Title
Essays in the Economics of Transportation and the Environment

Author
Mahmassani, Amine

Publication Date
2018

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA,
IRVINE

Essays in the Economics of Transportation and the Environment

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Economics

by

Amine Mahmassani

Dissertation Committee:
Professor David Brownstone, Chair
Associate Professor Michael McBride
Economist Dr. Kevin Roth

2018

DEDICATION

To my mother Millicent Kushner Mahmassani
who will always be my comfort

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I would like to thank committee members David Brownstone, Kevin Roth, and Michael McBride for their guidance and support throughout my graduate studies. It was truly a privilege to learn from them, and to have had them invest so much time and energy in my success. I wish to thank my committee chair, David Brownstone, in particular for being such a strong advocate for me during my time in program. I would also like to thank Kevin Roth and Damon Clark for how much I got out of the courses they taught. Additional thanks to Jan Brueckner, who seems to genuinely care about everyone in the Econ program.

I would like to thank my friend, classmate, and collaborator, Si-Yuan Kong, for all his support. I would not have made it through the program without him. I also wish to thank my classmates Mario Silva, Debapriya Chakraborty, Karina Hermawan, Michael Guggisberg, Sierra Garrett, Ian Burn, Jason Ralston, Jessica Monnet, Paul Stroik, and Michael Shin for their friendship and help during the program. Additional thanks to my friend Sinnott Murphy, who I knew was always there for me even when we weren't in touch.

I would like to thank John Sommerhauser, Pat Frazier, Adam Cook for going above and beyond to help with any sort of administrative issue I encountered. Also thanks to Melissa Valdez for working so hard to make sure the job market ran smoothly.

I am grateful to my Uncles Ben Kushner and Adam Nagourney for making sure I had a birthday dinner to look forward to every year. Special thanks to my father, Hani Mahmassani, the most amazing person I know, for how much he gives without ever asking for anything in return. Special thanks to my brother Ziad Mahmassani for his unconditional love and companionship. Special thanks to Kelsey Heider, for being my best friend and partner.

I would like to thank the Institute of Transportation Studies and director Steve Ritchie for their generous funding of my research through Pacific Southwest Region UTC fellowships. I would also like to thank the University of California Center on Economic Competitiveness in Transportation for their generous financial support of my research (UCCONNECT Award # 65A0529, Task #TO018 and #TO050). I am also grateful to the Center for Economics and Public Policy, along with Steve Borowski, for funding my first research experiment.

CURRICULUM VITAE

Amine Mahmassani

RESEARCH AND TEACHING INTERESTS

Transportation Economics, Environmental Economics, and Experimental Economics

EDUCATION

Ph.D. Candidate, Department of Economics, University of California - Irvine, 2018

M.S., Transportation Technology and Policy, University of California – Davis, 2010

B.S., Electrical Engineering, University of Maryland – College Park, 2005

RESEARCH

Working Papers

"Do truthful-bidding mechanisms improve the allocation of drivers to express lanes?," 2017

"Learning Equilibria in Congested Multi-State Networks," 2017

"The Effect of Ambient Air Pollution on the Frequency of Motor Vehicle Collisions," 2017

Publications

"An Experimental Study of Route Choice and Incident Management Using a Real-Time 2D Driving Simulator with Variable Message Signs," co-authored with Si-Yuan Kong, David Brownstone, and Michael McBride, *TRB 96th Annual Meeting Compendium of Papers - 2017*

"Experimental Studies for Traffic Incident Management," co-authored with Si-Yuan Kong, David Brownstone, and Michael McBride, *trid.trb.org - 2016*

"State of charge: Electric vehicles' global warming emissions and fuel-cost savings across the United States", co-authored with Don Anair," *Union of Concerned Scientists – 2012*

Research Experience

Graduate Student Researcher on Experimental Studies for Traffic Incident Management, supervised by David Brownstone and Michael McBride funded by two consecutive UC Connect Faculty Research Grants (Spring 2015 – Spring 2017).

Vehicles Analyst/Engineer at Union of Concerned Scientists, studying climate impacts of electric vehicle charging and unconventional oil production (January 2011 – August 2012).

