Essays in Environmental Economics

by

Kathleen Courtney Foreman

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Agricultural and Resource Economics

in the

Graduate Division
of the
University of California, Berkeley

Committee in charge:

Professor Maximilian Auffhammer, Chair
Professor Meredith Fowlie
Professor Severin Borenstein

Spring 2013

Essays in Environmental Economics

Abstract

Essays in Environmental Economics

by

Kathleen Courtney Foreman
Doctor of Philosophy in Agricultural and Resource Economics

University of California, Berkeley

Professor Maximilian Auffhammer, Chair

Policy makers are increasingly choosing market-based policies over command and control options. In this dissertation, I explore two instances of policy choices in environmental economics: a relatively novel market-based approach to handle congestion in transportation policy; and, government provision and control in the arena of water policy.

The first chapter estimates the traffic volume and travel time effects of the recently implemented road congestion pricing on the San Francisco-Oakland Bay Bridge. I employ both a difference-in-differences and regression discontinuity approach to analyze previously unexploited data for the two years spanning the price change and obtain causal estimates of the hourly average treatment effects of the policy. I find evidence of peak spreading in traffic volume and significant decreases in travel time during peak hours. I also find suggestive evidence of substitution to a nearby bridge that is not subject to congestion pricing. In addition, I show significant decreases in travel time variability. Using my results, I calculate own- and cross-price elasticities for trips due to the toll change and include back-of-the-envelope calculations for the welfare effects of the policy.

The second chapter I explore the impact of government water delivers in California's Central Valley. California's agricultural sector receives large quantities of irrigation water from the federal and state water projects. In recent years, there have been significant restrictions on these deliveries due to droughts and regulation to protect endangered species. This chapter empirically tests the hypothesis that higher deliveries to water districts in a given county lead to higher agricultural employment and cropped area and provides point estimates of this effect and uncertainty around the estimates. The results show robust evidence of a statistically and economically significant impact of irrigation water deliveries on employment and area. This effect is robust to different definitions of employment, alternate control groups, and different windows of data.

To Patricia and Glen for a lifetime of goodness.
To Sarah, David, and Buzz for constant kinships through thick and thin.
To Christine for teaching me how to read and how to be.

Contents

List of Figures

List of Tables

Acknowledgments

I would like to thank my dissertation chair, Maximilian Auffhammer, for his unwavering support, wise mentorship, and unending patience–he is a true all-star among advisors.

I would also like to thank my other committee members, Meredith Fowlie and Severin Borenstein for their invaluable feedback and advice.

I thank the co-authors of my second chapter: Maximilian Auffhammer (again) and David Sunding for their guidance and support.

I also thank Michael Anderson, Peter Berck, Lucas Davis, Jeff Perloff, Elizabeth Sadoulet, Sofia Villas-Boas, Catherine Wolfram, Kyle Emerick, Marshall Burke, Lydia Ashton, Jessica Rider, Anna Spurlock, as well as the participants of the Environmental and Resource Economics Seminar at the Agricultural and Resource Economics of the University of California Berkeley, and conference participants at the WEAI and Camp Resources for their helpful comments.

I thank Jeff Gerbracht from the Metropolitan Transportation Commission for providing me with data and other toll-related information.

Chapter 1

Crossing the Bridge: The Effects of Time-Varying Tolls on Curbing Congestion

In 2010, urban road congestion cost the US $101 billion, including nearly two billion gallons of wasted fuel and 4.8 billion hours of travel delay[1]. To address the problem, the US Department of Transportation has called for measures such as road congestion pricing to be implemented to help alleviate the problem. While theory suggests that congestion pricing should work, there is little empirical evidence as to the magnitude of change we should expect from and the effectiveness of pricing programs. In this paper, I provide point estimates of the causal effects on traffic volume and travel time from the recent implementation of congestion pricing on one of the main toll bridges in the San Francisco Bay Area. I also show reductions in travel time variability, calculate the ranges for price elasticities, and conduct back of the envelope calculations of the welfare impacts of the policy.

Pricing traffic congestion is based on the economic theory of negative externalities: the driver of a vehicle imposes costs on all the other drivers on the road and, due to the open-access nature of government-funded roads, the driver does not internalize those costs; this results in an inefficiently high number of drivers using the road. In theory, the optimal corrective policy is a tax equal to the marginal social cost imposed by the additional driver, which makes drivers internalize the external costs. The optimal fee should therefore be a time-varying fee that depends on real time traffic conditions. The pricing increase on the San Francisco Bay Bridge was primarily implemented to raise funds to pay for bridge maintenance and seismic retro-fitting, but it is in fact a peak time pricing program. While it falls short of being an optimal fee, the time-varying congestion pricing on the bridge studied in this paper does reflect the higher marginal social damage during times of heavier traffic, and it is one of the few attempts to implement time-varying congestion pricing on roads in the United States. There is a wide body of literature focused on the theory of externalities and congestion, but existing empirical studies of road congestion are scarce due partially to limited instances of actual implementation.

In this paper, I test the hypothesis that congestion pricing decreases traffic volume and results in a shorter travel time during peak hours on the bridge. In addition, I investigate the possibility that marginal drivers switch to traveling during off-peak times or substitute to an alternate route that lacks congestion pricing, and that travel time variability on the Bay Bridge during peak hours decreases. If the pricing plan is successful in altering behavior, it could have important consequences for use in policy-making to stem the costs associated with road congestion. Furthermore, because road use is linked to other outcomes (e.g. air and noise pollution, traffic accidents, and petroleum use and importation), policies that affect driver behavior could have consequences in other areas[2]

[1]Schrank, Lomax and Eiseleand (2011) and USDOT (2008)

[2]For example, Currie and Walker (2011) demonstrate that road congestion contributes to poor infant health, and easing road congestion can reduce prematurity and low birth weight. Levy, Buonocore, and von Stackelberg (2010) estimate that the value of small particulate matter related mortality attributable to congestion is $31 billion annually. Schrage (2006) shows the incident and costs of accidents are higher with higher congestion. Dickerson, Peirson and Vickerman (2000) show that the accident externality increases substantially with high traffic volumes. Treiber and Kesting (2008) show that traffic congestion

I use two methods for estimating driver response: difference-in-differences and regression discontinuity. Using a difference-in-differences approach and two years of data, I find that the 50% increase in toll price during peak hours in my study results in statistically significant decreases of 312 vehicles per hour (4.4%) during peak hours as well as increases in traffic volume hours just before and just after peak periods. In addition, I find a statistically significant decrease in travel time of between two to six minutes (10% to 24%) during peak hours. Finally, I find evidence consistent with a small amount of morning peak substitution to an alternate bridge that lacks congestion pricing. In order to estimate short-term effects of the policy, I use a regression discontinuity approach to analyze two months of data. I find the immediate response was a decrease in peak traffic volume of 639 vehicles per hour (9%), and a decreased travel time of 4.4 minutes (21%), both of which are statistically significant at the 1% level. In addition, I find statistically significant decreases in the standard deviation of travel time during peak hours of up to 5.4 minutes.

I use my point estimates of decreases in traffic volume to calculate price elasticities, and I find the peak hour own-price elasticities to be around -0.084, and ranging from 0 for mid-peak hours to -0.2 for shoulder peak hours. I find the off-peak response to peak-hour price increases to be around 0.034, and ranging from 0 for non-shoulder off-peak hours to 0.38 for off-peak shoulder hours. I do back-of-the envelope welfare calculations, and find the value of time saved to be between $1-28 million annually; the value of fuel savings to be around $1 million annually; and, the amount of CO_2 abated to be 2240 tons per year, which is equivalent to taking 480 commuting cars off the road.

This paper is organized as follows: section 1.1 provides some background on existing work; section 1.2 explains the empirical design; section 1.3 describes my data; section 1.4 shows my results; section 1.5 discusses the policy implications and welfare impacts; and, section 1.6 suggests directions for future work and concludes.

1.1 Background

The economic literature is rich with theory on congestion pricing that dates back to Vickrey (1963) who originated the theory of policies to handle road congestion and pointed out the need for different prices during peak and off-peak hours. Small (2005) noted the open-access problem inherent in our publicly-provided (and under-priced) road systems, and suggest that charging an optimal fee to access the roads could encourage a more efficient allocation. Baumol and Oates (1988) suggest that the optimal solution to congestion is to impose a Pigouvian tax (equal to marginal damage) on the drivers, who are the producers of congestion externalities. Small (1983) investigates the distributional impacts of congestion tolls using a modal choice equilibrium model and survey data, and concludes that “in almost all cases, the net result is benefits for all income groups”. Boardman (1977) model the speed and flow relationship of traffic and test the

can increase fuel consumption by 80% and travel time by up to a factor of 4.

model using data from a limited access highway. They conclude that congestion prices must vary by time of day to encapsulate the hour-dependent private and social costs of driving. Using numerical simulations of a spatial general equilibrium model, Anas and Rhee (2006) compare congestion pricing and urban boundaries, essentially comparing a price regulation to a standard, as a means of reducing congestion and urban sprawl, and find that congesting pricing is first best and that using urban boundaries actually reduce welfare. Pricing negative congestion externalities is going to become more important with the introduction of increasingly stringent Corporate Average Fuel Economy (CAFE) standards, which will likely decrease the private per mile cost of driving, and consequently increase congestion.

McFadden (1974) emphasizes the importance of proper planning and policy in developing urban transportation systems, and he notes that travel decisions have many dimensions, including purpose, timing, and mode, all of which are important considerations for pricing road congestion. McFadden (1974) also provides evidence that travel time is valued linearly and increasing in the wage rate; additionally, he presents suggestive evidence that more salient costs associated with driving, such as tolls, appear to be weighted more heavily than less salient driving costs, such as mileage and maintenance costs. However, recent work done by Finkelstein (2009) shows that collecting tolls electronically can have the effect of rendering tolls less salient, and that the adoption of electronic toll collection makes the short run toll price elasticity of driving more inelastic.

Elsewhere in the economics literature, non-pecuniary attempts to decrease traffic and congestion have been shown to be ineffective. For example, Davis (2008) shows that a one-day-a-week ban on driving a vehicle in Mexico City, which was a policy aimed at reducing air pollution that has been replicated in several locations in South America, did not decrease air pollution, and resulted in more, and disproportionately higher-emissions, vehicles being driven. The author concludes that the restrictions were unsuccessful in prompting drivers to reduce usage of private vehicles, which suggests that if such a policy were implemented to curb congestion, it would be equally unsuccessful. Additionally, expanding road capacity is found to be an ineffective way to reduce congestion, as found by Duranton and Turner (2011), who echo previous findings that peak traffic grows to fill the maximum available capacity, in short: if you build it, they will come.

Time-varying congestion pricing on roads had previously been implemented in only four other locations in the US (see table A.1): in 1995 in Orange County, California; in 1998 in San Diego, California; in 1998 in Lee County, Florida; and, in 2006 a short test run was done in Oregon. In San Diego, non-HOV drivers can use the HOV lanes of Interstate 15 for a price that varies dynamically with traffic demand with the goal of keeping the HOV lanes flowing at full speed, and the time saved is estimated to be about one hour for each round trip made by a vehicle. In Orange County, lanes running parallel to State Route 91 for 10 miles can be used by drivers by paying a predetermined price by time of day and day of week, the amount of which is adjusted every three months and is based on previous demand. The reports on the effectiveness of this program maintain that the speeds in priced lanes during peak hours are about 60-65 mph, while speeds in unpriced lanes are 15-20mph. The Florida pricing scheme was introduced on two Lee County bridges and offers a 50% discount on travel in pre- and post-peak periods; the program is reported to have shifted 5% of travel from peak to off-peak times. Since the Bay Bridge pricing policy was introduced, a few other areas in the US have introduced various forms of congestion pricing, including Seattle and San Jose.

The limited number of implemented congestion pricing schemes means there are not very many opportunities for economists to study the impacts, and in fact, there is little in the empirical economics literature about the instances of congestion pricing. One notable exception is Small (2005), which uses survey data and revealed and stated preference models to study drivers' preferences around the pricing of route 91 in Orange County. The authors model drivers' decisions around acquiring an Electronic Toll Collection (ETC) tag, using priced or unpriced lanes, and carpooling or driving solo. They find that people have a median value of travel time of \$21.46/hr, and a median value of travel time reliability of \$19.56/hr. Another study, Brownstone et al (2003), uses revealed preference data from Interstate 15 in San Diego and finds the willingness to pay to reduce travel time to be \$30 per hour.

Congestion pricing on roads has been implemented in several locations outside of the US, including London, Singapore, Stockholm, Norway, and Milan. Reported outcomes from policies in these cities include reductions in vehicle miles traveled (VMT), decreases in traffic and travel times, reductions in accidents and emissions, and increases in vehicle speeds.

The reports on these policies are not in the economics literature and rely mainly on comparing the unconditional means of outcome variables before and after policy changes to evaluate the impact and success of congestion pricing. For example, the report on Milan's policy claims the policy resulted in a reduction of 12% in traffic in 2008, the year the policy was implemented. However, concurrent factors (such as the global financial crisis and oil prices rising above \$100/barrel) were not taken into account. In subsequent years, Milan found that traffic increased back to pre-policy levels, which brings into question the effectiveness of the policy. The report does not attempt to disentangle the traffic reductions due to the policy and the reductions due to confounding factors.

A recent report on the Bay Area toll changes (Deakin and Frick (2011)) uses a small data set (as low as 8 observations in some cases) and compares the hourly means before and after the policy change (sometimes comparing observations of before and after in different months of the year). The report finds that there was a small and statistically insignificant decrease in overall morning peak traffic of 1.4% and evening peak decrease of 0.1%. Broken down by hour, they find statistically significant increases in traffic volume during some off-peak shoulder hours, but they find mixed results of increases and decreases in traffic volume during peak hours, which seems to suggest that the pricing policy might not be working as intended to decrease congestion during peak hours. They also find decreases in travel time during some morning peak hours, and a small decrease in travel time during the evening peak, although this is somewhat perplexing, given that the travel time decreases happen during some of the hours where they find both insignificant increases and decreases in traffic volume.

A simple comparison of means before and after policy changes and the overall dearth of evidence of the causal effects of road congestion pricing is unhelpful for policy-makers hoping to use congestion pricing to curb traffic problems. To address this lack of evidence, I employ large and previously unexploited data sets and use an empirical econometric design to estimate the causal effect of congestion pricing on traffic volume and travel time. My work builds on the established theory of congestion pricing and fills some of the existing knowledge gap in the effects of the practice of the few congestion-targeted policies that have been implemented.

1.2 Empirical Design

1.2.1 Quasi-Experimental Setup

In the San Francisco Bay Area in California, there are three bridges that link communities in the East Bay with the San Francisco peninsula. The San Francisco-Oakland Bay Bridge (hereafter Bay Bridge) is the main bridge connecting downtown San Francisco with smaller cities in the Easy Bay (see figure 1.1) and is used heavily by drivers commuting for work and traveling for leisure activities; it is used by roughly 124,000 vehicles per weekday in each direction. The two smaller bridges, the San Mateo and the Dumbarton Bridges, lie 18 and 25 miles south of the Bay Bridge, respectively. The San Mateo Bridge carries about 46,000 vehicles per weekday per direction, while the Dumbarton Bridge carries about 31,000 vehicles. Besides those three bridges, the main option for people crossing the Bay is to take the light rail system (Bay Area Rapid Transit, or BART). Very few people take the ferry, and two roundabout routes could be taking two bridges (the Richmond and Golden Gate Bridges) or driving south around the Bay. These latter two options would require most drivers to pay two tolls, drive over 50 miles out of their way, or both. Thus, I assume they are not reasonable substitutes for the Bay Bridge and I consider the ways that drivers can substitute away from the three bridges I study in this paper to be limited.

On all three bridges, tolls are only collected on vehicles crossing westward onto the San Francisco peninsula. At all times, there are dedicated lanes for Electronic Toll Collection tag holders, known in the Bay Area as "FasTrak" tags. Using a FasTrak lane decreases the delay from toll collection by allowing drivers to slow to 25mph instead of stopping completely to pay the toll. Also available at all times are cash lanes, which allow drivers to pay with cash or pay with their FasTrak tags. Cash lanes move much slower than FasTrak lanes, as drivers who are using cash must stop completely to pay the toll, thus decreasing vehicle throughput. During peak hours only (weekdays from 5am to 10am and from 3pm to 7pm), there are dedicated lanes for carpools[3]. During off-peak hours, the

[3]Carpools are three or more people on the Bay Bridge and two or more people on the San Mateo and Dumbarton Bridges. Other vehicles that are allowed to use the carpool lanes are motorcycles and

Figure 1.1: Map of the Bay Area Bridges

Table 1.1: Summary of toll charges on weekdays before and after July 1, 2010

	Before		**After**			
			Peak[a]		Off-peak	
	Regular	Carpool[b]	Regular	Carpool[b,c]	Regular	Carpool[d]
Bay Bridge	$4	$0	$6	$2.50	$4	N/A
San Mateo Bridge	$4	$0	$5	$2.50	$5	N/A
Dumbarton Bridge	$4	$0	$5	$2.50	$5	N/A

[a] Peak hours are weekdays from 5am to 10am and from 3pm to 7pm.
[b] Three or more people required in carpools on the Bay Bridge and two or more people required in carpools on the San Mateo and Dumbarton Bridges.
[c] FasTrak tag required to access carpool lanes after the change.
[d] There are no carpool lanes during off-peak hours.

carpool lanes are either closed, or they revert to bus, FasTrak or cash lanes. There is no toll for any vehicle on the eastbound trip.

Prior to July 1, 2010, the toll was the same for all three bridges at all times of the day and was a flat $4/vehicle for vehicles with 2 axles (vehicles with more axles face higher rate). Carpools using the carpool lanes during peak hours were allowed to use the bridge toll-free.

On July 1, 2010, the tolls on all seven Bay Area Bridges were increased to raise funds to pay for maintenance, transport projects, and seismic retro-fitting. The tolls on the San Mateo and Dumbarton Bridges[4] increased uniformly by 25% from $4 to $5 per vehicle, with carpools now required to pay a toll during peak hours, albeit at a reduced rate of 50% of the regular toll: $2.50 per vehicle. In addition, since the change, carpools are required to carry FasTrak tags in order to take advantage of the lower rate. There is no carpool discount during off-peak hours. Also on July 1, 2010, the Bay Bridge changed its rates, but implemented non-uniform congestion pricing: $6 per vehicle during weekday peak hours; $4 per vehicle during weekday off-peak hours; and, $5 per vehicle on weekends. The carpool vehicles on the Bay Bridge are treated the same as on the other bridges, that is, they are now required to carry a FasTrak tag and pay $2.50 per vehicle during peak hours, with no discount off of the $4 toll during off-peak hours. Table 1.1 shows a summary of the tolls on weekdays before and after the change.

