



# Market-Priced Parking in Theory and Practice

---

MICHAEL MANVILLE AND DANIEL G. CHATMAN

**O**ne of the first lessons of economics is that price controls lead to shortages, and shortages lead to queues. Street parking vividly illustrates this principle. Many cities keep valuable street spaces free or underpriced, and as a result they fill up quickly, creating shortages at busy times. These shortages then create moving queues as drivers circle the block, or “cruise,” searching for spaces. Cruising, in turn, creates congestion and pollution.

The textbook answer to this problem is simple: remove the price control and let the market set the price for curbside parking. The “right” price will keep one or two spaces open but no more. Just as a private firm wants its inventory to sell briskly without being exhausted, so too should cities keep parking spaces well-used but never completely full. With most but not all spaces occupied, any driver willing to pay can find a spot, reducing cruising without creating underuse.

This approach to street parking is sometimes called performance pricing, because instead of choosing a price and seeing what happens to occupancy, the city chooses a performance standard (e.g., one or two spaces always open) and lets the price adjust to achieve it.

---

*Michael Manville is Assistant Professor in the Department of Urban Planning at UCLA's Luskin School of Public Affairs (mmanvill@ucla.edu). Daniel G. Chatman is Associate Professor in the Department of City and Regional Planning at the University of California, Berkeley (dgc@berkeley.edu).*

Performance pricing for parking is similar to congestion pricing for roads: both use prices to “clear the market” and prevent the overuse of scarce infrastructure. Like road congestion pricing, performance-priced parking is rare. Most cities prefer to keep roads and parking free, even though cities that have experimented with congestion tolls have seen remarkable results. When London implemented congestion pricing in 2004, the price of driving into central London went from 0 to £5. Traffic volumes fell 25 percent the very first day of tolling. Results were similar in Singapore, where traffic volumes fell 44 percent in the first year of tolling, and in Stockholm, where traffic fell over 10 percent. Vehicles

in toll lanes on California’s performance-priced State Route 91 zip along unencumbered by congestion, even as vehicles in the nearby free lanes sit mired in traffic. In all cases, as the price goes up, congestion goes down. Could market-priced parking do the same thing?

In 2011, San Francisco decided to find out, by creating a market-priced parking pilot program, called *SFpark*, in its downtown. *SFpark*’s explicit goals were to reduce cruising (its slogan was “live more, circle less”), increase the speed and reliability of transit, and make walking and cycling safer. For researchers, *SFpark* provided a real-world test of performance pricing. Would raising the price for parking nudge occupancy down and vacancy up in one of America’s densest and most congested cities?



#### ABOUT SFPARK

Prior to *SFpark*, meter rates in San Francisco were like those in most cities. They varied by neighborhood, but not by time of day or day of week. Prices were rarely high enough to generate turnover, and often much lower than off-street rates. In the downtown, the highest on-street price was \$3.50 an hour, while the median off-street price was \$10 an hour. This disparity created curb shortages and gave drivers strong incentives to cruise. The San Francisco Municipal Transportation Agency (SFMTA) compounded this problem by rarely changing the rates. And when the SFMTA did raise prices, it usually did so to raise revenue, not to improve parking. There was no fixed timetable for reviewing meter rates, nor any formula for changing them. And, of course, increasing rates was rarely popular and often laborious, because most of the meters were older, coin-operated devices.

*SFpark* changed these conditions. Using modern equipment, the program made prices more responsive to demand, and made price changes more transparent and predictable. And unlike many public initiatives, which get launched with fanfare and then fade from view before anyone can scrutinize them, *SFpark*’s planners displayed an admirable commitment to openness and analytical rigor. The SFMTA selected eight “treatment” neighborhoods and four control neighborhoods. In both areas, it replaced thousands of coin-operated meters

with digital “smart” meters that allowed credit card and remote payment. The agency also placed magnetic sensors in the pavement to measure parking occupancy. Together the sensors and meters relayed information wirelessly to the SFMTA, allowing the agency to correlate prices with occupancy. All these data were available to the public.