Graduate Student Researcher on Potential Design, Implementation, and Benefits of a Feebate Program for New Passenger Vehicles in California supervised by David Bunch at University of California – Davis (2007 - 2009)

GRANTS, AWARDS, AND FELLOWSHIPS

Grants and Fellowships

Pacific Southwest Region University Transportation Center Graduate Fellowship, 2017

Co-wrote proposal for UCCONNECT Award # 65A0529, Task # TO050 with another student, 2016-2017

Co-wrote proposal for UCCONNECT Award # 65A0529, Task # TO018 with another student, 2015-2016

Center for Economics and Public Policy Borowski Award, Fall 2013

Awards

Ken Small Award for Best PhD Student in Transportation and Urban Economics, 2016-2017

CONFERENCE AND SEMINAR PRESENTATIONS

Conferences

Southwest Experimental and Behavioral Economics Workshop, 2017, Learning Equilibria in Congested Multi-State Networks

Economic Science Association World Meeting, 2017, Learning Equilibria in Congested Multi-State Networks (accepted)

Transportation Research Board Annual Meeting, 2017, Learning Equilibria in Congested Multi-State Networks

Seminars

University of California at Irvine, Department of Economics, Transportation, Urban & Environmental Seminar, Fall 2017

University of California at Irvine, Institute for Mathematical Behavioral Sciences Seminar, Spring 2013

ABSTRACT OF THE DISSERTATION

Essays in the Economics of Transportation and the Environment

By

Amine Mahmassani

Doctor of Philosophy in Economics

University of California, Irvine, 2018

Professor David Brownstone, Chair

This thesis uses applied econometrics and traffic experiments to identify environmental and behavioral factors that contribute to externalities in traffic networks, as well as evaluate mechanisms designed to address them.

The first chapter examines whether exposure to ambient fine particulate matter (PM 2.5) increases the likelihood of getting into a vehicle collision. PM 2.5 has been shown to affect alertness and cognition, which may in turn impair driving ability. Variation in daily AQI level from PM 2.5 was exploited to identify a possible causal effect on daily car accident rates in nearby cities. This approach yielded no evidence of a causal effect on vehicle accidents, perhaps due to endogeneity of PM 2.5 with other factors correlated with accident frequency. An alternative instrumental variables approach exploited exogenous shifts in wind direction relative to nearby coal power plants - a significant point source of PM 2.5. This specification found that a one-standard deviation in PM 2.5 AQI increases the car accident rate by 13.2 percent.

The second chapter investigates if the presence of multiple states in traffic networks adversely impacts the speed at which users learn route-choice equilibria. To answer this question, data were generated from several sessions of a repeated binary route-choice experiment with human subjects. Exogenous random state changes were introduced as discrete, varied reductions in roadway capacity. The

sessions were comprised of either a "simple" network treatment with only two states, or a "complex" network treatment with five states. Reinforcement learning models estimated from the experimental data show that learning was significantly impaired in the complex five-state treatment but not the simple two-state treatment. Simulations based on the learning behavior estimated from each treatment showed that the impaired learning from the five-state treatment resulted in disproportionately slower (and sometimes non-existent) equilibrium convergence compared to learning with two-states.

This third chapter demonstrates the workability of a truth-telling mechanism for efficiently allocating freeway capacity. I conduct a traffic experiment using an interactive multi-user driving simulator in which I allocate human subject drivers to freeway lanes using an optimal tolling scheme where users reveal their valuation of the road through a Vickrey-Clarke-Groves mechanism. I find that the mechanism generally elicits truthful values of time from subjects. However, there are also significant and persistent deviations from truth-telling caused largely by difficulty in understanding the complexity of the mechanism as well as stochasticity in travel time outcomes. Nevertheless, I show that the mechanism dominates alternatives under a plausible set of assumptions.

CHAPTER I

THE EFFECT OF AMBIENT AIR POLLUTION ON THE FREQUENCY OF MOTOR VEHICLE COLLISIONS

I.I INTRODUCTION

Motor vehicle crashes cause substantial economic losses in the United States; in 2010 alone the economic costs from crashes totaled $277 billion dollars in vehicle damage, congestion, and injury/death (Blincoe et al. 2014). Although significant resources are devoted by governments, transportation agencies, and industry to improve auto transportation safety, nobody to date has considered the possible role that air pollution might play in the causation of automobile accidents. Drivers are regularly exposed to many acutely harmful pollutants such as particulate matter (PM) and nitrogen oxides (NOx), and although the most clinically significant known effects of these pollutants are respiratory and cardiac related, there is a sparse literature linking PM 2.5 exposure to central nervous system inflammation (Stanek et al. 2011). Furthermore, a growing body of literature provides evidence of negative short-term productivity and cognitive outcomes associated with the inhalation of PM less than 2.5 microns in width (PM 2.5). For example, Chang et al. (2014) find PM 2.5 to be responsible for declines in the productivity of indoor pear packers, and Lavy et al. (2012) find declines in test scores resulting from exposure to the pollutant.