The non-uniform change in pricing provides a quasi-experimental setting that allows me to empirically test the hypothesis that peak time congestion pricing reduces traffic volume and travel time, and travel time variability.

---

vehicles with DMV-issued Clean Air decals. I refer to them collectively as "carpools".

[4]The same toll increase took place on four others: Antioch, Benicia, Carquinez, and Richmond Bridges.

1.2.2 Empirical Method

I test the hypothesis that congestion pricing affects traffic volume and travel time over the Bay Bridge during peak hours by using a difference-in-differences (DD) approach with day of week and week of sample fixed effects. Using my hourly-level data, I compare traffic volumes and travel times during each hour of the day from 1:00am to 11:00pm before and after July 1, 2010 on the Bay Bridge. I use the midnight hour before and after the policy change as a control. This allows me to identify the causal effects of the congestion pricing on peak and off-peak (non-midnight) traffic volume and travel time.

In my DD estimating equation 1.1, Y_{hdw} represents either traffic volume or travel time during hour h on day of the week d in week of sample w. $after_w$ is a binary indicator equal to one if the week of the observation is after the policy change on July 1, 2010. When $after$ is interacted with a coefficient for each hour of the day (1-23), as in $\boldsymbol{\beta_h} after$, the resulting $\boldsymbol{\beta_h}$'s are my coefficients of interest, that is, my estimates of the hourly average treatment effect of the pricing policy. My estimating equation is:

$$Y_{hdw} = \alpha_0 + \boldsymbol{\alpha_h} + \beta_0 after_w + \boldsymbol{\beta_h} after_w + \boldsymbol{\gamma_d} + \boldsymbol{\theta_w} + \varepsilon_{hdw} \tag{1.1}$$

Equation 1.1 includes a set of fixed effects for: hours of the day from 1:00am to 11:00pm, $\boldsymbol{\alpha_h}$; day of the week, $\boldsymbol{\gamma_d}$ (excluded category is Monday); and, week of sample, $\boldsymbol{\theta_w}$. For my week of sample, my excluded categories are January 1-7, 2010 and 2011. I do not use the data from the partial weeks at the beginning and end of the sample, so the first full week of sample is July 2-8, 2009 and the last is June 18-24, 2011. I also do not use the data from the week the policy change took place, as it was part way through a week. Thus, I use a total of 102 weeks of data, 51 before the policy change and 51 after. Finally, ε_{hdw} represents an unobserved disturbance.

Applying a difference in difference approach is predicated on two key identifying assumptions: stable unit treatment value assumptions (SUTVA); and, absent treatment, traffic volume and travel time during peak and off-peak (non-midnight) hours have similar trajectories as the midnight hour. For example, populations changes or weather shocks might affect traffic volumes and travel times in both non-midnight and midnight hours similarly. Presuming these assumptions hold, I am able to infer causality from my point estimates. I discuss these assumptions and test their validity in the results section.

1.2.3 Predictions from Theory

The most elementary economic principle at work here is the law of demand, which predicts that the toll price increase during peak hours will result in a decrease in the quantity of vehicle trips demanded during those hours. The reductions in trips during peak hours could result from: overall trip reduction from trip avoidance or trip combining; shifting the timing of trips from peak to off-peak hours, in which the toll price did not change; shifting the trip to a lower-priced route, for example, the San Mateo Bridge; or, shifting modes from driving to public transit (e.g. BART, bus). While a 50% increase in the toll from \$4 to \$6 during peak hours seems like a large increase that might result in substantial decreases in the number of trips taken, the predicted result becomes less clear when one considers the total trip cost, the availability of substitutes, and the toll salience compared with other costs of traveling.

First, many drivers use FasTrak tags to pay their tolls, which decreases the salience of the toll and would thus dampen the price effects of the toll change. One would expect that drivers paying cash would be more responsive to the price change. Anecdotally, I have spoken to a fair number of people about the toll, and many of them could not even tell me how much it is and did not know that it had changed, which suggests that their behavior would not have been directly affected by the toll change.

Second, the total cost of driving across the bridge is much higher than the toll alone. For example, someone driving from Walnut Creek (a city in the East Bay that is located 23 miles from the San Francisco Central Business District) can expect the drive to cost anywhere between \$23 and \$100 per round trip[5], meaning that the cost increase per trip due to the toll was closer to a 2-5% increase in cost instead of a 50% increase. Thus, the price response might be much lower than expected if drivers are taking into account the full cost of driving.

Finally, the availability of substitutes will impact the marginal driver's response. The prices of BART and buses did not change during this period, so one would expect some substitution toward public transit for some people living closer to transit options who were right on the edge of taking transit beforehand. Also, for the Bay Bridge, the off-peak hours did not have a price increase, so one would expect that marginal drivers who were driving during peak shoulder hours (that is, the hours on the shoulder of the peak: 5:00am, 9:00am, 3:00pm, and 6:00pm) might substitute from the peak shoulder hour to the closest off-peak shoulder hour (that is, the hours on the shoulder of the off-peak: 4:00am, 10:00am, 2:00pm, and 7:00pm). These drivers are likely to change their time only if they were already traveling close the off-peak hours. For example, someone who used to drive at 5:10am might leave a little earlier and now travel at 4:50am. Conversely, it is less likely that a driver who used to travel at 5:50am would switch to driving at 4:10am. Finally, if a driver used to use the Bay Bridge and his point of origin is located right on the edge of the middle in between the Bay and San Mateo Bridges, he might substitute toward the San Mateo during the peak hours, since it is now one dollar cheaper.

Overall, it's unclear if the congestion pricing would work to reduce traffic volume, and if it does work what the magnitude of the reduction would be. The availability (or lack) of substitute travel times and routes, the decreased toll salience due to FasTrak, and the toll being just a small part of total cost of driving may dampen the effect of the toll price

[5]Based on the following calculation: toll = \$4-6; mileage = \$9-26 (based on a 23.5 mile trip in each direction and either a mileage rate reflecting only gas use of about \$0.20/mile, which approximates the marginal cost of driving, or the IRS's standard mileage rate of \$0.555/mile allowance that is meant to include gas, repairs, and other costs of operating a vehicle, which is closer to the average cost of driving); parking in the San Francisco Central Business District = \$5-\$45 per day; and time costs of 33 minutes to an hour each way = \$5.5-\$30 depending on the value of travel time used.

increase.

1.3 Data

In this paper, I employ two panel data sets. The first data set is previously unexploited in the economics literature and measures the traffic volume, reported in number of vehicles crossing a bridge per hour. The second data set measures the travel time, reported as the median trip time to cross a bridge in a given hour. I do not observe vehicle- or driver-level data. I use data from the year before and the year after the toll change, resulting in an hourly data set for each bridge spanning from July 1, 2009 to June 30, 2011. After excluding weekends and holidays, as well as the few hours when the Bay Bridge was closed for repair, I end up with around 10,000 hourly observations per bridge, per outcome variable.

1.3.1 Traffic Volume

The traffic volume data come from the San Francisco Bay Area's Metropolitan Transportation Commission, and comprise the total number of vehicles crossing the three Bay Area Bridges controlled by the Bay Area Transportation Authority (the Bay, San Mateo, and Dumbarton Bridges) by hour, by lane, and include the lane type (described below). My outcome variable of interest is the number of vehicles that cross a given bridge (counted at the toll plaza), in a given hour, on a given date.

Lane Type Each of the bridges have three types of lanes for the purposes of toll payment: carpool-only, FasTrak-only, and cash lanes (lanes in which drivers can pay with cash, or they can use a FasTrak, but most cash lane users pay with cash, so I refer to them as cash lanes). As of July 1, 2010, vehicles using the carpool lanes must carry a FasTrak for payment. The two carpool lanes on the San Mateo Bridge revert to FasTrak-only during off-peak periods. The solo carpool lane on the Dumbarton Bridge is closed during off-peak periods. There are four carpool lanes on the Bay Bridge during peak hours, and in the off-peak hours, two of them revert to cash lanes and the other two revert to bus-only lanes.

Data on traffic volume over time on the Bay Bridge can be seen graphically in figure 1.2. The monthly mean of hourly traffic volume fluctuates between 5000 vehicles per hour to 5600 vehicles per hour, depending on the season, with traffic generally higher in the summer and lower in the winter. The percent of vehicles using the carpool lanes is fairly consistent at just under 10%, with a visible drop around the time of the policy change. The percent of vehicles using FasTrak lanes is under 40% at the begining of 2006, but increases to almost 60% by mid-2011. The percent of vehicles using cash lanes changes from under 60% to around 40%. This is important, as FasTrak lanes have a higher throughput per hour, and the increasing percent of traffic using FasTrak lanes could affect

Figure 1.2: Bay Bridge, traffic volume over time

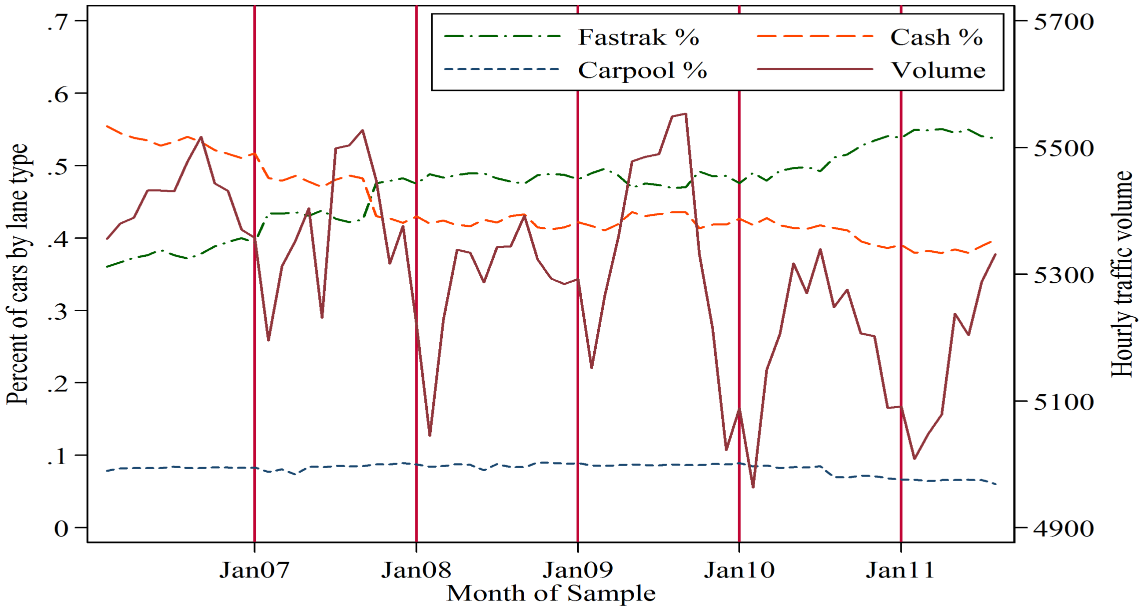

Jan 2006 - July 2011. Weekdays, Westbound. Excludes holidays.

my estimates; this is one reason I control for week of sample in my estimation. Analogous graphs for the San Mateo and Dumbarton Bridges are in the appendix; they show similar seasonality patterns and lane type usage trends.

Data on the traffic volume profile for weekdays on the three Bay Area bridges for the year before the policy change can be seen in figure 1.3. The vertical red lines in the graphs indicate the start (5:00am and 3:00pm) and end (10:00am and 7:00pm) of peak periods. The graph shoes the mean, fifth and ninety-fifth percentile of traffic volume. The solid blue lines show traffic volume on the Bay Bridge: during early morning hours, traffic volume is, on average, under about 1000 vehicles per hour. By the height of morning peak, traffic volume reaches over 8000 vehicles per hour on average. This decreases to about 6000 vehicles per hour in between the peaks, and then rises to over 7000 vehicles per hour in the evening peak, after which it tapers back down to about 2000 vehicles by 11:00pm. The spread between the fifth and ninety-fifth percentiles remains largely in the 1500-2000 range throughout most of the day, except for the morning pre-peak period, where it is mainly below 300.

On the San Mateo Bridge (red dashed lines), during early morning hours, traffic volume is very low. By the height of morning peak, traffic volume reaches around 4750 vehicles per hour on average. This decreases to about 2000 vehicles per hour in between the peaks, and then rises slightly in the evening peak, after which it tapers back down. The spread between the fifth and ninety-fifth percentiles is below 200 in the morning pre-peak period, then rises to 900 by 9:00am, then after 10:00am decreases and remains largely between

Figure 1.3: Mean, 5th and 95th percentiles of hourly traffic volume

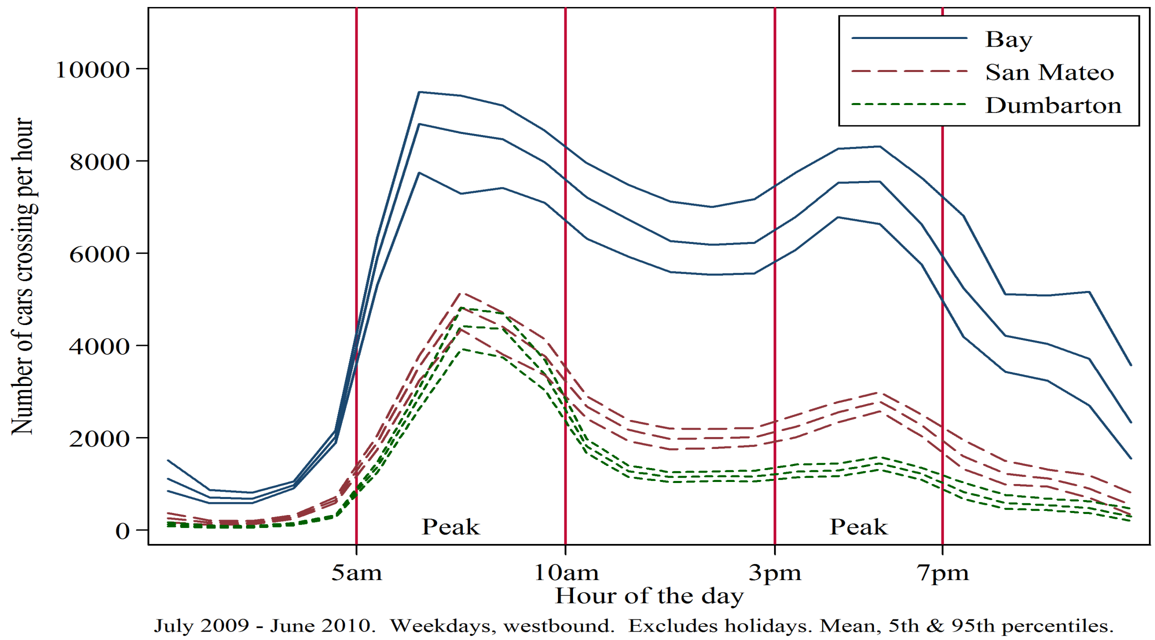

400-600. On the Dumbarton Bridge (green dotted lines), during early morning hours, traffic volume is very low. By the height of morning peak, traffic volume reaches over 4000 vehicles per hour on average. This decreases to just over 1000 vehicles per hour in between the peaks, and remains about there through the evening peak, after which it tapers back down. The spread between the fifth and ninety-fifth percentiles is below 100 in the morning pre-peak period, then rises to 1000 by 9:00am, then after 10:00am decreases and remains largely between 200-300.

The peakiness feature of the weekday volume profiles of the three bridges suggests that the bridges could be good candidates for congesting pricing during high traffic times.

1.3.2 Travel Time

The travel time data are from the Department of Transportation's Performance Measurement System (PeMS). These data provide the median time for vehicles carrying FasTrak tags to get from one FasTrak sensor to another, by hour, and provides an estimate of how quickly traffic is flowing. I use various sensor pairs to measure the time to westward on the Bay, San Mateo and Dumbarton Bridges[6]. My outcome variable of interest is the median time for FasTrak carriers to cross a bridge, in a given hour, on a given date.

[6] Median time to cross is measured from the sensor pairs: (1) Bay Bridge: the combination of four originating sensors on the East Bay and two destination sensors in San Francisco CBD, for a total of eight routes (2) San Mateo Bridge: SR 92 from Hayward to Foster City (3) Dumbarton Bridge: SR 84 from Fremont to US-101.

Figure 1.4: Mean, 5th and 95th percentiles of hourly travel time

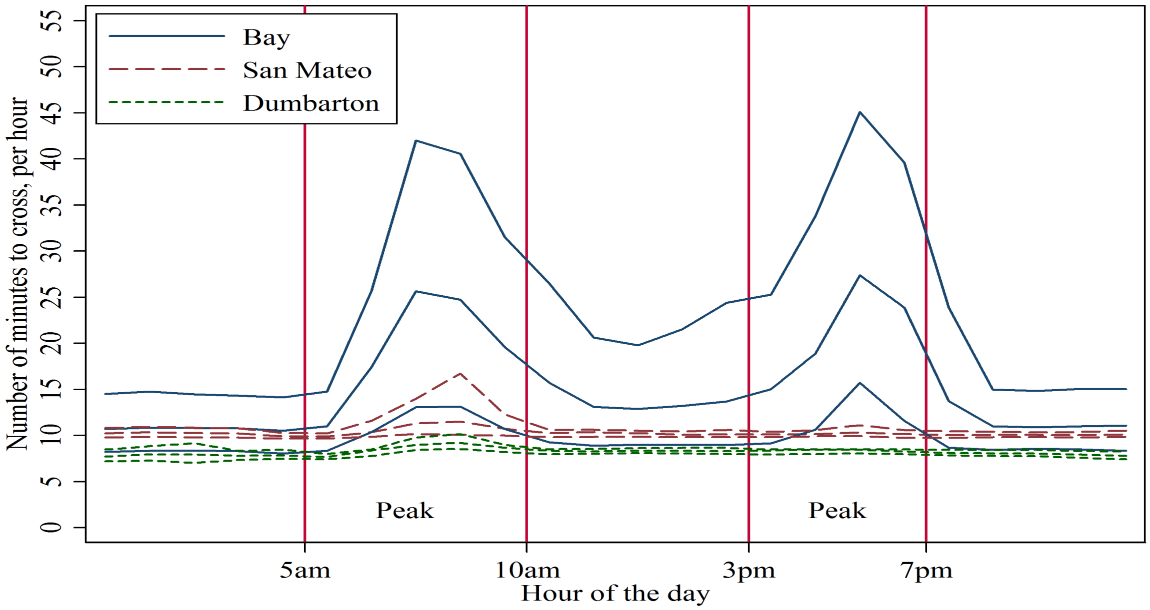

July 2009 - June 2010. Weekdays, Westbound. Excludes holidays. Mean, 5th & 95th percentile.