Once the new equipment was installed, the city began gathering data and also relaxed the parking time limits. On some blocks the city allowed parking for up to four hours, and on the remaining blocks it eliminated time limits altogether. Finally, in late spring 2011, the SFMTA used its new data to set meter rates in the treatment neighborhoods. The new rates varied by block, by time of day (morning, midday, and afternoon “timebands”), and by day of the week (weekday versus weekend). The price adjustments were based on the average occupancy for each timeband on each block over the course of six to eight weeks’ worth of sensor data. Prices for any of the three timebands on a block could rise or fall depending on the calculated occupancy levels (Table 1). Thus if a block was congested in the morning but vacant in the afternoon, the morning rate rose while the afternoon rate fell.

In short, *SFpark* replaced an opaque system of rates that changed infrequently and by whole neighborhoods with a more transparent system where prices changed over smaller units of time and space. It also provided something close to an experiment in priced parking. *SFpark* gave researchers the classic “before-and-after, within-and-without” research design: we could examine conditions on blocks that received variable priced parking before and after *SFpark*, and compare these to conditions on similar blocks that were never “treated” with variable pricing.

#### DID SFPARK WORK?

Performance pricing is intended to reduce cruising, and cruising is notoriously difficult to measure—it is hard to look at a car moving in traffic and know if it is searching for parking. However, cruising is caused by a shortage of street parking, and shortages *can* be measured, through occupancy and vacancy rates (which are simply the share of spaces that have vehicles in them, and the share that don’t). Thus one way to evaluate *SFpark* is to see if shortages became less common on treated blocks—if these blocks were more likely to have at least one open space.

Here is where things get tricky. *SFpark*’s meters and sensors can measure average occupancy. Drivers, however, can respond to price increases in ways that may not change average occupancy. As prices rise, more vehicles could park for shorter periods of time. This higher turnover could help local businesses, but need not alter average occupancy (and might even increase local traffic). Drivers could also respond to higher prices by carpooling. Carpooling would change *vehicle* occupancy but not necessarily change *parking-space* occupancy. And, of course, some drivers might respond to higher prices by choosing not to pay. When subway fares rise, some people pay more, some people ride less, and some people jump the turnstile. Drivers may be no different. If higher prices just encourage meter evasion or double parking, then price changes may have little impact on occupancy or vacancy.

*SFpark*’s meters and sensors could not track many of these changes. The SFMTA could not rely on its meter and sensor data to calculate vehicle turnover or the parking duration. Sensors also cannot tell if drivers are double-parking or carpooling, and cannot distinguish between types of nonpayment. Some nonpaying drivers are simply scofflaws, while others have credentials, such as disabled placards or government tags (acquired legally or illegally) that let them avoid payment.

