lags	CD	J	J P-Value	MBIC	MAIC	MQIC
1	0.9871842	124.8447	0.1279227	-625.7398	-91.15528	-293.919
2	0.8367188	56.40032	0 .9116846	-443.9893	-87.59968	-222.7755
3	**	6.09532	1.0000000	-244.0995	-65.90468	-133.4926

*Based on 1,043 observations, 186 panels, average number of T = 5.608
**Internal numerical error in Stata prevented calculation of CD for lag 3

The results from Table B.1 show that the MBIC, MAIC, and MQIC are minimized with the first-order PVAR model. The CD is also maximized with the first-order model, although the CD was not able to be computed for the third-order model due to numerical errors in Stata (perhaps due to the small magnitude of the determinant of the unconstrained covariance matrix of the dependent variable in this model). Hansen's J-statistic is declines as the number of lags increases, suggesting that model fit improves as lag order increases; however, this measure does not account for degrees of freedom in the model like the MMSC criteria and is thus considered less heavily in model selection (Abrigo and Love 2015). In all cases, though, the p-value of Hansen's J-statistic is insignificant at the 10% significance level, so for all three models I fail to reject that the null hypothesis that the over-identifying restrictions are valid.

The above discussion leads me to choose the first-order model as the preferred model throughout the paper.

B.2 Stationarity Testing

It is important to determine whether or not the key endogenous variables are stationary in an autoregressive model in order to understand the effect of shocks on the system and identify spurious relationships.

The data in this study are unbalanced (i.e. some panels have missing occupancy observations), which limits the panel root tests that may be applied. I implement the Fisher unit root test in Stata using the xtunitroot fisher program. The Fisher unit root test performs unit root tests on each panel individually and essentially aggregates the p-values to report a single statistic. The null hypothesis is that the all panels contain a unit root, and the alternative hypothesis is that at least one panel is stationary. The test performed on each panel is the Phillips-Perron test, which I choose since it uses Newey-West standard errors to control for the possibility of serial correlation, with a linear time trend. Before performing the test, I subtract the cross-sectional averages from the series, as recommended by Levin, Lin, and Chu (2002), to reduce the potential impact of cross-sectional dependence. Four lags of the dependent variables are used in the test since the average time period is three months and thus resembles quarterly data (Ryan-Collins 2016). Finally, to test the null hypothesis specified above, Choi (2001) recommends using the Z test for empirical applications based on simulation results supporting the Z test's ability to outperform alternative tests (i.e. P, L^*, P_m tests) with respect to size and power considerations.

The results of the above statistical procedure are presented for the endogenous variables $\vec{y}_{i,t}$ in Table A.2. The null hypothesis is rejected at the 5% level or less in all cases, so I reject the null hypotheses that the endogenous variables are non-stationary. I conclude that each series of ln(occupancy rates) is I(0), i.e. stationary.

Table B.2. Fisher Unit Root Test Results*

	Variable	Inverse Normal Statistic (Z)	P-Value
Weekdays	Morning occupancy	-13.8998	0.0000
	Afternoon occupancy	-12.1727	0.0000
	Evening occupancy	-14.1685	0.0000
Weekends	Morning occupancy	-24.6148	0.0000
	Afternoon occupancy	-25.5827	0.0000
	Evening occupancy	-22.3086	0.0000

*Based on 205 to 209 panels, average number of T=14.97; all occupancy variables are the natural logarithm of average occupancy rates

B.3 Stability Testing

Figure B.1 shows that the stability condition of the pVAR estimates is met since all eigenvalues of the companion matrix lie inside the unit circle, i.e. have a modulus less than one.

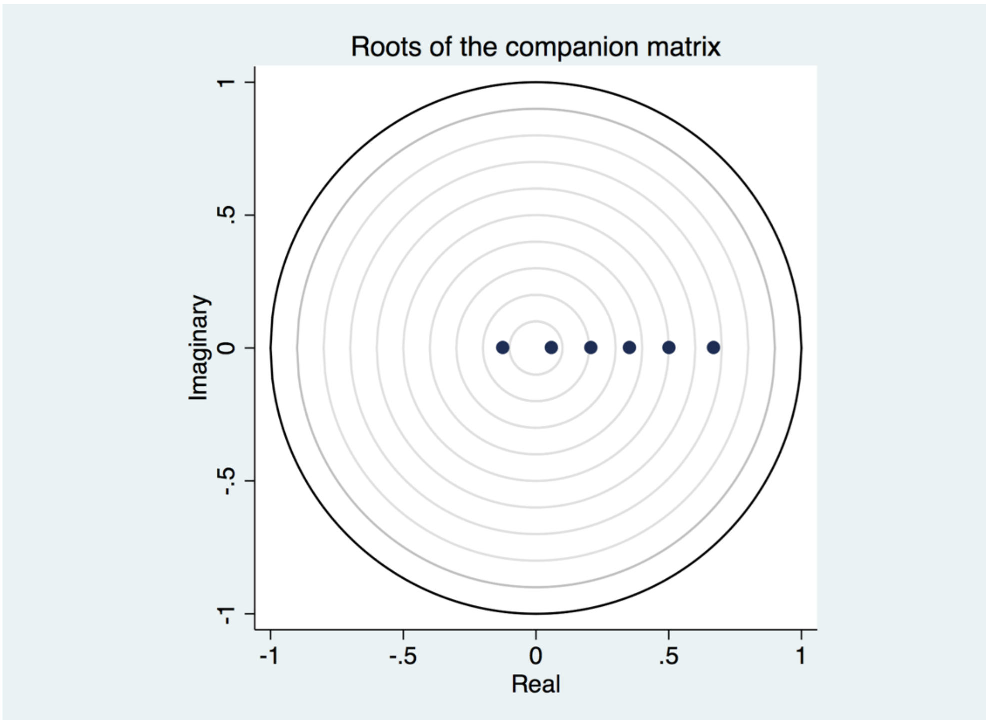

Figure B.1. Eigenvalues of the Companion Matrix

Appendix Works Cited

Abrigo, Michael R.M. and Inessa Love. "Estimation of Panel Vector Autoregression in Stata: a Package of Programs." University of Hawaii, Department of Economics. Feb. 2015. Web. 10 March 2016.

Andrews, Donald W.K. and Biao Lu. "Consistent model and moment selection procedures for GMM estimation with application to dynamic panel models." *Journal of Econometrics* 101 (2001): 123-164. *ScienceDirect*. Web. 13 March 2016.

Choi, In. "Unit root tests for panel data." *Journal of International Money and Finance* 20 (2001): 249-272. *ScienceDirect*. Web. 14 March 2016.

Levin, Andrew, Chien-Fu Lin, and Chia-Shang James Chu. "Unit root tests in panel data: asymptotic and finite-sample properties." *Journal of Econometrics* 108 (2002): 1-24. *ScienceDirect*. Web. 14 March 2016.

Ryan-Collins, Joshua Jacob. "Credit creation, monetary policy and the macroeconomy: three empirical studies." University of Southampton, School of Business. Feb. 2016. Web. 13 March 2016.