To be clear, my travel time data do not comprise the universe of vehicles crossing the bridge: they comprise only those vehicles carrying FasTrak tags traveling between my chosen sensor pairs, so I am only tracking the travel times of those vehicles. In principle, vehicles carrying FasTrak tags ought to be able to cross the bridges faster than non-tag holders because the FasTrak lanes have higher throughput and do not require vehicles to stop to use them, unlike the cash lanes. In addition, since the policy change, carpools are required to carry FasTrak tags, and they tend to cross the bridge even faster than FasTrak lane users because they do not need to stop at metering lights after the toll plaza, so if the number of carpools changes, it could bias my results. For example, if the number of carpoolers decreases, my results might be biased downward, since carpools tend to cross faster. However, the measurement I use is the hourly median travel time of the tracked FasTrak tags, and since carpools make up fewer than half of all FasTrak holders for an hour, any changes in the carpool should not affect my results. Moreover, part of the congestion issue with the Bay Bridge is the bottleneck that forms before vehicles can select into lane types. Hence, if a carpool or FasTrak user is stuck in the pre-bridge traffic, it will have a similar travel time to the cash lane user beside it, and the congested time of day is precisely the time that the policy is aimed at affecting. Finally, the percent of vehicles using FasTrak for payment is roughly 65%, so the travel time measurement I use captures the experience of two-thirds of the vehicles crossing.

A summary of the data on the fifth percentile, mean, and ninety-fifth percentile of median time to cross the three bridges on a weekday is shown in figure 1.4. During early morning (uncongested) hours, crossing the Bay Bridge takes (solid blue lines), on average, around 12 minutes (an average speed of 41 mph). By the height of morning peak, travel time has increased to almost 30 minutes on average (17 mph). This decreases to about 17 minutes in between the peaks, and then increases again in the evening peak, to over 36 minutes (under 14 mph), after which it tapers back down to the uncongested rate. The spread between the fifth and ninety-fifth percentiles is around 6 minutes in the morning pre-peak period, and then increases to 29 by 7:00am, decreases to about 11 by noon, and then increases again to 29 by 3:00pm, after which it decreases to about 6.5 by 8:00pm, where it remains steady for the rest of the day. The data for the San Mateo (dashed red lines) and Dumbarton Bridges (green dotted lines) show weekday travel times profiles of 10-12 minutes (60 mph) and 9-10 minutes (75 mph), respectively, with very little difference between peak and off-peak hours. For the San Mateo Bridge, the spread between the fifth and ninety-fifth percentiles is at or below 1 minute for most of the day, but spikes to 4, 7, and 2 minutes at 7:00am, 8:00am, and 9:00am, respectively. For the Dumabarton Bridge, the spread between the fifth and ninety-fifth percentiles is at or below 1 minute for most of the day, and never gets much higher than 2 minutes at any time of the day.

The weekday travel time profile graph suggests strongly that the Bay Bridge is an excellent candidate for congesting pricing, and also suggests that the San Mateo and Dumbarton Bridges are operating at or below capacity and would benefit less from time-varying congestion pricing. In addition, it shows wide variation in travel times on the Bay Bridge during most of the day, making trip times across the bridge quite unpredictable. If congestion pricing is able to decrease the travel time variance in addition to reducing travel time, there could be additional benefits to drivers, which is a further reason that the Bay Bridge is a good candidate for congestion pricing.

1.4 Results

1.4.1 Key Identifying Assumptions

As mentioned previously, the difference-in-differences approach is predicated on two key identifying assumptions: the SUTVA and similar evolution of outcomes for the treated and control groups absent treatment. I test the validity of those assumptions here.

The SUTVA, also referred to as the no interference assumption, requires that the control's outcomes are unaffected by the treatment; for this paper, it requires that the pricing change for the peak hours does not affect traffic during the midnight hours. It seems reasonable to believe the the midnight hour on the Bay Bridge is not a good substitute for, and was not affected by the toll change in the peak hours. Figures 1.5 and 1.6 reflect this assumption in the lack of change in traffic volume and travel time during the midnight hour (the lower, black line in the graphs) at the time of the policy change. Statistical test also show no significant shift in these variables in the post-policy year.

The assumption of similar trajectories for the treated and control groups is key to

Figure 1.5: Bay Bridge monthly mean traffic volume over time

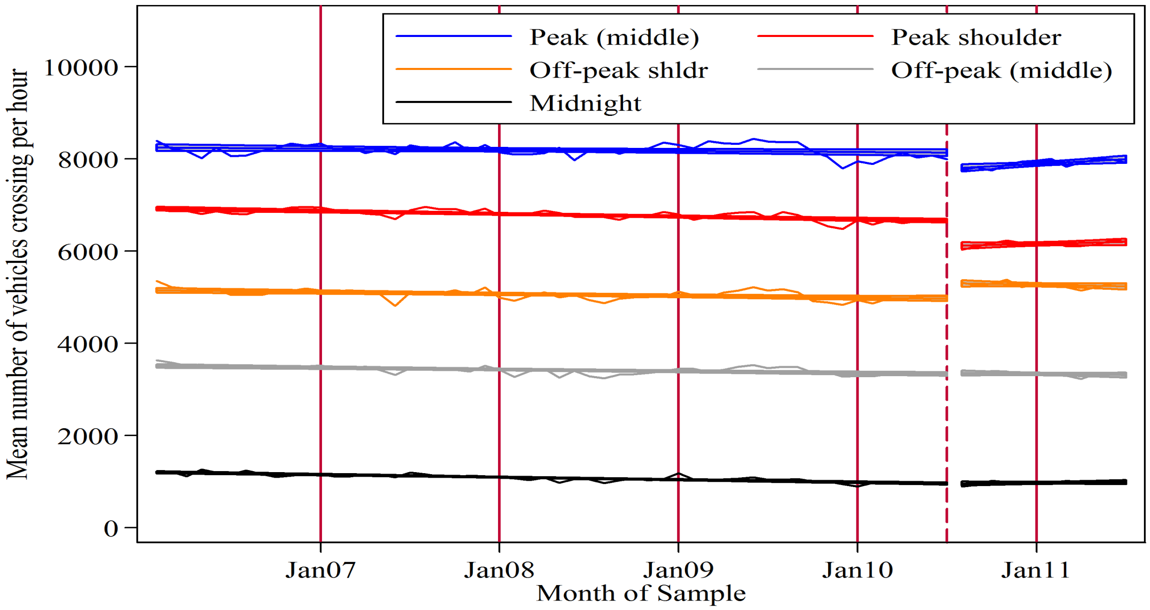

Jan 2006-July 2011. Weekdays, westbound. Excludes holidays. 95% CI. Seasonally adjusted by month.

the difference-in-differences estimation because any factors that are specific to either the treated or control hours but do not change over time, or factors that change over time equally for both time groups are netted out in the estimation. The reasonableness of this assumption is less obvious and requires testing. The evolution of peak and off-peak outcomes for the Bay Bridge from January 2007 to June 2010 (before the policy change) and from July 2010 to June 2011 (after the policy change) are shown as monthly means in figures 1.5 and 1.6.

The graph for traffic volume (figure 1.5) appears to show similar evolution of monthly means between the midnight hour and each of: peak (non-shoulder) hours, peak shoulder hours, off-peak shoulder, and off-peak (non-shoulder) hours. Statistical tests indeed indicate that the trajectories are not statistically different. This is true when testing the entire time period shown in the graph (January 2007 - June 2010), the three full years before the policy change (July 2007 - June 2010), and the full year before the policy change (July 2009 - June 2010), which is the data I use in my DD estimation. For comparison, the figure also shows the monthly means for the year following July 2010; the peak, peak shoulder, and off-peak shoulder hours appear to have been affected by the policy, while the midnight and off-peak non-shoulder hours appear to have been unaffected by the policy, which strengthens the SUTVA. I conclude that the traffic volume in the treated and untreated hours have similar trajectories, absent treatment, and thus these data are suitable for a difference-in-difference approach.

The graph for travel time (figure 1.6) appears to show similar evolution of monthly

Figure 1.6: Bay Bridge monthly mean travel time over time

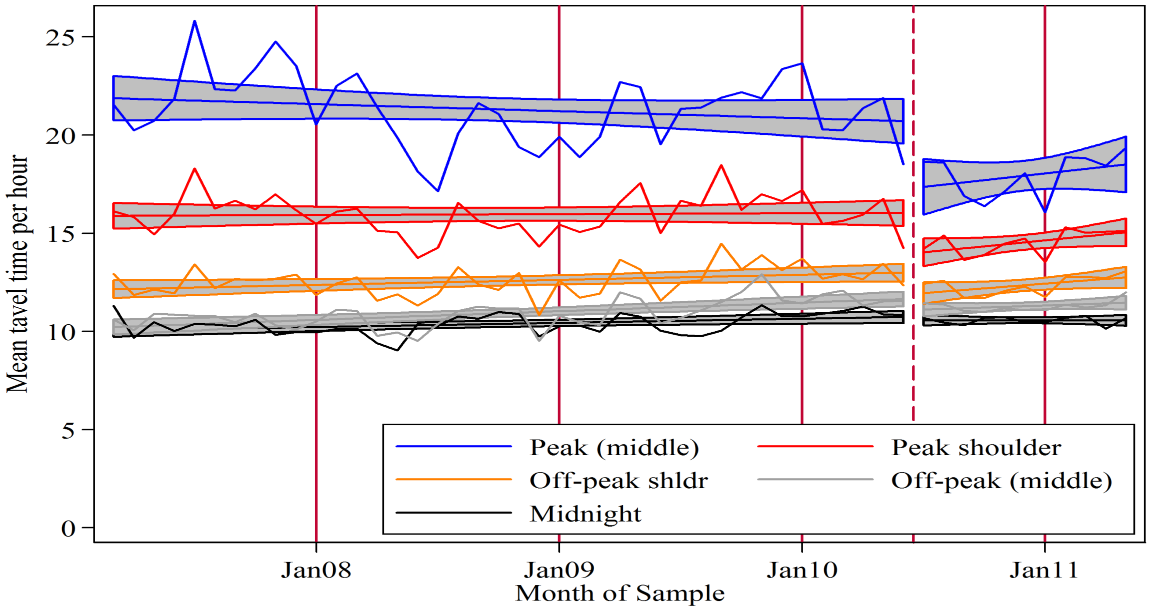

Mar 2007-July 2011. Weekdays, westbound. Excludes holidays. 95% CI. Seasonally adjusted by month.

means between the midnight hour and each of: peak (non-shoulder) hours, peak shoulder hours, off-peak shoulder, and off-peak (non-shoulder) hours. Statistical tests, weighted by the the number of vehicles using each route, show mixed results. For each of the three periods tested (the full sample from January 2008 to June 2010; the two years before from July 2008 to June 2010; and, one year before from July 2009 to June 2010), the midnight hour shows a statistically significant, but practically insignificant trend increase of 1.2 seconds. The off-peak shoulder trend is statistically not different from the midnight hour. The off-peak (non-shoulder) and peak shoulder hours show trends differing from the midnight trend by less than 3 seconds, significant at the 10% level. The peak (non-shoulder) hours have a trend that is different from the midnight hour by -7.2 seconds, significant at the 5% level. Overall, although there is mild statistical significance, the practical significance of the differences in trend are small enough to be considered insignificant, and I conclude that the travel time in the treated and untreated hours have similar trajectories, absent treatment, and thus these data are suitable for a difference-in-differences approach.

In summary, the validity of the SUTVA and the evidence shown above regarding common evolution of the outcome variables in treated and control hours suggest that employing a difference-in-differences approach using the midnight hour as a control is a tenable method for estimating the hourly average treatment effect of the recently implemented congestion pricing on the Bay Bridge.

Table 1.2: Traffic volume regressions results, by bridge

Dependent variable:	Bay		San Mateo		Dumbarton	
Traffic volume (vehicles/hr)	(1)	(2)	(3)	(4)	(5)	(6)
After July 1, 2010	-87.19	-5.68	-14.75	-11.63	-48.08	-10.59
	(41.63)**	(18.58)	(16.14)	(6.65)*	(14.40)***	(2.91)***
Peak hours		6476.53		2884.43		2280.49
		(26.00)***		(15.44)***		(19.11)***
Off-peak hours		2934.28		989.95		585.16
		(35.99)***		(10.82)***		(6.53)***
Peak hours, after		-312.00		32.46		-58.97
		(35.19)***		(22.41)		(26.91)**
Off-peak hours, after		65.23		-24.06		-24.53
		(50.08)		(15.28)		(8.99)***
Const.	5254.72	1111.80	1922.20	261.81	1321.93	124.33
	(31.28)***	(13.12)***	(11.54)***	(3.83)***	(10.39)***	(1.80)***
Off-peak mean	3879	3879	1169	1169	654	654
Peak mean	7428	7428	3157	3157	2370	2370
Obs.	11594	11594	11594	11594	11594	11594
F statistic	4.39	27732.7	.84	16638.35	11.15	9629.48

Significance is at the 10%=*, 5%=**, and 1%=*** levels and use Newey-West standard errors with 26 lags. All specifications include only weekdays from 1 July, 2009 to 30 June, 2011, excluding holidays and hours when the Bay Bridge was closed for repair.

1.4.2 Traffic Volume

Before running my main regression for traffic volume, I begin with more simplified comparisons for each bridge: comparing the traffic volume before and after the policy change; and, comparing pooled peak and pooled off-peak changes before and after, still using midnight as the control hour. The results, using Newey-West standard errors with 26 lags, are shown in table 1.2. Column 1 shows that when comparing the unconditional means of hourly traffic volume before and after the policy change, there is a statistically significant decrease of 87 vehicles per hour (1.7%) on the Bay Bridge. When separating the data into peak and off-peak hours[7] (column 2), the data show that after the change, the average peak hour traffic volume decreased by a statistically significant 312 vehicles per hour, a change of 4.2%, which is much larger than the statistically insignificant decreases found by Deakin and Frick (2011). The overall off-peak change was an increase of 65 vehicles per hour, but is not statistically significant.

The results for the San Mateo Bridge for before and after (column 3) show the expected direction (a decrease of 15 vehicles per hour) due to the increase in price from \$4 to \$5, although it is insignificant. When broken down into peak and off-peak hours (column

[7]Off-peak excludes the hour from midnight to 1:00am

4) the results for the San Mateo Bridge again show the expected sign for off-peak hours (decrease of 24 vehicles per hour), as well as evidence consistent with traffic substituting from the Bay Bridge to the San Mateo Bridge during peak hours (an increase of 33 vehicles per hour), although neither of these estimates is statistically significant.

The results for the Dumbarton Bridge for before and after (column 5) show the expected direction: a statistically significant decrease of 48 vehicles per hour for the increase in price from \$4 to \$5. When broken down into peak and off-peak hours (column 6) the results for the Dumbarton Bridge show a statistically significant decrease in both peak and off-peak hours: 59 and 24 vehicles per hour, respectively. This is consistent with little, if any, substitution toward the Dumbarton Bridge.

I then run my full estimation for my Bay Bridge data, estimating each hour independently and including day of week and week of sample fixed effects. The point estimates for the change in traffic volume in each hour caused by the congestion pricing changes can be seen graphically in figure 1.7. The graph shows small and statistically insignificant changes in the off-peak hours from 1:00am-3:00am, 11:00am-1:00pm, and 8:00pm-11:00pm. The estimates for the peak hours are all negative and mostly significant, indicating that the congestion pricing worked to decrease traffic volume during peak hours. The estimates for the peak shoulder hours (5:00am, 9:00am, 3:00pm, and 7:00pm) are, as expected, all large, negative, and significant, ranging from 400 to 550 fewer vehicles per hour (decreases of 6% to 8%). The estimates for the off-peak shoulder hours (4:00am, 10:00am, 2:00pm, and 7:00pm) are, as expected, large, positive and significant, ranging from 225 to over 400 more vehicles per hour (increases of 4% to 20%).

These results confirm what the long-standing theory predicts will happen when peak time congestion pricing is implemented: (1) overall traffic volume decreased–because of the higher price, people are making fewer driving trips across the bridge, and (2) some travelers substituted from driving during peak hours to driving during the less congested off-peak hours, resulting in peak-spreading.

The analogous graph for the San Mateo Bridge is figure 1.8. The graph shows small and mostly statistically insignificant decreases from 9:00am onward. The estimates for the remaining morning peak hours are positive, and the estimates for the hours between 5:00am-7:00am are statistically significant and range from 100 to 200 more vehicles per hour (around a 5% increase for each of those two hours). This lends further suggestive evidence that some morning travelers may be substituting away from the higher priced Bay Bridge toward the San Mateo Bridge. However, the number of drivers substituting bridges is likely to be small, as there was an overall decrease of 312 vehicles per hour during peak hours on the Bay Bridge, and an overall increase of only 32 vehicles per hour during peak hours on the San Mateo.