**TABLE 1**

Criteria for Parking Rate Changes, *SFpark*

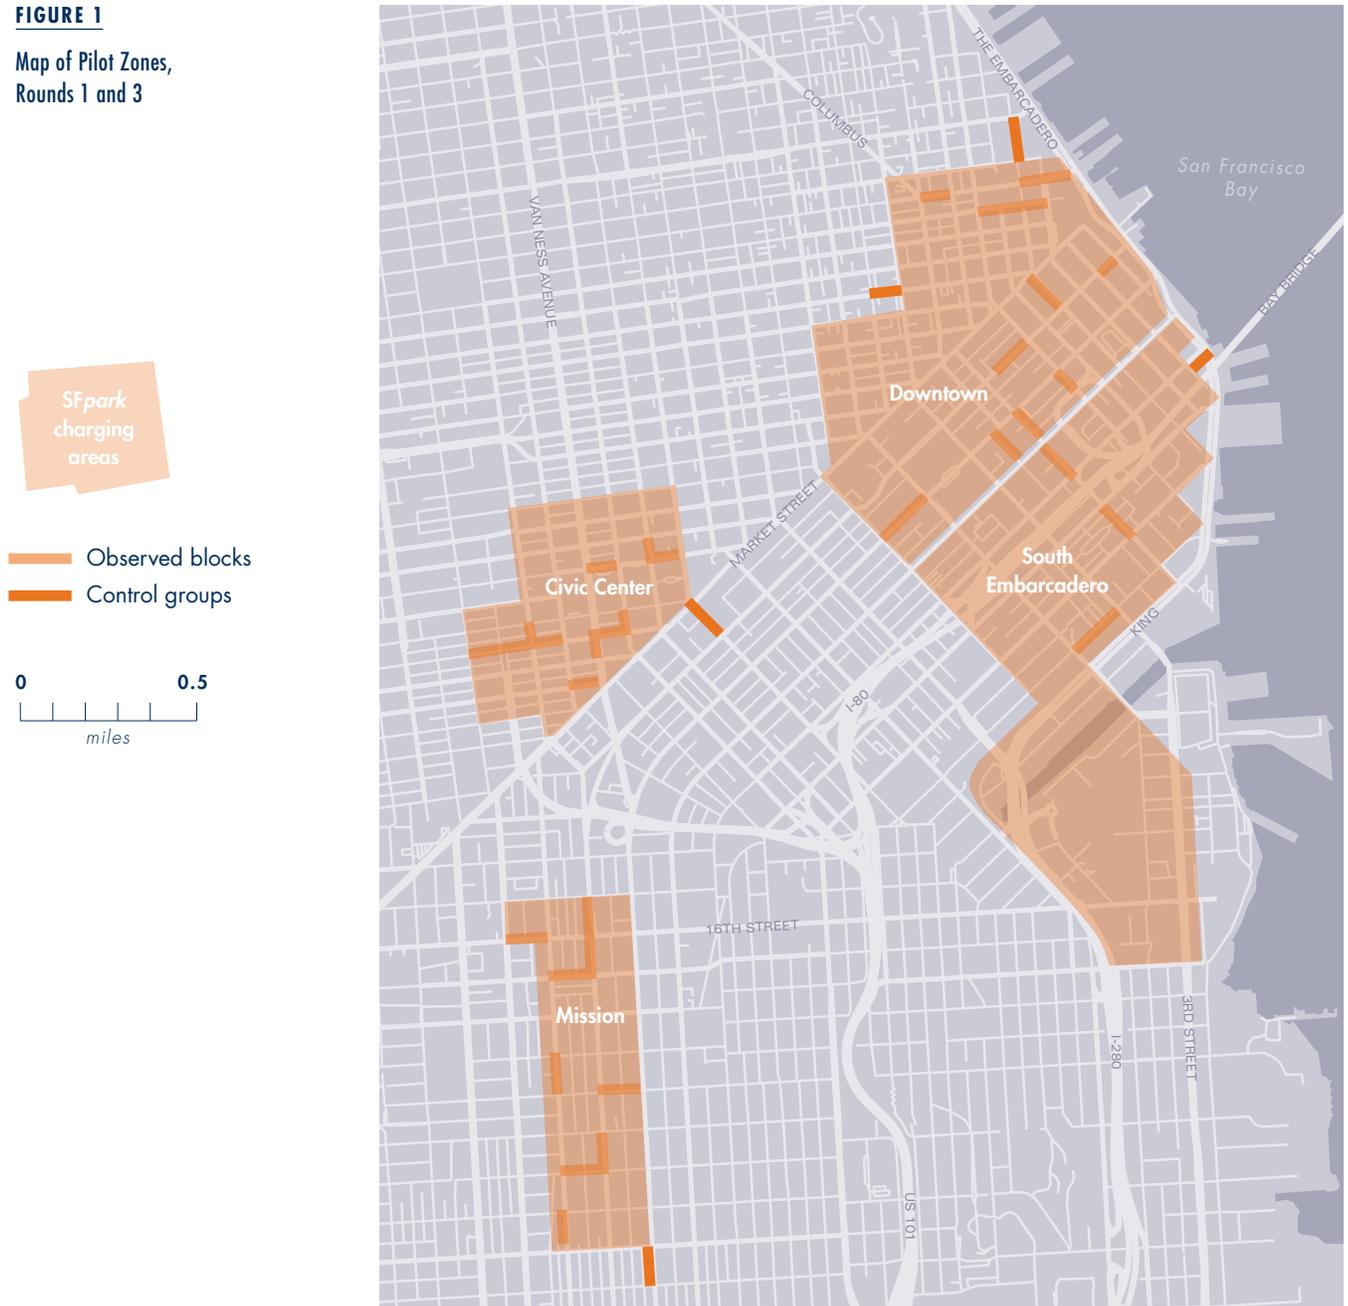
AVERAGE BLOCK-SIDE OCCUPANCY	RATE CHANGE PER HOUR
Under 30%	–\$0.50
30–60%	–\$0.25
60–80%	No change
80–100%	+\$0.25