The analogous graph for the Dumbarton Bridge (not shown) indicate small to moderate and mostly statistically significant decreases from 8:00am onward. The estimate for the 6:00am hour is small, positive, and significant, but represents only around a 2.5% increase in traffic volume in that hour. This is inconclusive, although one possible expanation is that the drivers substituting from the Bay Bridge Bridge to the San Mateo

Figure 1.7: Bay Bridge point estimates of traffic volume treatment effect

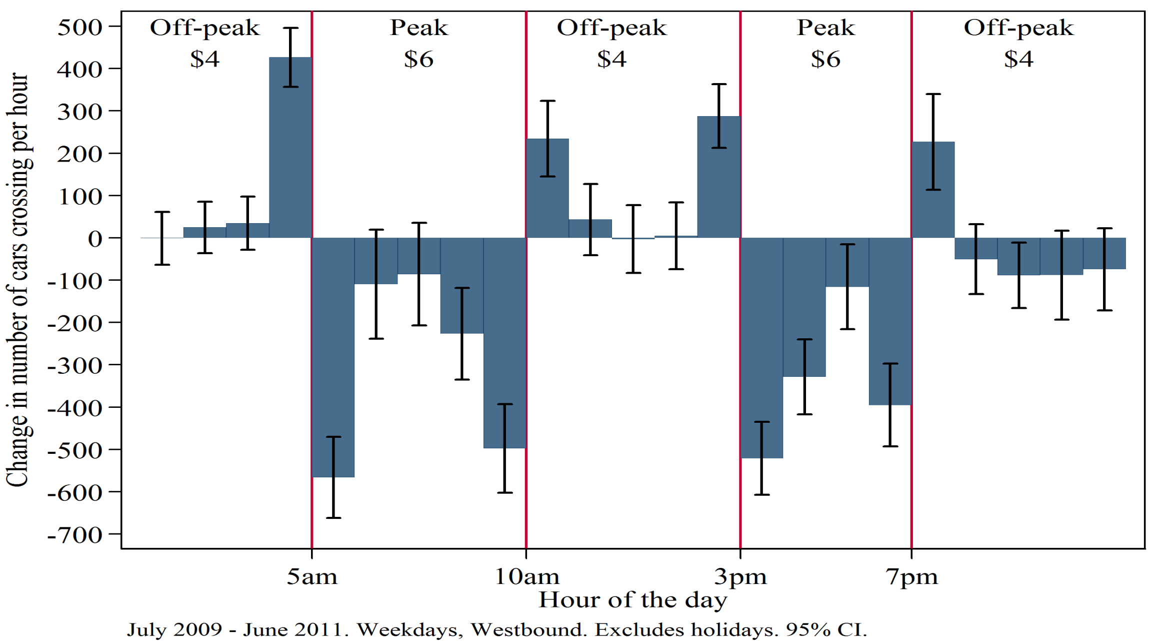

Figure 1.8: San Mateo Bridge point estimates of traffic volume treatment effect

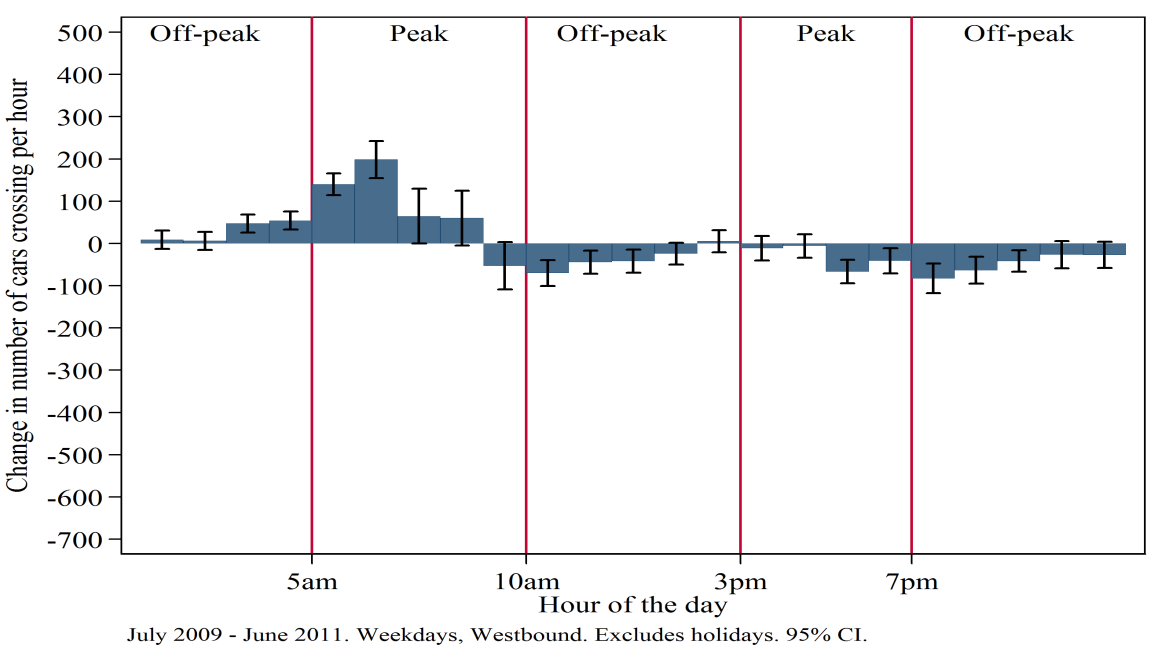

Bridge are crowding out some drivers who previously used the San Mateo Bridge, who are now substituting toward the Dumbarton Bridge.

1.4.3 Traffic Volume by Lane Type

Breaking the average treatment effects down by lane type yields some insight into what kinds of drivers are time-shifting versus route-shifting. Carpool lane usage on the Bay Bridge has decreased by between 300-500 vehicles per hour during all peak hours. This is unsurprising, and likely a direct result of two factors: the new requirement to carry a FasTrak in order to use the carpool lanes; and, the toll for carpoolers in peak hours going from \$0 to \$2.50. These two factors also have nuanced indirect effects.

The new requirement for carpoolers to use a FasTrak to pay for the toll is a barrier in itself because people must acquire a FasTrak tag, and besides actually going out and buying one, some drivers may be afraid to carry a FasTrak tag because they do not want their toll crossings to be observable by the toll authority. In addition, the FasTrak requirement means that toll violators are easier to catch: previously, with no toll, there was no need to catch toll-violators, so it was more difficult to catch those who were using the toll-free carpool lanes without the required number of people. Now, if a vehicle does not have a Fastrak and tries to use the carpool lane, a photograph is taken of the vehicle's license plate in order to collect the toll. This likely discouraged anyone who was previously crossing with fewer than the required number of people in the carpool.

The price increase from zero to a positive number, besides having a direct price effect on those who no longer deem the effort to organize a carpool worthwhile, also has likely impacted what is known in the Bay Area as "casual carpooling". The Bay Area has various known locations where people who would like a ride across the bridge can go and wait. Those who are driving over the bridge and wish to take advantage of the carpool discount and designated lanes can stop by and pick up the people who are waiting for a ride. With no toll, the riders just got into the car and were driven across the bridge. Since the price increase from \$0 to \$2.50, there is now the social awkwardness of a positive price. Should each rider offer to pay one third (\$0.80) or one half (\$1.25) of the toll? Should the driver request part of the toll to be paid by riders or accept any offer of partial toll payment? Surveys of the casual carpool community have shown a decrease in casual carpools since the toll hike, and social awkwardness is one main reason stated for the decrease (see Deakin and Frick (2011)). Most designated pickup locations have developed a norm for what amount is considered acceptable for the rider to offer and the driver to accept in response to the toll hike. Again, the salience offers an insight into the curiosity of human behavior: there is little social awkwardness in not offering to help pay for gas or other costs of driving, but the toll, because it is paid for while the riders are in the car, causes driver and rider to behave differently.

Focusing only on the non-carpoolers, that is the FasTrak and cash lane users (see figure 1.9), there is time-shifting occurring around the shoulder hours and little change in the off-peak non-shoulder hours. However, there are increases during the mid-peak hours, which

Figure 1.9: Bay Bridge point estimates of traffic volume treatment effect for non-carpool lanes

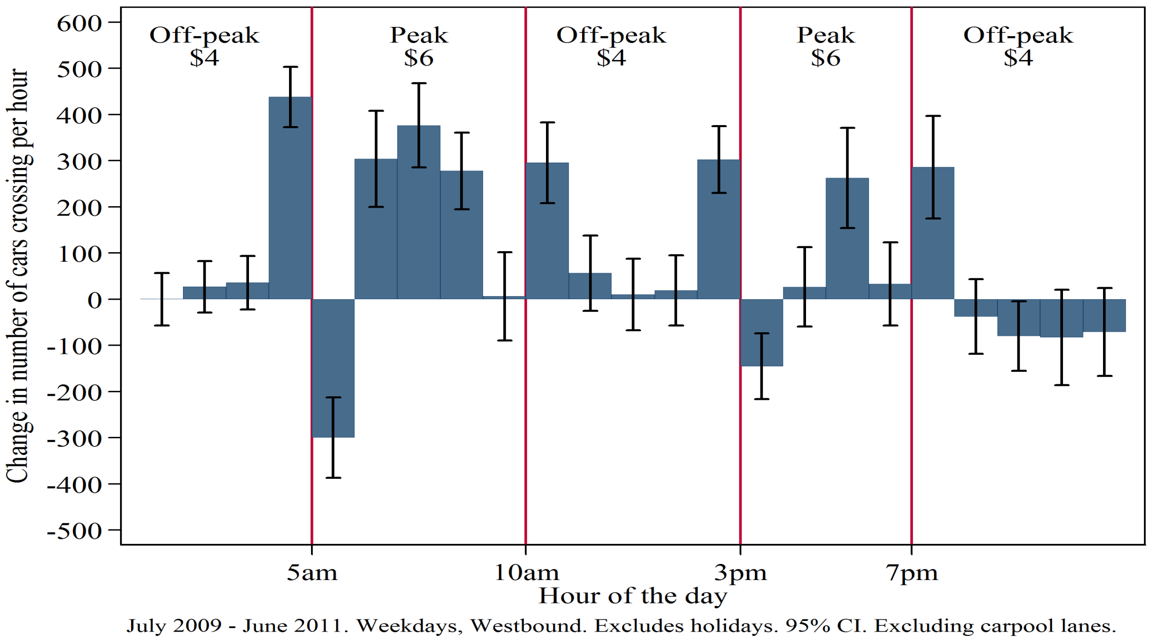

suggests that the FasTrak and cash lanes are absorbing some of the carpool decreases, leading to overall increases in some peak hours for the non-carpool lanes. This result is perhaps an unforeseen consequence of increasing the price of carpooling. Much of the non-shoulder hour decrease in non-carpool usage comes from the cash lanes, suggesting that cash-payers are more price sensitive, which is in line with economic theory about FasTrak tags and toll salience (see Finkelstein (2009)).

1.4.4 Travel Time

For my travel time results, I again begin with more simplified comparisons for each bridge: comparing the travel time before and after the policy change; and, comparing pooled peak and pooled off-peak changes before and after, still using midnight as the control hour. The results are shown in table 1.3. Column 1 shows that when comparing the unconditional means of hourly travel time before and after the policy change, there is a statistically significant decrease of 1.6 minutes (9.6%) on the Bay Bridge. When separating the data into peak and off-peak hours[8] (column 2), the data show that after the change, the average peak hour travel time decreased by a statistically significant 2.9 minutes (14%). The off-peak hour travel time decreased by 20 seconds (3%).

The results for the San Mateo Bridge for before and after (columns 3) show no change

[8]Off-peak excludes the hour from midnight to 1:00am

Table 1.3: Travel time regressions results, by bridge

Dependent variable:	Bay		San Mateo		Dumbarton	
Travel time (minutes)	(1)	(2)	(3)	(4)	(5)	(6)
After June 30, 2010	-1.60	-.12	.03	.04	.008	.17
	(.12)***	(.09)	(.09)	(.03)	(.08)	(.04)***
Peak hours		9.51		.41		.97
		(.12)***		(.12)***		(.10)***
Off peak, exl 12am		2.08		-.07		.47
		(.07)***		(.06)		(.04)***
Peak after change		-2.94		.02		-.17
		(.14)***		(.14)		(.11)
Off peak, after		-.34		-.07		-.16
		(.09)***		(.06)		(.05)***
Const.	16.64	10.74	10.45	10.23	8.58	7.75
	(.10)***	(.06)***	(.08)***	(.02)***	(.08)***	(.03)***
Off-peak mean	12.5	12.5	10.2	10.2	8.1	8.1
Peak mean	20.5	20.5	10.6	10.6	8.5	8.5
Obs.	77104	77104	9963	9963	9954	9954
F statistic	166.12	2724	.12	39.6	.008	144.13

Significance is at the 10%=*, 5%=**, and 1%=*** levels and use Newey-West standard errors with 26 lags.
All specifications include only weekdays from 1 July, 2009 to 30 June, 2011, excluding holidays and hours when the Bay Bridge was closed for repair.

in travel time due to the increase in price from \$4 to \$5. No change is also the case when the hours broken down into peak and off-peak hours (column 4). This suggests than any extra traffic the San Mateo Bridge may have experienced as a result of drivers substituting from the Bay Bridge was easily absorbed by the San Mateo Bridge.

The results for the Dumbarton Bridge for before and after (column 5) show no change in travel time due to the increase in price from \$4 to \$5. When broken down into peak and off-peak hours (column 6) the results for the Dumbarton Bridge show a statistically significant (though practically small) decrease in off-peak hour travel time of 9.6 seconds. This suggests that there was little, if any, effect on travel time on the Dumbarton Bridge.

Turning to the results for the full set of hours, the estimates for the hourly average treatment effect of the toll change on travel time over the Bay Bridge can be seen in figure 1.10. The estimates for the overnight hours between 8:00pm and 5:00am are extremely small and largely insignificant. The estimates for the morning peak hours from 6:00am to 9:00am show decreases in travel time of 2 to 6 minutes (a 12% to 24% decrease in travel time). Additionally, I see a decrease in travel time during all evening peak hours of between 2 to over 4 minutes, as well as small (less than one minute) in travel savings during some between-peak hours. These findings suggest that the congestion pricing

Figure 1.10: Bay Bridge point estimates of travel time treatment effect

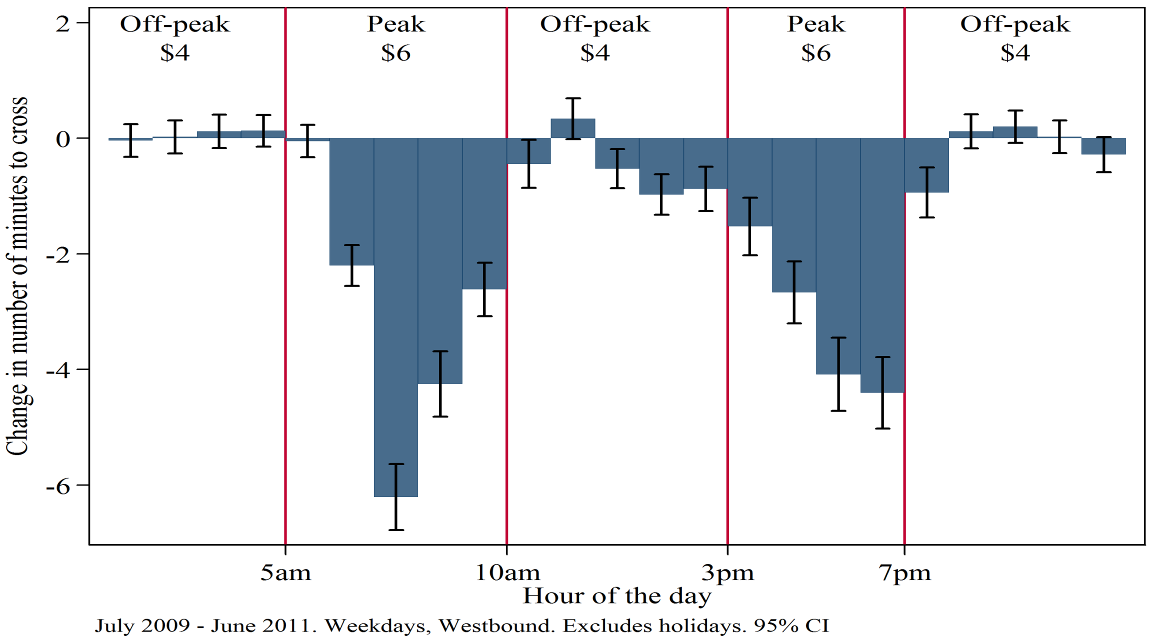

worked, as intended, to decrease traffic congestion and allow vehicles to travel at higher speeds during peak hours.

The estimates for the San Mateo and Dumbarton Bridges (not shown) indicate that travel times did not change on those bridges: all estimates are under 1 minute in magnitude and almost all are statistically insignificant. This is consistent with the hypothesis that these two bridges are being used at or below capacity and do not require peak time pricing of congestion. In addition, this again suggest that any extra traffic the San Mateo Bridge may have experienced as a result of drivers substituting from the Bay Bridge was easily absorbed by the San Mateo Bridge.

Travel Time Variability In addition to ascertaining the decrease in expected travel time, I also calculate the change in the variability of travel time, since travelers care not only how long a trip takes on average, but also the reliability of their expected trip time. Figure 1.11 shows that the standard deviation of median travel time decreased in most hours throughout the day on the Bay Bridge. The decrease in standard deviation can be as much as 4.3 minutes (7:00am hour) or 5.6 minutes (3:00pm hour). The decreases are significant at the 1% or 5% level, except for hours 3:00am, 11:00am, 8:00pm, 9:00pm, and 10:00pm, none of which are peak hours. This holds true for both Levene's tests centered either at the mean or the median.

Figure 1.11: Bay Bridge standard deviation of travel time

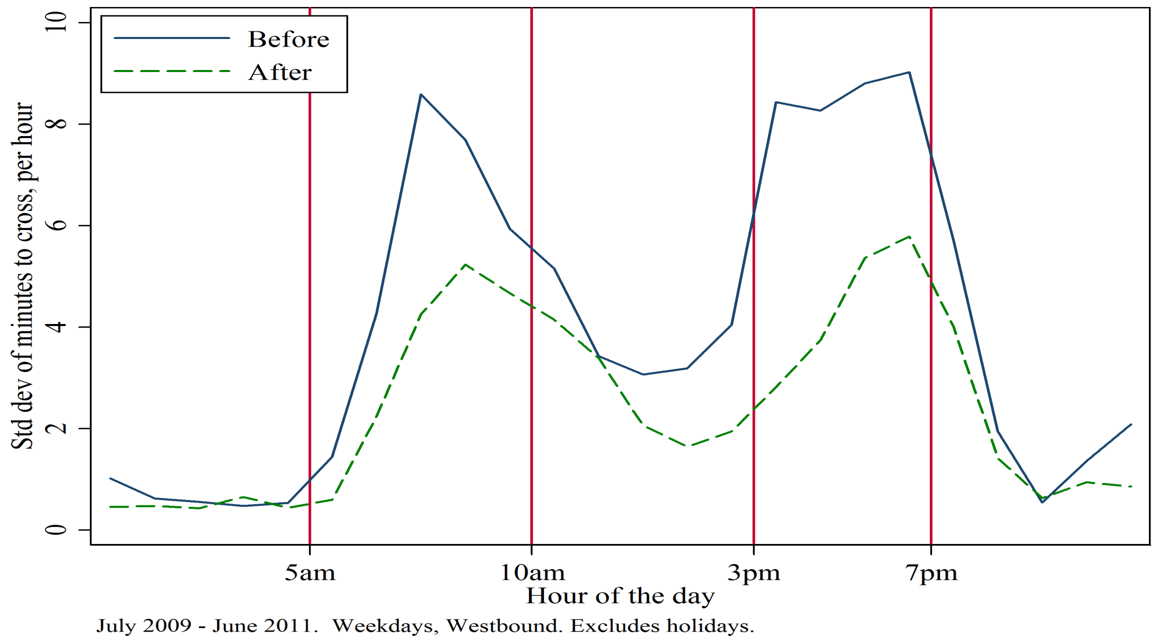

1.4.5 Robustness Checks

I run the same regressions but with the log of traffic volume (instead of level) and get qualitatively and quantitatively very similar results. This robustness check suggests that my estimated treatment effect is not being driven by the combination of a small level change in my midnight travel outcome and a large level change in my outcome in my treated hours. These results can be seen in the appendix, in figure A.3.