## EVALUATING SFPARK'S PRICE CHANGES

In our study of *SFPark*, we wanted to observe all of these behaviors. The best way to do so was to pay research assistants to stand on the streets all day and have them watch drivers park. (Yes, it was tedious; we paid well.) We selected about 40 block sides in the treatment zones and 9 “control” block sides nearby (Figure 1). Because we were interested in pricing’s impact on cruising, we concentrated on blocks where occupancy was often high. We then observed each block three different times, usually a week or two after *SFPark* announced price changes. Student surveyors watched and recorded cars parking while pricing was in effect, typically from 7 or 9 am until 6 pm. This continuous observation allowed us to collect not only arrival and departure times for vehicles at individual meters,

**FIGURE 1**

Map of Pilot Zones,  
Rounds 1 and 3



but also data on vehicle occupancy, double parking, and nonpayment. We observed 13,431 parking sessions during three rounds of observation over one year.

We found that when prices rose on a block its average occupancy rate fell. This result was encouraging—exactly as *SFpark* had intended. Average occupancy, however, is only one way to measure parking availability, and may not be the best one, particularly if the average occupancy gets measured over the course of many weeks (as it did in *SFpark*). A potentially better metric is minimum vacancy: the share of minutes that a block has at least one space open. When we analyzed minimum vacancy rates, we found that price changes had no effect. We also found no statistical association between price changes and carpooling, or price changes and vehicle turnover.

How can we make sense of these results? Nonpayment seems to be part of the answer, but not a huge part. The larger issues, we think, are twofold. First is the crucial difference between average occupancy and minimum vacancy. *SFpark* raised prices only if the average occupancy was over 80 percent. Our favorite way to think about this, which we described for CityLab when we first conducted this analysis, is as follows: suppose you have a block with ten spaces and observe it for three hours, meaning there are 1,800 total possible minutes of parking on the block. If 1,200 of those minutes are occupied, the average occupancy rate is 67 percent, and the price should not change. But this figure indicates nothing about how those 1,200 minutes are distributed. They could be spread evenly across the three hours, implying that three spaces are always empty, or they could be two straight hours of zero vacancy followed by one hour of complete vacancy.

Now think about how this disparity between average occupancy and minimum vacancy could widen as occupancy gets calculated over longer periods of time. A block with an average morning occupancy of 67 percent for the month could contain hundreds of hours with no vacancies at all. The pricing mechanism can achieve the “right” average occupancy without attaining a consistent minimum vacancy. This is a problem, because drivers search for vacancies, not average occupancies.