In another robustness check, I expand my control hours from midnight alone to include all hours from 11pm-3am and the results do not change substantially. This robustness check suggest that my estimated treatment effect is not driven by choosing a single unaffected hour. These results can be seen in the appendix, in figure A.4.

Additionally, if I relax assumptions from the difference-in-differences and use a triple difference approach instead (using the Dumbarton Bridge as the additional control), then I find similar results to my difference-in-differences. My triple difference estimates suggest that my treatment estimates are not driven by the assumption of treated and midnight hours evolving similarly; I only need to assume that, to the extent that the outcomes evolve differently, they evolve similarly on the two bridges. These results can be seen in the appendix, in figure A.5 and A.6.

I also use the last toll price change, in 2007, as a falsification test to see how drivers responded to a flat toll increase from \$3 to \$4. The results show small and barely significant decreases in most hours, with some moderate and barely significant increases in traffic volume in the morning peak hours. This is a very different result from the effects of the congestion pricing put into practice in 2010, which strengthens my findings of peak-spreading from the time-varying congestion pricing. These results can be seen in the appendix, in figure A.7.

1.4.6 Alternative Methodology: Regression Discontinuity

As an alternative to the DD approach, I employ a regression discontinuity (RD) design to capture the short-run effects of the congestion pricing to see if they are much different from the longer-run effects captured in the DD. For the RD, I use daily averages of the outcome variable by hour category (e.g. peak, off-peak) and use data from one month before and one month after the change. For comparison, I show the same data for the two years before and the year after the policy change. As can be see in figure 1.12, there is very little, if any, change in peak hour traffic volume in the 2008, 2009, and 2011 periods. However, the 2010 graph shows a large drop in peak hour traffic volume of 639 vehicles per hour (significant at the 1% level). There is also a statistically significant increase in the morning off-peak shoulder hours of 289 vehciles/hour (not shown). All other hour categories show insignificant changes on 1 July, 2010 on the Bay Bridge.

Interestingly, although the immediate decrease was large (and larger than the decrease shown by the DD: 639 compared to 312), there seems to be evidence of a fairly quick recovery of peak hour traffic volume in 2010, as shown by the upward sloping fit line in the second half of the 2010 figure. In addition, the 2011 graph shows the traffic volume seems to be below the 2008 and 2009 levels, but above the July 2010 levels, suggesting that traffic volume ended up half way between the 2008 and 2009 levels of around 7500 vehicles per hour during peak hours and the July 2010 levels of around 7000 vehicles per hour. One explanation for this could be that the initial response of drivers to a $2 toll increase was to react strongly to it, but after they saw the decreased travel time benefits of the policy or just figured that $2 was not a large increase after all, they started to increase their trips again, albeit to lower levels that the pre-policy change period.

Similarly, I show the RD results for travel time in figure 1.13, which shows a small, but significant, decrease in peak hour travel time of 4.4 minutes per hour in 2010. The other years show little, if any discontinuity at the 1 July date. As above, the other hour categories have small (under 20 seconds) and mostly insignificant changes in travel time.

1.5 Policy Implications and Welfare Impacts

From a policy maker's perspective, there are important lessons to be learned from this analysis. First of all, congestion pricing seems to work fairly well at decreasing traffic volume, travel time, and travel time variability during higher-priced hours: drivers seem to peak spread by shifting the time of their travel to off-peak hours; and they avoid making unnecessary trips. Some broader implications and welfare impacts are discussed below.

Figure 1.12: Bay Bridge graphs of Regression Discontinuity

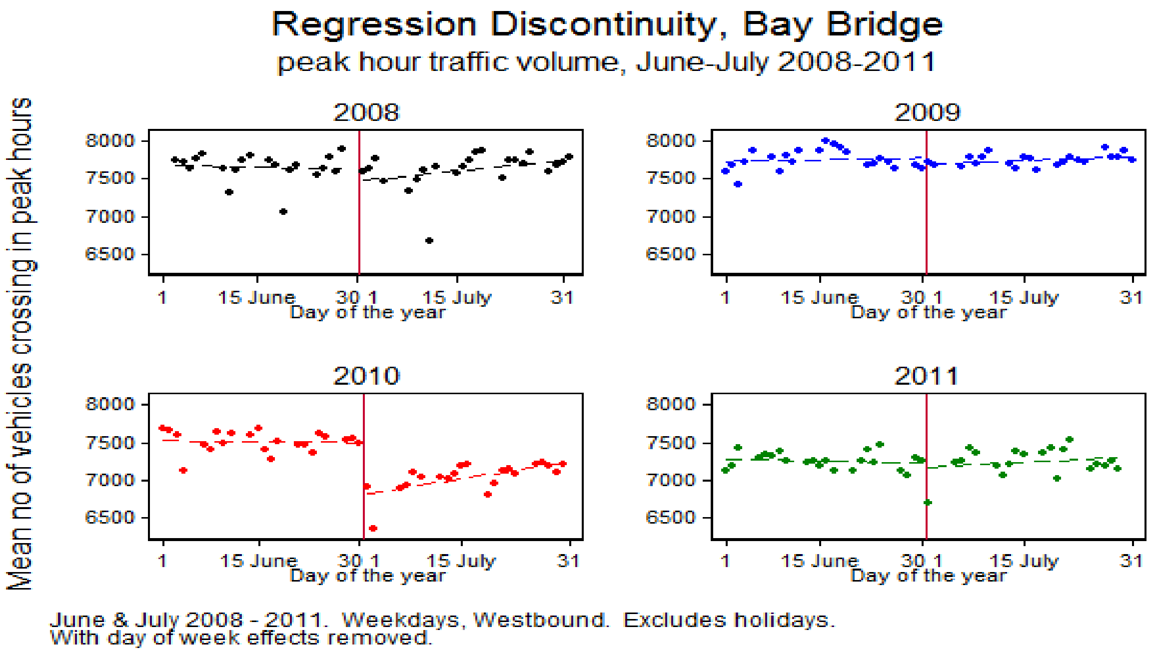

Figure 1.13: Bay Bridge graphs of Regression Discontinuity

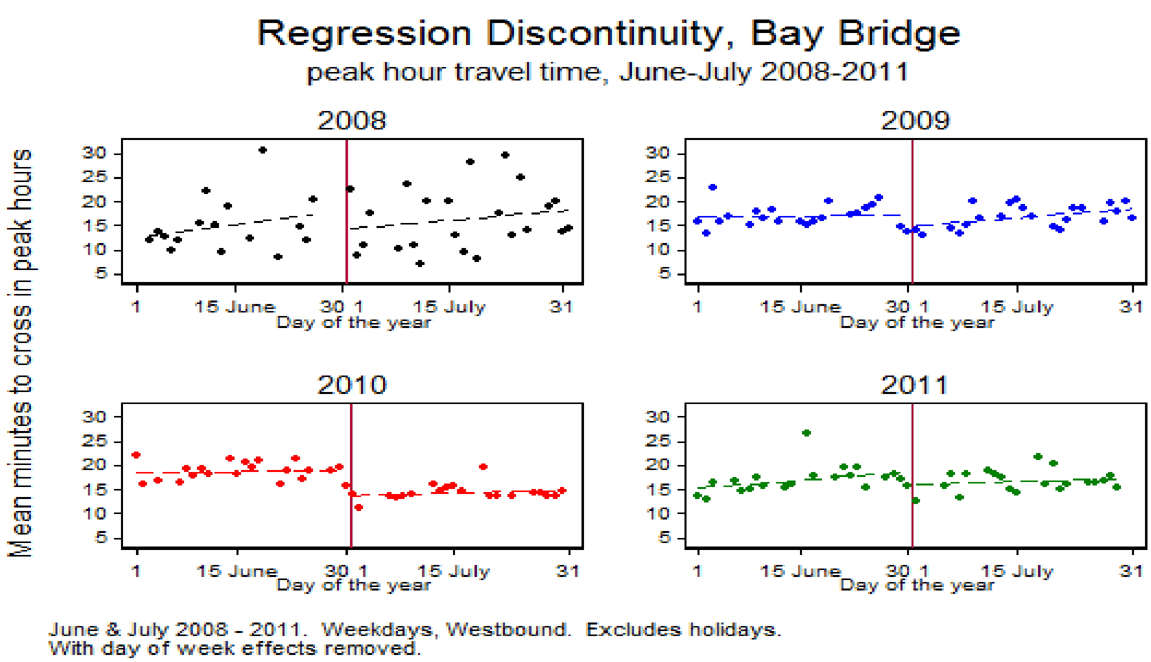

1.5.1 Price Elasticities

One potentially useful measure of drivers' responses to congestion pricing are the price elasticities of demand for driving. Both own- and cross-price elasticities can be calculated. For peak hour travel over the Bay Bridge, the overall own-price elasticity is -0.084, meaning that a 10% increase in the Bay Bridge toll decreases across-bridge trips by under 1%. Or, as actually happened, a 50% increase in the toll during peak hours decreased peak trips by about 4.4%. If elasticities are broken down by hour, then the range of own-price elasticities for peak hours is 0 to -0.20, the largest of which is during the 5:00am hour. For off-peak hour travel over the Bay Bridge, treating the peak hours as substitutes for off-peak hours, the overall cross-price elasticity is 0.034, meaning that a 10% increase in the Bay Bridge peak hour toll increases across-bridge trips during off-peak hours by one third of one percent. Or, as actually happened, a 50% increase in the toll during peak hours increased off-peak trips by about 1.5%. If elasticities are broken down by hour, then the range of cross-price elasticities for off-peak hours is 0 to 0.38, the largest of which is during the 4:00am hour.

My estimate of the overall own-price elasticity of demand is in line with previous similar estimates. Finkelstein (2009) estimates the price elasticity of tolls to be around -0.061.

These elasticities might seem small; a 50% increase in toll prices sounds large, but it is only a $2 increase, and the toll, while salient, is only a small portion of the average total costs of driving. Other costs of driving include gas, parking, maintenance, and time costs. As stated above, for a round trip from a mid-distance communting city, such as Walnut Creek, these costs can add up to between $23 and $100, so the toll change is a smaller percentage change of costs if all costs of driving are considered. If the full costs ($23-100) are considered, then peak own-price elasticities are more in the range of -1.05 to -2.1, and off-peak cross-price elasticities are more in the range of 0.42 to 0.84. Average total costs of driving are important in the long run decisions drivers make about their daily commute, living and working locations, and other considerations. More pertinent in the short run is the marginal cost of driving, which, again considering a trip from Walnut Creek might be between $24 and $90, and would mean the peak own-price elasticities are lower, at around -0.50 to -1.89, and off-peak cross-price elasticities are more in the range of 0.20 to 0.75.

Although these estimates of elasticities are internally valid and can provide some guidance to urban planners, their external validity for other congested areas may depend on the availability of substitutes (time-shifting, alternative routes, efficacy of public transit), income, travel time, and numerous other factors. In addition, the use of these elasticities for anything other than traffic volume predictions should be done very carefully. As with time-varying electricity pricing, shifting driving from one part of the day to another is not always a one-for-one transfer in terms of time savings, emissions, and the like. Perhaps a more useful analysis is a categorization of the welfare impacts of the congestion pricing.

1.5.2 Time Savings

Using my point estimates of the causal impacts of the toll price change, I construct estimates of the value of time saved due to the time-varying nature of the congestion pricing. Because my travel time data only track FasTrak users, I can only generate bounded estimates of the benefits from travel time savings. For the lower bound, I assume that the only people who have saved time due to the congestion pricing are those that I can track who are using FasTraks, then I have a daily flow of those users on my eight routes of about 30,000 between the hours of 4am and 8pm, and on average each of those vehicles saves 1.7 minutes, which sums to annual savings of about 210,000 hours. The total value of the hours saved depends on the willingness to pay for travel time savings. I used two estimates: \$5 and \$30 per hour[9], which gives annual savings from decreased travel times of between \$1 and \$6.4 million. Some reasons my lower bound might be too low include: I only have FasTrak data for eight routes; many vehicles have more than one person in the vehicle; an increased number of people are using the FasTrak lanes; and, some people put their FasTrak tag away when not in use (thus they would not be measured). In addition, I haven't estimated the travel time savings for the return trip (which could double these estimates) or any travel time savings from spillovers of decreased bridge traffic to other roadways, which would decrease congestion and reduce travel time for those road users.

If I assume that every vehicle crossing the bridge saves 1.7 minutes of travel time on average, then my daily flow is 124,000 vehicles (as measured by total traffic count instead of FasTrak), then annual travel time savings are almost 900,000 hours, which translates into a value of savings between \$4.4 and \$27.6 million. This is likely to be an overestimate of travel time savings because those in the cash lanes have not likely experienced as much of a decrease in travel times, although it is not necessarily an overestimate because I don't consider travel time savings from the return trip or spillovers. Additionally, there are likely to be gains from decreased travel time variability, which are valued at around \$19.56 per hour (see Small (2005)).

1.5.3 Fuel Savings

The relationship between fuel efficiency and speed suggests that driving at speeds of around 50mph is optimal in terms of fuel usage[10]. Decreasing congestion can decrease fuel usage per mile driven by increasing speeds toward the optimum. Based on my travel time estimates and the distance of each route to cross the Bay Bridge, I find that average miles per hour increased from the low 40's to the high 40's, resulting in increased fuel efficiency, and a lower bound on fuel savings of 230,000 gallons per year. With gasoline

[9] If minimum wage in the Bay Area is \$10/hour and people value leisure time at half that of working time, so \$5 per hour. If the estimates from Brownstone et al (2003) are used, then WTP for travel time savings is more in the range of \$30 per hour.

[10] I use the emissions factors from the California Air Resources Board's EMFAC model, which have estimates that are specific to the Bay Area.

prices currently hovering right around \$4/gallon, this results in savings to drivers of nearly \$1 million per year. As before, this is likely to be a lower bound, both because I do not consider the return trip, and because traffic is more likely to be stop-and-go with congestion, which would increase fuel usage, and thus there would be increase fuel savings from decreased congestion. This is small compared to an estimate of total fuel used to cross the Bay Bridge every year (about 10 million gallons[11]), but is nonetheless a significant positive savings, and again is likely to be a lower bound.

1.5.4 Emissions Abatement

Similarly to the speed-mpg relationship, the relationship between CO_2emissions and speed suggests that driving about 50mph emits the lowest amount of CO_2 per mile driven[12]. Decreasing congestion can decrease CO_2 emissions per mile driven by increasing speeds toward the optimum. Based on my travel time estimates and the distance of each route to cross the Bay Bridge, I find a lower bound on CO_2abated of 2240 tons per year. This is equivalent to taking about 470 commuting cars off the road every year[13]. If I apply a value on carbon dioxide of \$20/tCO_2, then those savings have generated benefits of nearly \$45,000 annually by reducing CO_2 emissions. This is a lower bound both because I do not consider the return trip and because the price I use for carbon is on the lower end of what many economists think the price for carbon ought to be.

Other air pollutants (e.g. NO_X, volatile organic compounds, ozone, CO, and particulate matter) that are emitted from vehicles have complicated emissions structures and mix unevenly to produce pollution due to weather patterns and other factors[14]. This makes the effects on these other pollutants difficult to estimate, but if traffic volume deceased overall, then the effect would have been to decrease these pollutants as well.

1.5.5 Other Welfare impacts

Decreasing traffic volume and travel time is likely to have other positive welfare impacts, such as decreased noise pollution, reduced water pollution, lower infrastructure damage, less vehicle maintenance required, beneficial psychological effects from being stuck in traffic less (reduced road rage), and potentially decreased traffic accidents.

Since the time-varying congestion pricing on the Bay Bridge was designed to raise revenue for bridge maintenance and the prices were chosen to bring in roughly the same revenue as would have been brought in from a flat increase from \$4 to \$5, it seems like the

[11]Roughly 124,000 vehicles per day * 250 work days per year * 8 miles * 0.041 gallons per mile (at a speed of around 40mph) = 10 million gallons.

[12]Again, I use the emissions factors from the California Air Resources Board's EMFAC model, which have estimates that are specific to the Bay Area.

[13]Calculated as: 12,000 VMT per vehicle/year *0.0003989 tCO\$_2\$/VMT = 4.8tCO_2 per vehicle per year. Divide total tCO_2 saved by 4.8CO_2 to get 470 vehciles.

[14]See Auffhammer and Kellogg (2011)

benefits from implementing time-varying pricing are just generating additional consumer surplus.

1.6 Conclusions

This paper analyzes the effectiveness of the recently implemented congestion pricing on the San Francisco-Oaklsand Bay Bridge. Using hourly data on traffic volume and travel time, and employing a difference-in-differences empirical econometric approach, I find that the policy change decreased traffic volume across the Bay Bridge by and average of 312 vehicles per hour during peak hours. My analysis indicates that drivers took fewer trips overall and some drivers substituted from the peak hours to the off-peak shoulder hours, which resulted in decreases in travel time of 2-6 minutes (7-20% decreases) during peak hours. I find suggestive evidence of drivers substituting toward the San Mateo Bridge, which is an alternative route that lacks congesting pricing. I also show significant decreases in travel time variability during peak hours.

Using hourly data on traffic volume and travel time, and employing a regression discontinuity empirical econometric approach, I find that the policy change decreased traffic volume across the Bay Bridge by and average of 639 vehicles per hour during peak hours. There were also immediate average decreases in travel time of 4.4 minutes during peak hours.

In addition, I calculate peak-hour own-price elasticities on the Bay Bridge. If only the toll price is considered, then the own-price elasticity is -0.084 overall, and -0.20 for a peak shoulder hour. If the overall cost of driving is considered, then the elasticities are more in the range of -1.05. For the off-peak-hour cross-price elasticities on the Bay Bridge, if only the toll price is considered, then the cross-price elasticity is 0.034 overall, and 0.38 for an off-peak shoulder hour. If the overall cost of driving is considered, then the elasticities are more in the range of 0.42.

Using back-of-the-envelope calculations, I find a lower bound of the annual welfare impacts of this program to include: travel time savings of between 210,000 to 900,000 hours; fuel savings of around 230,000 gallons; and, emissions abatement of around 2240 tons of CO_2. The value of these savings is in the range of $6.4 to $35 million.

My future work to extend the current analysis will include integrating data on public transit ridership to estimate inter-modal substitution and cross-price elasticities between driving and two modes of public transit.

Chapter 2

Estimating the Impact of Reducing Agricultural Water Supplies On Farm Employment and Fallowing for California's Central Valley

2.1 Introduction

California has experienced nine large-scale multi-year droughts since 1900. The two most recent of these occurred in 2000-2002 and 2007-2009. These droughts, and less pronounced dry spells, have decreased water inflow to the San Francisco Bay Delta, the largest estuary in the western US. Besides being an important habitat for a number of threatened and endangered species, the Delta is the origin of water purveyed by the State Water Project (SWP) and the federal Central Valley Project (CVP) to 23 million Californians, including many farmers in the agricultural heartland of California - the San Joaquin Valley. Legislative protection of a specific species of fish, the so-called Delta Smelt, combined with these decreased inflows have significantly reduced water deliveries to farmers. In 2009, only 40% of SWP and 10% of CVP water delivery requests were filled leading to a severe shortage of surface irrigation water. In one of the most sensitive areas supplied by the CVP, the west side of the San Joaquin Valley, contractors have only received their full allocation for three years since 1990 and have received more than three quarters of their full allotments for only eight years since 1990. During the latest multi year drought (2007-2009), it was reported that famers abandoned large orchards and vineyards due to these water shortages (DWR, 2012). While some farmers do have backstop groundwater available, not all areas of California have good groundwater resources. Farmers in the Central Valley have decried the decrease in water available for crop irrigation. Congressman Devin Nunes testified that up to 80,000 jobs may be lost as a result of the water delivery restrictions, which, if true, would amount to 63% of the 126,000 jobs held in the farming sector in the whole of the Central Valley in 2007[1]. He further went on to blame some of the unemployment rate of 20-50% in his congressional district on the decrease in government water deliveries. A common slogan has been that the drought is "government-imposed".

This paper addresses the question of whether irrigation water deliveries from the CVP and SWP are positively correlated with agricultural employment and cropped area, which would suggest that shortages may indeed lead to decreased farm employment in agricultural counties in California. The subject of government water delivery effects on farm employment in California is relatively novel to academia, and it has been largely ignored in the peer-reviewed literature. In one of the few studies looking at this issue, Howitt, MacEwan and Medellin-Azuara (2009a) estimates job losses from a 90% shortage in irrigation water to be in the range of 60,000-80,000, agricultural jobs, which they state is between 20-26% of direct and indirect agricultural employment in the Central Valley. The authors use a modified version of the Statewide Agricultural Production Model (SWAP), which is a calibrated optimization model, to obtain revenue estimates. They then input the revenue estimates into REMI (a regional economic model) to get estimates for job losses. In another study, Michael (2009a) uses IMPLAN, an input-output model for economic analysis, to produce estimates of 6,000 jobs lost due to decreases in government

[1]Calculated using the BEA's NAICS data on farm employment in 2007.

water deliveries. When he includes both indirect and induced job losses, Michael (2009a) puts an upper bound of 11,700 on the total job losses. He further points out several interesting facts about employment in the San Joaquin Valley. First, he notes that, for the last ten years, there has been a farm labor shortage in California, which suggests that reduced water deliveries could merely be decreasing the shortage of workers and not causing net unemployment. He also notes that, despite the drought, farm employment has increased by 3,200 jobs since 2007, while non-farm industries (such as construction) have experienced a decrease of 40,400 jobs, most likely due to the recession and subprime mortgage crisis (the San Joaquin Valley had the highest foreclosure rates in the nation). He also shows that the unemployment rate in the San Joaquin Valley has been high for some time, even during wet years with full water deliveries, so it is unlikely that the water delivery decrease is the main driver of the high unemployment rate in the Valley. He suggests that his estimates are an upper bound on the true number of jobs lost.

Howitt, MacEwan and Medellin-Azuara (2009b) revised their previous estimates of job losses from 80,000 down to 21,000. The authors use two methods to produce these numbers: a "bottom-up" approach designed to estimates how changes in water affect changes in yields, which in turn affect employment; and, a "top-down" approach that uses aggregated employment data. They warn that using the top-down approach fails to account for intra-county variation in water distribution, which can cause job losses to be concentrated in smaller areas. Michael (2009b) subsequently revised downward his estimates of job losses from reduced water availability to 8,500 jobs in all sectors. None of Howitt, MacEwan and Medellin-Azuara (2009a), Howitt, MacEwan and Medellin-Azuara (2009b), Michael (2009a) nor Michael (2009b) use econometric techniques in combination with observed data on deliveries and employment.

The contribution of this paper is to empirically estimate the response of employment and cropped area to changes in irrigation deliveries, while controlling for unobservable confounders constant at the county level (e.g. soil quality and groundwater availability) as well as common shocks affecting all counties at the same time (e.g. business cycles). We have constructed a detailed dataset which matches both employment and irrigation water deliveries to California counties. We find a statistically and economically significant impact of water deliveries on agricultural employment and cropped area. Our estimates are close to those of Michael and suggest that between 5,000 and 9,000 jobs would be lost during a year with a 90% shortage in deliveries. We further show that this effect is mostly due to area fallowed during shortage years. Our results are robust to a number of specifications and different definitions of the control group. We further show that employment in counties with better groundwater resources is less sensitive to deliveries.

The remainder of this paper is structured as follows: section 2.2 develops our empirical model, section 2.3 describes the dataset used for estimation. Section 2.4 contains the empirical results and discussions, and section 2.5 concludes.

2.2 A Simple Empirical Model

The analysis begins with the assumption that Nit is the agricultural employment in county i during year t, and D_{it} are the deliveries from the federal and state water projects to the water districts in county i in year t. Without controlling for any other observable and unobservable confounders, a basic statistical model estimating the correlation between N_{it} and D_{it} is given by:

$$N_{it} = a + b_1 D_{it} + e_{it} \tag{2.1}$$

.

The identifying assumption required in order for Ordinary Least Squares to provide consistent estimates of b_1 is that $E[e_{it}|D_{it}] = 0$. Any factor not included in this simple model, which is correlated with deliveries would violate this assumption. Examples that come to mind are soil quality and availability of groundwater. One could explicitly control for observable confounders by including them in a regression as given in equation 2.2 below:

$$N_{it} = a + b_1 D_{it} + b_2 Z_{it} + e_{it} \tag{2.2}$$

.

The vector Z_{it} contains potential observable confounders at the county level, which may be correlated with the deliveries variable and vary over time. Failing to control for these confounders will lead to biased and/or inefficiently estimated coefficients. Since one may not observe any or all confounders varying at the county level over time, one can estimate a model of the type:

$$N_{it} = a_i + d_t + b_1 D_{it} + e_{it} \tag{2.3}$$

where the a_i are the county specific time invariant confounders and the d_t are the shocks common to all counties. The identifying assumption then becomes $E[e_{it}|D_{it}, a_i, d_t] = 0$. This assumption would be violated if one failed to include any confounders that are correlated with deliveries over time within a county.

One could estimate this equation on a sample containing just the counties receiving deliveries or a sample of counties receiving deliveries and include, as a control group, counties that do not receive deliveries. In the material that follows, we show that the estimation results are robust to using either sample. In the first sample, the identifying source of variation is within county time series variation. For the larger sample it is within county variation relative to the control group county variation, which identifies the coefficient of interest b_1.

2.3 Data

The data used in this analysis are comprised of an annual panel data set covering the years 1980 to 2009. Counties included in the dataset, which receive irrigation water from either the CVP or SWP are the following: Fresno, Kern, Kings, Merced, San Joaquin, Stanislaus and Tulare. As a control group, to capture the effects of general changes in the agricultural economy, we use six California counties that do not receive Delta water deliveries: Madera, Imperial, Monterey, Sutter, Yolo and Yuba. The data period covered by our analysis evidences significant variation both in employment and water deliveries. It also includes two of the largest droughts in the recent past (1987-1992 and 2007-2009).

County-level employment data are publicly available, and we obtained them from the Bureau of Economic Analysis (BEA)[2]. We distinguish between direct farm employment, and total agricultural employment. Direct farm employment includes anyone who works in the direct production of agricultural commodities, including crops and livestock (SIC codes 01 – 02; NAICS code 111 - 112)[3]. Total agricultural employment is the sum of direct farm employment and employment in the agricultural services sector (SIC code 07; NAICS code 113 - 115). The agricultural services sector includes farm labor contractors.

The data we used from 1980 – 2000 are categorized in the Standard Industrial Classification (SIC) system. In the 1990s, a new classification system (North American Industrial Classification System (NAICS)) was introduced, in part to facilitate accounting under the North American Free Trade Agreement. The SIC data series was discontinued in 2000. In that year, the BEA shifted to reporting sectoral employment based on the SIC industry classification to reports based on the NAICS classification. The BEA provides a concordance to match industry descriptions between the two coding systems. As we control for year fixed effects in our preferred specification, if there are year-to-year differences in employment that are due to the new classification, our method implicitly controls for these differences.

Government water delivery data include both state deliveries from the State Water Project (SWP) and federal deliveries from the Central Valley Project (CVP). The state water delivery data come from the California Department of Water Resources' Bulletin 132 and the Kern County Water Agency.[4,5] The federal water deliveries data are from

[2]Bureau of Economic Analysis, Regional Economic Accounts, Local Area Personal Income, Table CA25-Total employment by industry, accessed at http://www.bea.gov/regional/reis/default.cfm?selTable=CA25, June 1, 2011.

[3]Bureau of Economic Analysis, Local Area Personal Income Methodology, Appendix: Concordance between BEA industry descriptions and SIC codes, accessed at http://www.bea.gov/regional/pdf/lapi2008/appendix.pdf, February 25, 2011.

[4]California Department of Water Resources, State Water Project Analysis Office, Bulletin 132, Appendix B. Years 1995-2007 accessed at: http://www.water.ca.gov/swpao/bulletin.cfm, July 21, 2009. Years 1973-1994: PDF copies received via electronic communication with DWR, October 14, 2009. Years 2008-2010: Microsoft Excel tables received via electronic communication with State Water Contractors, February 24, 2011.

[5]Kern County Water Agency, SWP Supply and Delivery Summary. Years 1970-2008, received via elec-

the Bureau of Reclamation.[6] We used a Geographic Information System to allocate water deliveries to counties. We first took the intersection of the boundaries of each of the water districts and counties. We then calculated the acreage of the district-county intersection and divided that by the acreage of each of the districts. We multiplied this ratio by the water deliveries in each water district and summed the share of water deliveries in the district-county intersection over counties. Thus, water deliveries are allocated to the county level according to the share of acres of each water district that falls within each county.[7] Annual deliveries are reported in acre-feet.

The data set also includes harvested acres by county. These data come from the Agricultural Commissioners' Offices of Fresno, Imperial, Kern, Kings, Madera, Merced, Monterey, San Joaquin, Stanislaus, Sutter, Tulare, Yolo and Yuba counties for the years 1980 through 2009.[8] We consider land allocated to a subset of crops: almonds, avocados, broccoli, cotton, grapes, hay, lemons, lettuce, oranges, pistachios, rice, strawberries, tomatoes and walnuts. We use a subset of crops for this analysis because acreages are more consistently defined for these individual crops than for total harvested acreage. For example, some counties include rangeland in total area statistics in some years, but not in other years. The crops in our analysis account for roughly two-thirds of total harvested acreage in the San Joaquin Valley.

Table 2.1 displays average employment, water deliveries and harvested acreage by county from 1980 to 2009. Fresno County has the highest number of total employed workers and the highest number of employed farm workers, while Kings County has the lowest in both categories. Fresno County has the largest area harvested while Stanislaus County has the smallest. Kern County is second in terms of harvested acreage. Fresno County has the highest average level of federal and state water deliveries from the Delta. These large differences across counties show the importance of controlling for unobservable differences across counties via a fixed effects estimation strategy.

---

tronic communication with KCWA, September 29, 2009. Year 2009 received via electronic communication with KCWA on February 24, 2011.

[6]US Bureau of Reclamation Mid-Pacific Region Central Valley Operations, Report of Operations Monthly Delivery Tables. Years 1985-2009 accessed at: http://www.usbr.gov/mp/cvo/deliv.html, October 21, 2009. Years 1970-1984: PDF copies received via electronic communication with USBR, November 5, 2009.

[7]Cal-Atlas Geospatial Clearing House, boundaries of "Federal," "State" and "Private" water districts accessed at: http://www.atlas.ca.gov/download.html, May 26, 2009. Boundaries of Counties obtained from ESRI ArcGIS basemap layers.

[8]Source: Various County Crop Reports.

Table 2.1: Summary Statistics by County

County	Total Employment	Direct Farm Employment	Total Ag Employment	Acres Harvested	Total Delta Deliveries
Fresno	363	29	58	844	1029
Kern	279	17	40	597	985
Kings	45	5	8	300	285
Merced	78	10	14	282	144
San Joaquin	228	13	20	261	37
Stanislaus	178	12	18	200	100
Tulare	153	18	37	419	11

Note: Employment is in thousands of jobs, acres harvested is in thousands of acres, and water deliveries is in thousands of acre-feet.

2.4 Estimation Results

2.4.1 Employment Results

We first estimate the effects of water deliveries on employment using the basic model, which does not control for confounders, using the data from 1980 to 2009 on the sample of the seven counties that receive Delta deliveries. As described above, the specification is $N_{it} = a + b_1 D_{it} + e_{it}$. We denote this specification Model (1). The estimated coefficient of b_1 is 0.00981 and is significantly different from zero at the 10% level. The standard errors in this estimation are clustered at the county level to control for serial correlation. In Model (2) we control for time invariant county characteristics by including county fixed effects to the basic model. The specification of this Model (2) is $N_{it} = a_i + b_t D_{it} + e_{it}$, and under this specification the estimated coefficient of b_1 drops to 0.00414. This estimate is statistically different from zero at the 10% level.

In Model (3) we further control for shocks affecting each county in a given year via year fixed effects, or $N_{it} = a_i + d_t + b_1 D_{it} + e_{it}$. Under this specification the estimated coefficient of b_1 drops to 0.00374 and is statistically different from zero at the 10% level with clustered standard errors. Finally, we use the same model and expand the sample to include control counties that do not receive deliveries from the Delta in my Model (4) in Table 2.2. The control areas help to separate out the influence of macroeconomic trends and other broad changes in the agricultural sector. In this model, the coefficient drops slightly to 0.00341 and is statistically significant at the 10% level.

Next we consider the influence of Delta deliveries on total agricultural employment by county. Table 2.3 displays the results of this analysis. The models correspond to those discussed above, with the exception that the dependent variable is the sum of direct farm employment and agricultural service sector employment. As before, Model (4) is our preferred specification since it is estimated on a dataset including a group of control counties. As before, the estimated coefficient on deliveries is positive, this time with a value of 0.00450. It is significant at the 1% level. To put this estimated coefficient into

Table 2.2: Ordinary Least Squares Regression Results of Country Level Farm Employment on Delta Deliveries

Model	(1) Farm Emp.	(2) Farm Emp.	(3) Farm Emp.	(4) Farm Emp.
Delta water deliveries	0.00981	0.00414	0.00374	0.00341
	(0.00479)*	(0.00194)*	(0.00187)*	(0.00166)*
Constant	11,216	13,314	14,360	10,427
	(2,560)***	(716.7)***	(582.3)***	(805.3)***
County FEs	No	Yes	Yes	Yes
Year FEs	No	No	Yes	Yes
Control Counties	No	No	No	Yes
Observations	210	210	210	385
R-squared	0.340	0.070	0.543	0.456
Number of Counties	7	7	7	13

Note: The dataset used in these regressions is comprised of annual data for the period 1980 – 2009. The dependent variable in each regression is farm employment at the county level. Deliveries are acre-feet of water delivered to the districts within a county by the Central Valley Project or State Water Project. Standard errors are clustered at the county level. Coefficients are significantly different from zero at the 1% (***), 5% (**) or 10% (*) level.

perspective, it indicates that reducing water deliveries by 222.22 acre-feet results in the loss of one farm job.

The statistical models above confirm that reductions in water deliveries from the Delta have a statistically significant effect on farm employment in the San Joaquin Valley. To illustrate the size of this effect, we calculate the implied reduction in farm employment from the 2009 water export restrictions as compared to 2005. This calculation is performed using our preferred Model (4) in Tables 2.2 and 2.3. Considering just direct farm employment, the reduction in 2009 deliveries causes an estimated loss of 6,884 jobs, which is equivalent to a 7% decline in employment. Our county-level model therefore is consistent with the proposition that reductions in water supplies in 2009 caused economically and statistically significant losses in employment in the agricultural production sector.

Reductions in water deliveries in 2009 caused even larger losses in total agricultural employment, which includes farm labor contractor employees as well as farm employees. Using the estimated coefficient in Model (4) of Table 2.3, we conclude that the reduction in 2009 deliveries causes an estimated loss of 9,091 jobs. This estimate is significantly different than zero at the 1% level.

Table 2.3: Ordinary Least Squares Regression Results of Country Level Total Agricultural Employment on Delta Deliveries

Model	(1) Total Ag. Emp.	(2) Total Ag. Emp.	(3) Total Ag. Emp.	(4) Total Ag. Emp.
Delta water deliveries	0.0246	0.00467	0.00457	0.00450
	(0.00810)**	(0.000510)***	(0.00127)**	(0.000886)***
Constant	19,674	27,448	25,345	21,569
	(6,065)**	(199.1)***	(1,633)***	(766.6)***
County FEs	No	Yes	Yes	Yes
Year FEs	No	No	Yes	Yes
Control Counties	No	No	No	Yes
Observations	193	193	193	352
R-squared	0.431	0.028	0.508	0.388
Number of Counties	7	7	7	13

Note: The dataset used in these regressions is comprised of annual data for the period 1980 – 2009. The dependent variable in each regression is total agricultural employment at the county level, defined as the sum of direct farm employment and employment in the agricultural services industry. Deliveries are acre-feet of water delivered to the districts within a county by the Central Valley Project or State Water Project. Standard errors are clustered at the county level. Coefficients are significantly different from zero at the 1% (***), 5% (**) or 10% (*) level.

2.4.2 Area Under Cultivation Results

While the models above do not formally test the mechanism of how changes in water deliveries influence job losses, one would expect that acreage planted to crops would decrease if deliveries are short, which would lead to lower labor requirements to service this smaller area. We therefore test whether deliveries are correlated with total acreage cropped in the seven counties in our sample receiving deliveries.

Below we find that there is a strong and statistically robust relationship between water deliveries and area harvested in the San Joaquin Valley. The model specifications are the same as those used in the models explaining direct farm employment and total agricultural employment, only that we use total area cropped in acres as the left hand side variable. The estimated coefficient on deliveries in Model (4) is 0.0797, which is significantly different from 0 at the 1% level with clustered standard errors. This finding suggests that increasing Delta exports in a given year significantly increases the amount of land under cultivation in the relevant counties.

The estimated coefficient in Model (4) of Table 2.4 suggests that over the historical record from 1980 to 2009, a reduction in water deliveries of 12.55 acre-feet causes one additional acre to be fallowed. Recall that the model is estimated based on plantings of a subset of crops accounting for roughly two-thirds of total harvested acreage in the San Joaquin Valley. Accounting for the crops not in the sample, and assuming the same acreage response to changes in water deliveries, it follows that a reduction in water deliveries of around 8.37 acre-feet would result in an extra acre of total fallowing.