The second problem is that *SFpark* was not an example of “true” congestion pricing, in that prices did not closely match changes in demand. Compared to standard approaches for pricing parking, *SFpark* was certainly using a market mechanism. Compared to most other markets, however, *SFpark* remained tightly controlled. Look back at Table 1: the SFMTA limited the size and frequency of price changes. Rates changed once every eight weeks, and rates could neither increase by more than 25 cents per hour nor decrease by more than 50 cents per hour each time. Finally, the agency imposed a price floor of 25 cents, and a price ceiling of \$6.00 per hour. So a block that started out \$1.00 below its optimum level would take eight months to reach its market-clearing price (assuming nothing else changed) and blocks where the price should have been \$6.50, or zero, would never reach their correct prices.

*SFpark*, in short, was an example of price-controlled performance pricing. Because the price had a cap, it may not have risen enough to actually create consistent vacancies in some areas. On blocks with high parking demand, rather than “clearing the market,” rising prices might have simply attracted drivers who were willing to pay more. As a result, in high-demand areas, rising prices may have changed the composition of parkers rather than created more vacancies. In principle this problem could be solved over time, if the price catches up to demand. But because there are caps on both the price level and the size of price changes, that is a big *if*. We cannot measure the queue on blocks without vacancies, but if they are large, prices may not be able to rise to clear them.

The pricing mechanism can achieve the “right” average occupancy without attaining a consistent minimum vacancy. This is a problem, because drivers search for vacancies, not average occupancies.



We do not fault the designers of *SFpark* for these decisions. It is always easier to criticize a program after the fact than it is to design and deliver that new program in the first place. And trying to overcome the obstacles we list certainly has its own challenges. For example, winning permission to let prices truly float would have been difficult, perhaps even impossible. Nor is it obvious that more frequent or larger price changes would be administratively possible, or even desirable. A price that truly keeps at least one space vacant might fluctuate a lot. With more frequent or larger price changes, the benefits of increased vacancy might be outweighed by the unpredictability such a system could create for drivers. Drivers who arrive at their regular spots and find that the price had doubled might get discouraged and circle the block looking for a better deal—exactly the behavior *SFpark* was designed to prevent.

The primary takeaway from our research is that performance pricing will always have to navigate a tension between the effectiveness of a price (does it actually create vacancies?), the stability of a price (how often does it fluctuate?), and the political acceptability of the price (is it so high that the public revolts, leading to no pricing at all?). Because this balance is most difficult to strike in the highest-demand areas, which are the areas most likely to generate cruising, the benefits of pricing programs may not be as large as were originally hoped. Nevertheless, the benefits *are* large, and *SFpark* was a beginning, not an end. Policymakers and academics alike should work to expand and improve upon San Francisco's valuable work. ♦

This article is adapted from "Theory versus Implementation in Congestion-Priced Parking: An Evaluation of *SFpark*, 2011–2012," published in *Research in Transportation Economics*.

#### FURTHER READING

Daniel G. Chatman and Michael Manville. 2014. "Theory versus Implementation in Congestion-Priced Parking: An Evaluation of *SFpark*, 2011–2012," *Research in Transportation Economics*, 44(1): 52–60.

Eric Jaffe. 2014. "Does San Francisco's Smart Parking System Reduce Cruising for a Space?" *CityLab*, June 25.

Michael Manville. 2014. "Parking Pricing," in Stephen Ison and Corinne Mulley (Eds.), *Parking: Issues and Policies (Transport and Sustainability, Volume 5)*, Emerald Publishing: 137–155.

Adam Millard-Ball, Rachel Weinberger, and Jeffrey Hampshire. 2014. "Is the Curb 80% Full or 20% Empty? Assessing the Impacts of San Francisco's Parking Pricing Experiment," *Transportation Research Part A: Policy and Practice*, 63: 76–92.

Gregory Pierce and Donald Shoup. "Getting the Prices Right: An Evaluation of Pricing Parking by Demand in San Francisco," *Journal of the American Planning Association*, 79(1): 67–81.