It is also instructive to estimate the amount of fallowing caused by the water delivery reductions of 2009 as compared to 2005. Model (4) indicates that roughly 240,000 acres were fallowed in the San Joaquin Valley in 2009 as a result of the water delivery reductions. This figure is calculated by multiplying the change in deliveries between these two years by the coefficient on deliveries in Model (4) and then adjusting for the fact that Model (4) is based on a subset of crops accounting for two-thirds of total acreage.

2.4.3 Heterogeneity of Estimated Effects

There is evidence that the existence of groundwater stocks reduced the economic impacts of the drought in certain regions of California, for example in Kern and Kings Counties. To illustrate this point, we rerun our analysis of direct farm employment allowing for the influence of water deliveries on farm jobs to vary between Kern and Kings Counties and the rest of the San Joaquin Valley. This formulation with a county-specific treatment effect allows one to compare the influence of surface water deliveries on farm employment between areas with differential access to groundwater.

Table 2.5 displays the results of this analysis. The model formulation and variable definitions are exactly as in Table 2.5, with the addition of an interaction term on deliveries that allows deliveries to have a different effect on farm employment in Kern and Kings Counties than in the rest of the sample. This finding is consistent with the mitigating

Table 2.4: Ordinary Least Squares Regression Results of County Level Area Harvested on Delta Deliveries

Model	(1) Acres Harv.	(2) Acres Harv.	(3) Acres Harv.	(4) Acres Harv.
Delta water deliveries	0.404	0.0741	0.0827	0.0797
	(0.0959)***	(0.0135)***	(0.0314)**	(0.0226)***
Constant	265,077	387,186	389,448	287,329
	(49,924)***	(4,989)***	(17,007)***	(16,531)***
County FEs	No	Yes	Yes	Yes
Year FEs	No	No	Yes	Yes
Control Counties	No	No	No	Yes
Observations	210	210	210	385
R-squared	0.683	0.093	0.246	0.251
Number of Counties	7	7	7	13

Note: The dataset used in these regressions is comprised of annual data for the period 1980 – 2009. The dependent variable in each regression is harvested acreage for a subset of crops at the county level. Deliveries are acre-feet of water delivered to the districts within a county by the Central Valley Project or State Water Project. Standard errors are clustered at the county level. Coefficients are significantly different from zero at the 1% (***), 5% (**) or 10% (*) level.

effect of groundwater reserves.

While groundwater extraction can serve as an important buffer against reduced surface water deliveries, reducing the amount of groundwater in storage has economic costs. Because groundwater reserves are a stock as opposed to a flow, there are two economic costs associated with depleting them. The first cost is the expenditures of capital and energy required to bring groundwater to the surface. The second type of cost is user cost. Unlike pumping lift costs, user cost does not entail an actual monetary outlay, but rather relates to the fact that groundwater pumping in any given year increases the cost of future pumping and, if groundwater reserves are finite, may limit the amount of groundwater that can be extracted in the future. Also included in user cost is the risk of land subsidence caused by overdraft.

The costs of groundwater overdraft can be large. Pumping lift costs depend on groundwater elevations and other factors such as pump efficiency and energy prices. User costs also depend on specific hydrogeologic conditions as well as contractual and legal limitations on groundwater extraction. We should note that Kern County has a highly developed system of groundwater banking, with a careful accounting system and rules limiting the amount of water that can be withdrawn from storage. Taking account of all these factors, and also the amount of groundwater that was withdrawn in the San Joaquin Valley during the drought to compensate for reduced surface water deliveries, we conclude that the economic costs of reliance on groundwater are large, easily reaching into the hundreds

Table 2.5: Ordinary Least Squares Regression Results of County Level Farm Employment on Delta Deliveries Including Interaction Term of Deliveries with Kern and Kings County

Model	(1) Farm Emp.	(2) Farm Emp.	(3) Farm Emp.	(4) Farm Emp.
Delta water deliveries	0.0156	0.00643	0.00584	0.00436
	(0.00252)***	(0.000997)***	(0.000581)***	(0.000567)***
Interaction	-0.0117	-0.00461	-0.00433	-0.00557
(Deliveries and Kern or Kings)	(0.00288)***	(0.000998)***	(0.000748)***	(0.000680)***
Constant	11,207	13,301	14,245	10,490
	(2,428)***	(188.3)***	(886.8)***	(506.1)***
County FEs	No	Yes	Yes	Yes
Year FEs	No	No	Yes	Yes
Control Counties	No	No	No	Yes
Observations	210	210	210	385
R-squared	0.547	0.092	0.561	0.472
Number of Counties	7	7	7	13

Note: The dataset used in these regressions is comprised of annual data for the period 1980 – 2009. The dependent variable in each regression is farm employment at the county level. Deliveries are acre-feet of water delivered to the districts within a county by the Central Valley Project or State Water Project. Standard errors are clustered at the county level. Coefficients are significantly different from zero at the 1% (***), 5% (**) or 10% (*) level.

of millions of dollars annually. These withdrawals helped mitigate the effects of reduced surface deliveries, but the costs of groundwater extraction, as well as the long-term limitations of groundwater overdraft, are real and should be borne in mind when considering the effect of future reductions in surface deliveries to the San Joaquin Valley.

2.5 Conclusions

In this paper we show evidence of an economically and statistically significant effect of irrigation water deliveries from California's state and federal water projects on both county level employment as well as area cropped. We show that for a shortage similar to that experienced in 2009, California's agricultural employment in these counties would be lowered by roughly 9000 jobs. Further, we show that the likely mechanism through which this operates is the fallowing of large swatch of farmland. For the same reduction in irrigation water deliveries we estimate that 240,000 acres would be fallowed.

Importantly, we find that employment and area cropped in areas with good groundwater resources appear to be less sensitive to irrigation water deliveries than areas with less favorable groundwater. This should not be regarded as good sign, as withdrawal of groundwater has both private as well as higher social costs, which should be taken into account when evaluating the full costs of irrigation water shortages.

Bibliography

[1] Alex Anas and Robin Lindsey. Reducing Urban Road Transportation Externalities: Road Pricing in Theory and Practice. *Review of Environmental Economics and Policy*, 5(1):66 – 88, 2011.

[2] Alex Anas and Hyok-Joo Rhee. Curbing excess sprawl with congestion tolls and urban boundaries. *Regional Science and Urban Economics*, 36(4):510 – 541, 2006.

[3] Henrik Andersson, James Hammitt, Gunnar Lindberg, and Kristian Sundstrom. Willingness to Pay for Car Safety: Sensitivity to Time Framing. Working Papers 2008:8, Swedish National Road & Transport Research Institute (VTI), 2008.

[4] Richard Arnott and Kenneth Small. The Economics Of Traffic Congestion. (256), December 1993.

[5] Maximilian Auffhammer and Richard T. Carson. Forecasting the path of China's CO2 emissions using province-level information. *Journal of Environmental Economics and Management*, 55(3):229 – 247, 2008.

[6] Maximilian Auffhammer and Ryan Kellogg. Clearing the Air? The Effects of Gasoline Content Regulation on Air Quality. *American Economic Review*, 101(6):2687–2722, October 2011.

[7] Orange Country Transportation Authority. The Model of Success: the 91 Express Lanes, Annual Report. Technical report, Orange Country Transportation Authority, 2006.

[8] William J. Baumol and Wallace E. Oates. *The Theory of Environmental Policy*. Cambridge University Press, 2nd edition, 1988.

[9] Antonio Bento, Jonathan Hughes, and Daniel Kaffine. Carpooling and Driver Responses to Fuel Price Changes: Evidence from Traffic Flows in Los Angeles. *Working Paper*, 2012.

[10] Stuart H. Berness and James P. Quirk. Appropriative Water Rights and the Efficient Allocation of Resources. *American Economic Review*, 69(1):25–37, March 1979.

[11] Anthony E. Boardman and Lester B. Lave. Highway congestion and congestion tolls. *Journal of Urban Economics*, 4(3):340 – 359, 1977.

[12] David Brownstone, Arindam Ghosh, Thomas F. Golob, Camilla Kazimi, and Dirk Van Amelsfort. Drivers' willingness-to-pay to reduce travel time: evidence from the San Diego I-15 congestion pricing project. *Transportation Research Part A: Policy and Practice*, 37(4):373 – 387, 2003.

[13] Jan K. Brueckner. Urban Growth Boundaries: An Effective Second-Best Remedy for Unpriced Traffic Congestion? 2005.

[14] Meghan R. Busse, Christopher R. Knittel, and Florian Zettelmeyer. Pain at the Pump: The Differential Effect of Gasoline Prices on New and Used Automobile Markets. (15590), December 2009.

[15] Agricultural Issues Center. Water Supply and Demand. *University of California Agricultural Issues Center*, November 2009.

[16] Janet Currie and Reed Walker. Traffic Congestion and Infant Health: Evidence from E-ZPass. *American Economic Journal: Applied Economics*, 3(1):65–90, January 2011.

[17] Lucas W. Davis. The Effect of Driving Restrictions on Air Quality in Mexico City. *Journal of Political Economy*, 116(1):pp. 38–81, 2008.

[18] UC Davis. Watershed Information. Technical report, 2010.

[19] Elizabeth Deakin and Karen Frick. Bay Bridge Toll Evaluation: Final Report. Technical report, Global Metropolitan Studies and the University of California Berkeley, November 2011.

[20] Olivier Deschenes and Michael Greenstone. Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US. Working Paper 13178, National Bureau of Economic Research, June 2007.

[21] Comune di Milano. Monitoraggio Indicatori Ecopass. Technical report, Comune di Milano, December 2008.

[22] Andrew P. Dickerson, John Peirson, and Roger Vickerman. Road Accidents and Traffic Flows: An Econometric Investigation. *Economica*, 67(265).

[23] Gilles Duranton and Matthew A. Turner. The Fundamental Law of Road Congestion: Evidence from US Cities. *American Economic Review*, 101(6):2616–52, September 2011.

[24] Molly Espey. Explaining the Variation in Elasticity Estimates of Gasoline Demand in the United States: A Meta-Analysis. *The Energy Journal*, 17(3):pp. 49–60, 1996.

[25] Gordon J. Fielding. Private Toll Roads: Acceptability of Congestion Pricing in Southern California. (qt749118j3), January 2001.

[26] Amy Finkelstein. E-ZTax: Tax Salience and Tax Rates. *Quarterly Journal of Economics*, 123:3:969–1010, 2009.

[27] Meredith Fowlie. Emissions Trading, Electricity Restructuring,and Investment in Pollution Abatement. *American Economic Review*, 100:837–869, 2010.

[28] Jacqueline Geoghegan. *The Road Not Taken: Environmental Congestion pricing on the San Francisco-Oakland Bay Bridge*. PhD thesis, UC Berkeley, 1995.

[29] J.A. Hewitt and W.M. Hanemann. A discrete continuous choice approach to residential water demand under block rate pricing. *Land Economics*, 71(2):173–192, 1995.

[30] Wolfgang S. Homburger, Jerome W. Hall, William Reilly, and Edward C. Sullivan. *Fundamentals of Traffic Engineering*. Institute of Transportation Studies, UC Berkeley, 15th edition, 2001.

[31] Richard Howitt, Duncan MacEwan, and Josue Medellin-Azuara. Economic Impact of Reductions in Delta Exports on Central Valley Agriculture. *Agricultural and Resource Economics Update*, 12(3), Jan/Feb 2009.

[32] Richard Howitt, Duncan MacEwan, and Josue Medellin-Azuara. Measuring the Employment Impact of Water Reductions. Technical report, September 2009. Department of Agricultural and Resource Economics and Center for Watershed Sciences, University of California, Davis.

[33] Jonathan Hughes, Christopher Knittel, and Daniel Sperling. Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand. *The Energy Journal*, 29:93–114, 2008.

[34] Christopher R. Knittel. Reducing Petroleum Consumption from Transportation. (17724), January 2012.

[35] Jonathan Levy, Jonathan Buonocore, and Katherine von Stackelberg. Evaluation of the public health impacts of traffic congestion: a health risk assessment. *Environmental Health*, 9(1):65, 2010.

[36] C. Robin Lindsey, Vincent A.C. van den Berg, and Erik T. Verhoef. Step by Step: Revisiting Step Tolling in the Bottleneck Model. (10-118/3), November 2010.

[37] Robin Lindsey. Do Economists Reach A Conclusion on Road Pricing? The Intellectual History of an Idea. *Econ Journal Watch*, 3(2):292–379, May 2006.

[38] Robin Lindsey. Cost recovery from congestion tolls with random capacity and demand. *Journal of Urban Economics*, 66(1):16–24, July 2009.

[39] Louie Nan Liu and John F. McDonald. Efficient Congestion Tolls in the Presence of Unpriced Congestion: A Peak and Off-Peak Simulation Model. *Journal of Urban Economics*, 44(3):352–366, November 1998.

[40] Louie Nan Liu and John F. McDonald. Economic efficiency of second-best congestion pricing schemes in urban highway systems. *Transportation Research Part B: Methodological*, 33(3):157–188, April 1999.

[41] Daniel McFadden. The Measurement of Urban Travel Demand. *Journal of Public Economics*, 3:303–328, 1974.

[42] Jeffrey Michael. Employment Impacts of Reduced Water Supplies to San Joaquin Valley Agriculture. *Eberhardt School of Business: Business Forecasting Center, University of the Pacific*, December 2009.

[43] Jeffrey Michael. Unemployment in the San Joaquin Valley in 2009: Fish or Foreclosure? *Eberhardt School of Business: Business Forecasting Center, University of the Pacific*, August 2009.

[44] Department of Water Resources. Bulletin 132: Management of the California State Water Project. December 2007.

[45] Department of Water Resources. California's Drought Update. Technical report, March 2010.

[46] Ian W.H. Parry. Pricing Urban Congestion. *Resources for the Future Discussion Paper*, 2008.

[47] Andrea Schrage. Traffic Congestion and Accidents. (419), 2006.

[48] David Schrank, Tim Lomax, and Bill Eisele. Urban Mobility Report. Technical report, Texas Transportation Institute, September 2011.

[49] Steve Sexton. Paying for Pollution? How General Equilibrium Effects Undermine the "Spare the Air" Program. *Job Market Paper*, 2010.

[50] Kenneth A. Small. The Scheduling of Consumer Activities: Work Trips. *The American Economic Review*, 72(3):pp. 467–479, 1982.

[51] Kenneth A. Small. The incidence of congestion tolls on urban highways. *Journal of Urban Economics*, 13(1):90 – 111, 1983.

[52] Kenneth A. Small and Xuehao Chu. Hypercongestion. *Journal of Transport Economics and Policy*, 37(3):319–352, September 2003.

[53] Kenneth A. Small, Clifford Winston, and Jia Yan. Uncovering the Distribution of Motorists' Preferences for Travel Time and Reliability. *Econometrica*, 73(4):1367–1382, July 2005.

[54] Lester A. Snow. Lester A. Snow. Technical report, May 2009.

[55] David Sunding, Ajami Newsha, Steve Hatchet, David Mitchell, and David Zilberman. Economic Impacts of the Wanger Interim Order for Delta Smelt. *Berkeley Economic Consulting*, December 2008.

[56] Martin Treiber and Arne Kesting. Impact of traffic congestion on fuel consumption and emissions. *Networks for Mobility*, 2008.

[57] USDOT. Lee County Variable Bridge Tolls. Technical report, USDOT, November 2006.

[58] USDOT. Congestion Pricing: A Primer, 2008.

[59] USDOT. Transportation Vision 2030. Technical report, USDOT, January 2008.

[60] J. Vernon Henderson. Peak shifting and cost-benefit miscalculations. *Regional Science and Urban Economics*, 22(1):103–121, March 1992.

[61] William S. Vickrey. Pricing in Urban and Suburban Transport. *The American Economic Review*, 53(2):pp. 452–465, 1963.

[62] William S. Vickrey. Congestion Theory and Transport Investment. *The American Economic Review*, 59(2):pp. 251–260, 1969.

[63] Philip A. Viton. Equilibrium Short-Run-Marginal-Cost Pricing of a Transport Facility: The Case of the San Francisco Bay Bridge. *Journal of Transport Economics and Policy*, 14(2):pp. 185–203, 1980.

[64] James M. Whitty. Oregon's Mileage Fee Concept and Road User Fee Pilot Program: Final Report, 2007.

[65] Frank A. Wolak. An Experimental Comparison of Critical Peak and Hourly Pricing: The PowerCentsDC Program. *Preiminary Draft*, 2010.

Appendix A

Appendix

A.1 Other Examples of Congestion Pricing

Table A.1: Examples of previously existing road congestion pricing

Location	Description of Pricing	Reported Outcomes[a]
San Diego, CA Interstate 15 since 1998	Non-HOV can pay to use HOV lanes. Price are adjusted dynamically with traffic in 25 cent increments.	Saves commuters 30 minutes on their commute each way.[b]
Orange County, CA State Route 91 since 1995	Four variably-priced lanes run parallel to unpriced lanes for 10 miles. Toll schedule is pre-set and adjusted every 3 months based on previous demand.[c]	In peak hours, speeds in priced lanes are 60-65mph and speeds in unpriced lanes are 15-20mph. During Friday peak (5-6pm), each priced lane serves twice as many vehicles as each unpriced lane.
Lee County, FL Two bridges since 1998	Two toll bridges offer 50% off toll charges for weekday travel during pre- and post-peak hours.	5% shift from peak to off-peak travel.[d]
London, England Central London since 2003	Per day charge of £8 for travel within a central London zone along with improved public transit.	Decrease of 15% in traffic in central London. Travel delays reduced by 30%. Waiting time for buses reduced by one-third.
Singapore since 1975	Variable time-of-day pricing on expressway system and peak pricing in morning rush hour.	Reduced traffic by 13% and increased vehicle speed by 22%.
Stockholm City center since 2006	Cordon pricing in the city center.	Decreased travel times, vehicle trips, accidents, and emissions. Increase in ridership on inner city buses.
Milan City Center since 2008	Charge of E2 to E10 to enter the city center on weekdays between 7:30am and 7:30pm.	12.3% reduction in traffic and 14-47% reduction in various atmospheric pollutants.[e]

[a] Unless noted, outcome estimates come from USDOT (2008).
[b] Estimates from Orange Country Transportation Authority (2006)
[c] For example, as of January 1, 2013, the toll for going east on Fridays from 3:00-4:00pm is set at $9.55; and going west on Mondays - Thursdays from 7:00-8:00am is $4.75. There are special tolls for major holidays. Source: http://www.octa.net/91_schedules.aspx
[d] Estimates from USDOT (2006)
[e] Vehicle miles traveled.

Figure A.1: San Mateo Bridge traffic volume over time

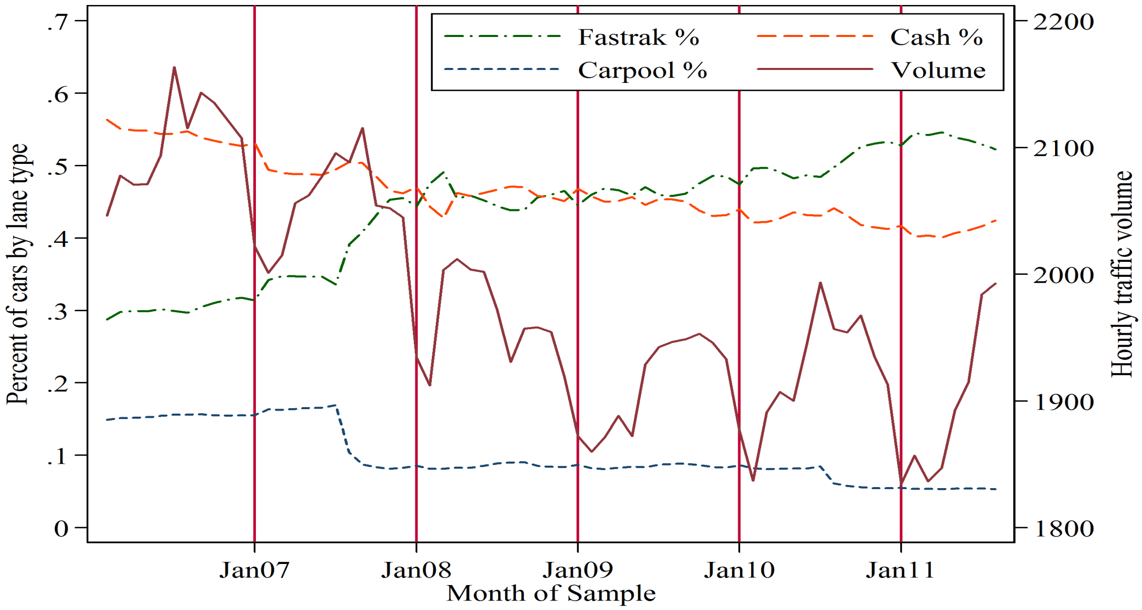

Jan 2006 - July 2011. Weekdays, Westbound. Excludes holidays.

A.2 Traffic Volume Data

Data on traffic volume over time on the San Mateo Bridge can be seen graphically in figure A.1. The monthly mean of hourly traffic volume fluctuates between 1825 vehicles per hour to 2175 vehicles per hour, depending on the season and year, and display a decreasing trend over time. Similar to the Bay Bridge, traffic is generally higher in the summer and lower in the winter. The percent of vehicles using the carpool lanes decreases in steps from about 18% to around 10% midway through 2007, when there were lane regonfigurations, to under 10% around the time of the toll change. The percent of vehicles using FasTrak lanes is under 30% at the begining of 2006, but increases to almost 50% by mid 2011. The percent of vehicles using cash lanes decreases from under 60% to around just over 40%.

The analogous graph for the Dumbarton Bridge is in figure A.2. The monthly mean of hourly traffic volume fluctuates between 1200 vehicles per hour to 1500 vehicles per hour, depending on the season and year, and display a decreasing trend over time. Similar to the Bay and San Mateo Bridges, traffic is generally higher in the summer and lower in the winter. The percent of vehicles using the carpool lanes decreases from below 10% to closer to 5% around the time of the toll change. The percent of vehicles using FasTrak lanes is around 30% at the begining of 2006, but increases to around 45% by mid 2011. The percent of vehicles using cash lanes changes from over 60% to around 45%.

Figure A.2: Dumbarton Bridge traffic volume over time

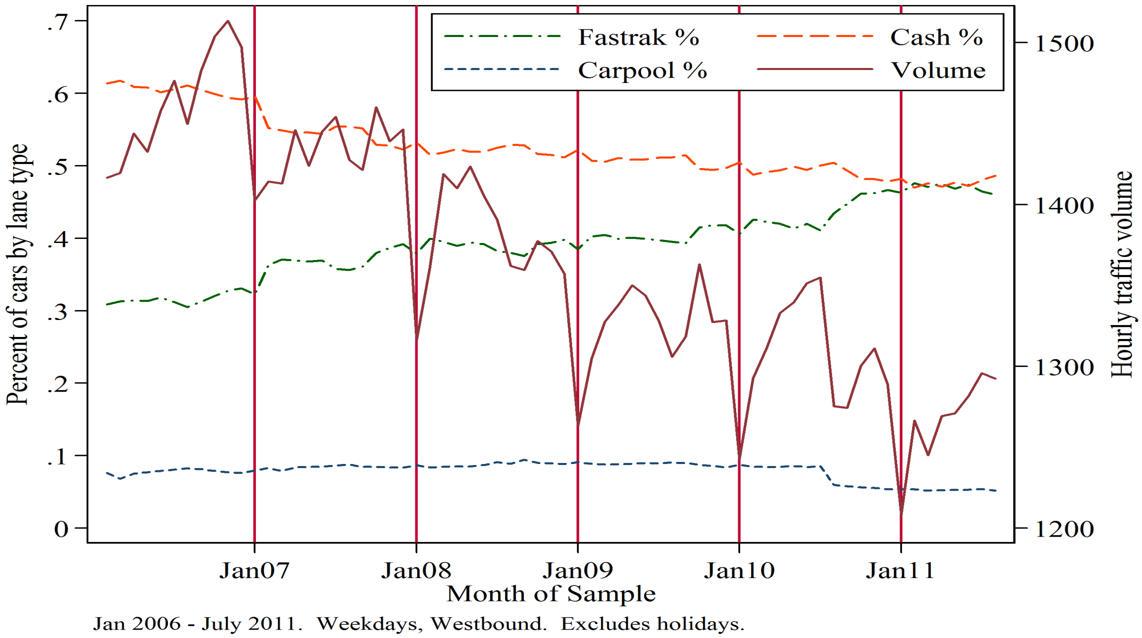

A.3 Robustness Checks

A.3.1 Using Logs Instead of Levels

My first robustness check is to use logs instead of levels of my outcome variables. The results are quantitatively and qualitatively the same as my estimates done in levels.

A.3.2 Using 11:00pm - 3:00am as the Control (instead of midnight)

My second robustness check is to use 11:00pm till 3:00am as a control (instead of the midnight hour). This is to expand the control to include more hours that might capture more of the changes that might have happened apart from the toll policy change (e.g. population growth, fuel prices, etc.). Again, the results are quantitatively and qualitatively similar to the results just using the midnight hour as a control.

A.3.3 Triple Difference Approach

The difference in differences approach requires the SUTVA to hold in addition to the treated and control hours having similar trajectories before and after the policy change (absent treatment). Using the Dumbarton Bridge as an additional control and taking a

Figure A.3: Bay Bridge point estimates of traffic volume treatment effect (logs)

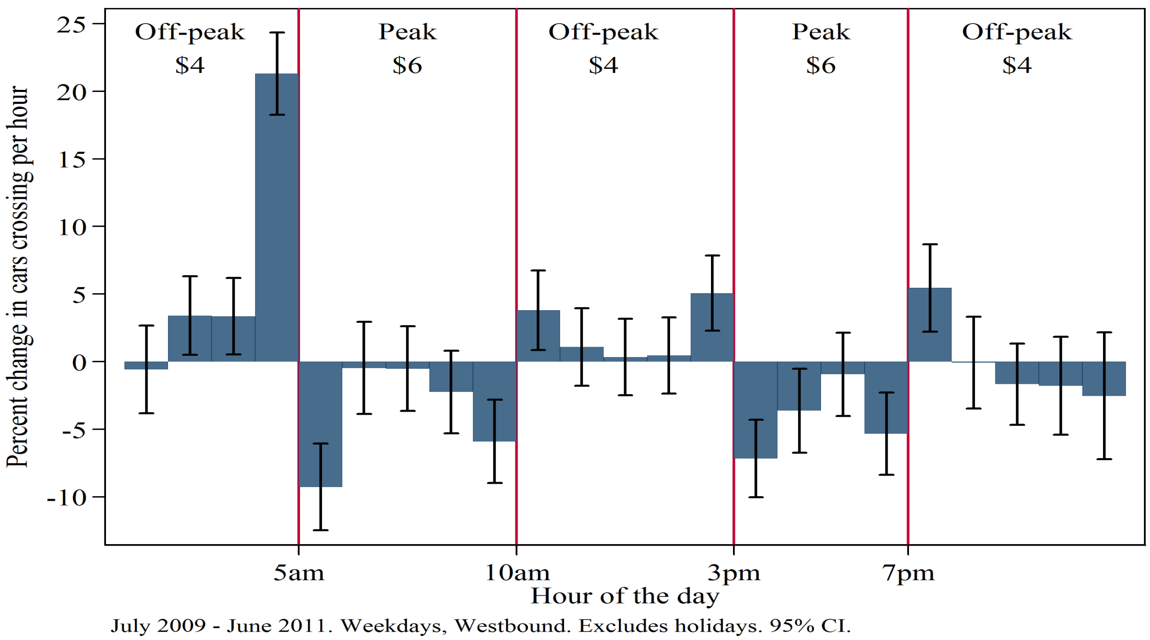

Figure A.4: Bay Bridge point estimates of average treatment effect using 11pm-3am as control

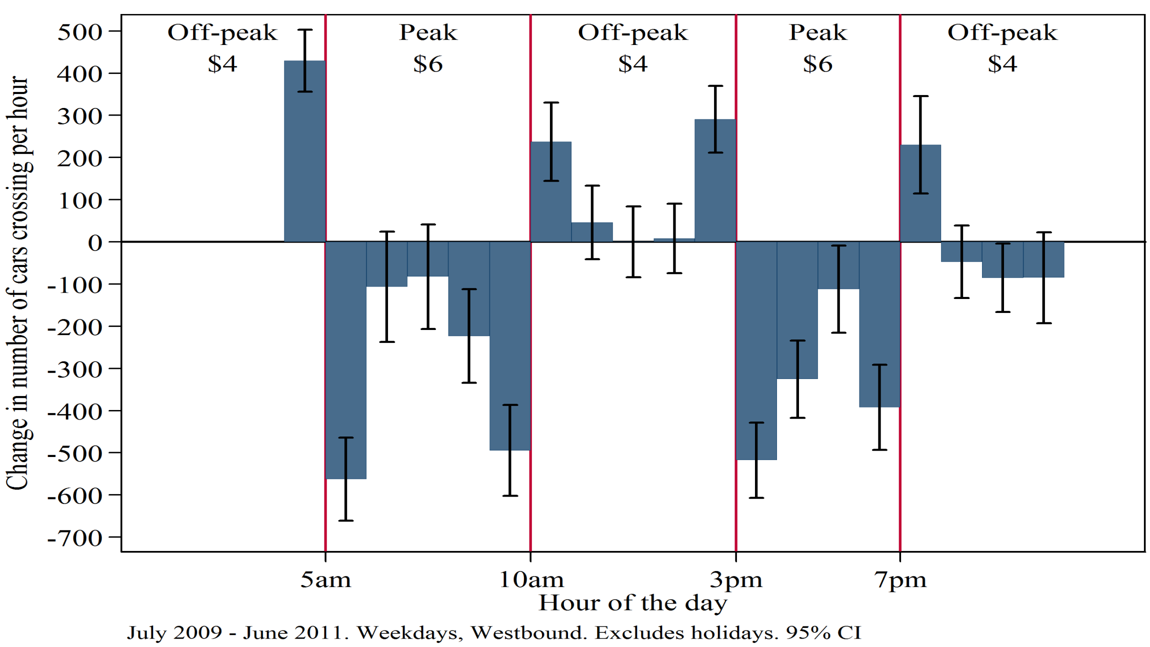

Figure A.5: Bay Bridge, triple difference of traffic volume treatment effect

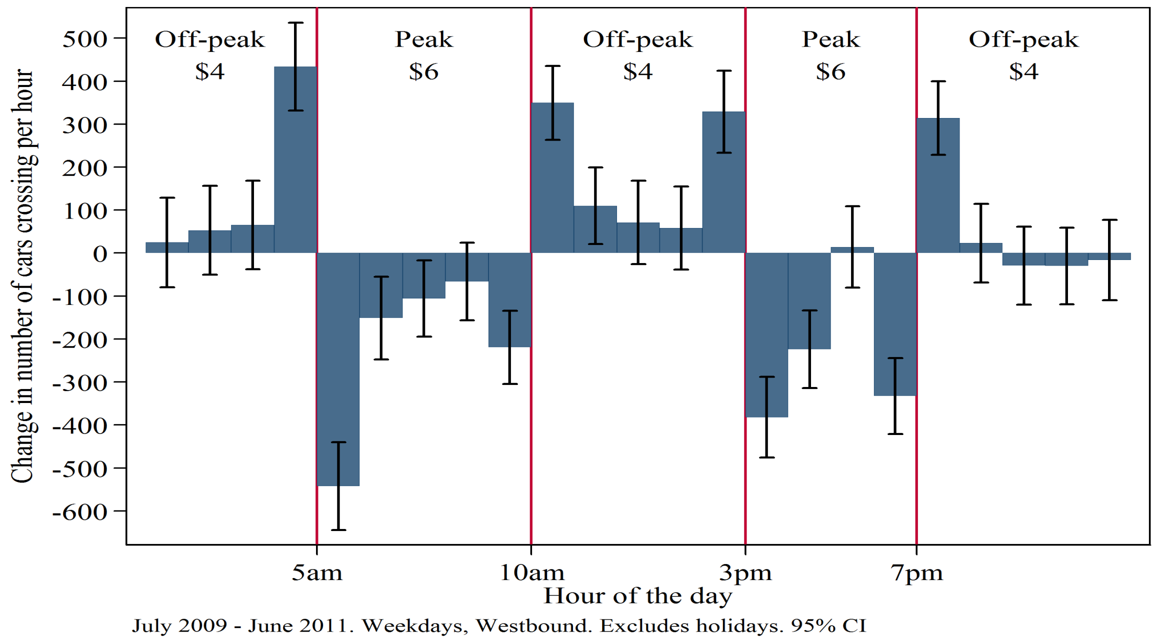

triple difference approach allows me to relax the assumption that requires similar trajectories; I now only require that, to the extent that outcomes for each hour evolve differently over time, the differences affect each hour on the Bay and Dumbarton Bridges similarly.

Thus, my estimating equation becomes:

$$\begin{aligned} Y_{hdwb} &= \alpha_0 + \boldsymbol{\alpha_h} + \beta_0 after_w + \boldsymbol{\beta_h} after_w + \gamma_d + \theta_w + \boldsymbol{\delta_0} Bay + \boldsymbol{\delta_h} Bay + \\ &\quad \tau_0 after_w Bay + \boldsymbol{\tau_h} after_w Bay + \phi_d Bay + \rho_w Bay + \varepsilon_{hdwb} \end{aligned}$$

Where the only changes from the difference-in-differences approach are the interactions of the terms with the binary indicator Bay, which is equal to one for the Bay Bridge, resulting in the τ_h's being my coefficients of interest, where hour of the day (1-23) is interacted with $after$ and Bay. The τ_h's have a causal interpretation of the impact of the congestion pricing on the change in traffic volume and travel time.

Traffic Volume The results for the triple difference for traffic volume on the Bay Bridge are shown in figure A.5. They show time-shifting behavior from peak shoulder hours to off-peak shoulder hours as well as decreases in peak and increases in off-peak traffic volume.

Travel Time The results for travel time on the Bay Bridge are shown in figure A.6. They show similar results to the DD approach.

Figure A.6: Bay Bridge, triple difference of travel time treatment effect

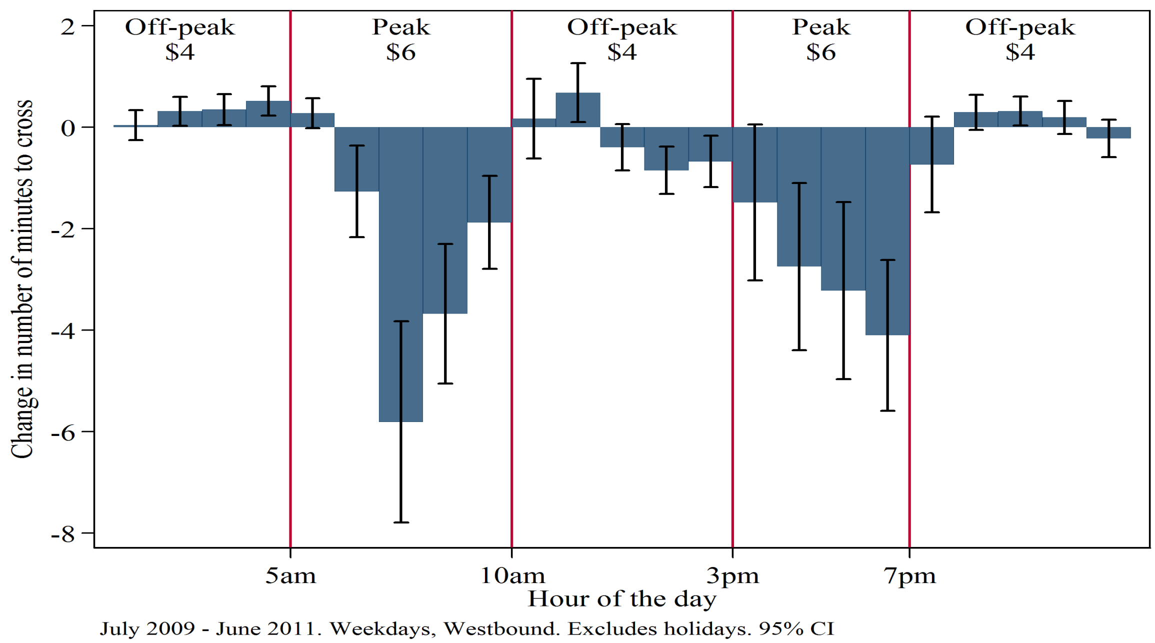

A.3.4 Falsification Tests Using a Previous Toll Change

The last toll change on the Bay Area bridges happened on January 1, 2007, when all prices increased by \$1 from \$3 to \$4 for all three of the bridges in my analysis. I use this change to test how traffic volume changed after that toll change. The results for the Bay Bridge are shown in figure A.7. As theory would predict with a price increase, most hours experience a decrease in traffic volume, although most estimates are statistically insignificant. The morning peak hours experience small and positive and sometimes statistically significant changes. There results suggest that morning travelers (probably commuters) are less sensitive to price changes and there might even be latent demand from drivers who think the travel time will decrease due to the higher toll.

Figure A.7: Bay Bridge, falsification test of traffic volume treatment effect using 2007 toll change