This work seeks to explore whether subtler impacts of exposure to ambient PM 2.5 such as fatigue and cognitive impairment are sufficient to increase the likelihood of getting into a vehicle collision. Driving is a cognitively demanding task that requires alertness, quick decision-making, and fast reaction times. Factors impairing mental alertness such as fatigue and drug/alcohol consumption are shown to be a major cause of motor vehicle accidents (National Highway Traffic Safety Administration 2003), so it is plausible that exposure to air pollution might also play a role. If this is the case, then pollution abatement has the potential to reduce the frequency of accidents and confer significant benefits in the form of avoided damages.

Driver Exposure to PM 2.5

For PM 2.5 to have an impact on accidents, drivers would need to be exposed to sufficiently high levels of the pollutant. In many cities, daily ambient PM 2.5 often attains an air quality index (AQI) - a piecewise linear function of a pollutant level intended to characterize its health effect - above 50, which according to the US EPA indicates “Air quality is acceptable; however, for some pollutants there may be a moderate health concern for a very small number of people who are unusually sensitive to air pollution” (Air Quality Index). PM 2.5 also occasionally reaches levels where the AQI exceeds 100, which implies that “members of sensitive groups may experience health effects. The general public is not likely to be affected.” Due to the regional nature of PM 2.5 (U.S. EPA 2004), large numbers of people are affected on days when elevated levels are recorded. In California in 2006, exposure to PM 2.5 levels exceeded the daily federal standard, corresponding to an AQI greater than 100, for hundreds of millions of person-days (California Department of Public Health 2012). These particulates readily enter indoor spaces due to their small size, and especially vehicles due to the high rate of air exchange while driving (Liu & Frey 2011). This ensures exposure for drivers both before and during the act of driving.

Although it is rare for the AQI from PM 2.5 to exceed 150, at which point “everyone may begin to experience health effects; members of sensitive groups may experience more serious health effects,” in-car PM 2.5 levels were on average double those found at monitoring sites (Adams et al. 2001). This is likely due to the contribution of vehicles and road dust to particulate matter concentrations, and makes it more plausible that drivers are regularly exposed to clinically significant PM 2.5 levels even when the ambient AQI would suggest otherwise. In fact, studies have shown significant elevations in several biomarkers of acute pulmonary and systemic inflammation following a realistic morning commute. These results indicate that a 2-hour on-roadway exposure to traffic air pollution can lead to sub-clinical changes consistent with airway inflammation and oxidative stress, though these effects may result from other on road pollutants besides PM 2.5 (Greenwald et al. 2012).

I.II DATA

PM 2.5

Because it is prohibitively expensive and time-consuming to measure the individual PM 2.5 exposure of enough drivers to perform a sufficiently-powered statistical analysis, PM 2.5 concentrations measured at outdoor ambient monitoring stations were used as a proxy for driver exposure. This proxy is plausibly suitable for this study because both on-road and in-vehicle pollutant concentrations were shown to be strongly correlated with concentrations from background PM 2.5 monitors (Li et al. 2013, Adams et al. 2001). Most likely, the correlation between on-road and ambient levels of PM 2.5 is weakest on the freeway and strongest on rural roads, because traffic on busier roads provides a greater localized contribution to particulate matter levels. This is supported by the finding of (Li et al. 2013) that on-road PM 2.5 concentrations were higher on freeways than other types of roads in the same area. For future work, the use of existing models from the literature of in-cabin exposure based on ambient levels, road conditions, and meteorological factors would significantly improve this analysis.

PM 2.5 monitor data are available on the AirData website maintained by the U.S. EPA (http://aqsdr1.epa.gov/aqsweb/aqstmp/airdata/download_files.html), and the monitoring itself is administered by the New Jersey Bureau of Air Monitoring to ensure compliance with National Ambient Air Quality Standards (NAAQS). Every one to six days, the stations sample continuously for a 24 hour period. These samples are analyzed to provide both a mean PM 2.5 concentration and AQI calculation for that day.

In the state of New Jersey, which is the focus of this analysis, monitors are placed in 17 different cities. The US EPA officially categorizes PM 2.5 to be a neighborhood scale pollutant, which means that measurements are consistent for a 1-10km radius, which is roughly the size of most cities. This facilitates the assumption that ambient levels are the same throughout the cities where the monitors are located, and thus that drivers within a city will all receive roughly the same exposure. Figure I.I and I.II plot PM 2.5 concentrations at two monitoring stations across time; the former shows monitors in nearby cities, while the latter shows monitors in far-away cities.

FIGURE I.I

Variation of daily PM 2.5 AQI in 2006 for Two Nearby Cities: Elizabeth and Newark.

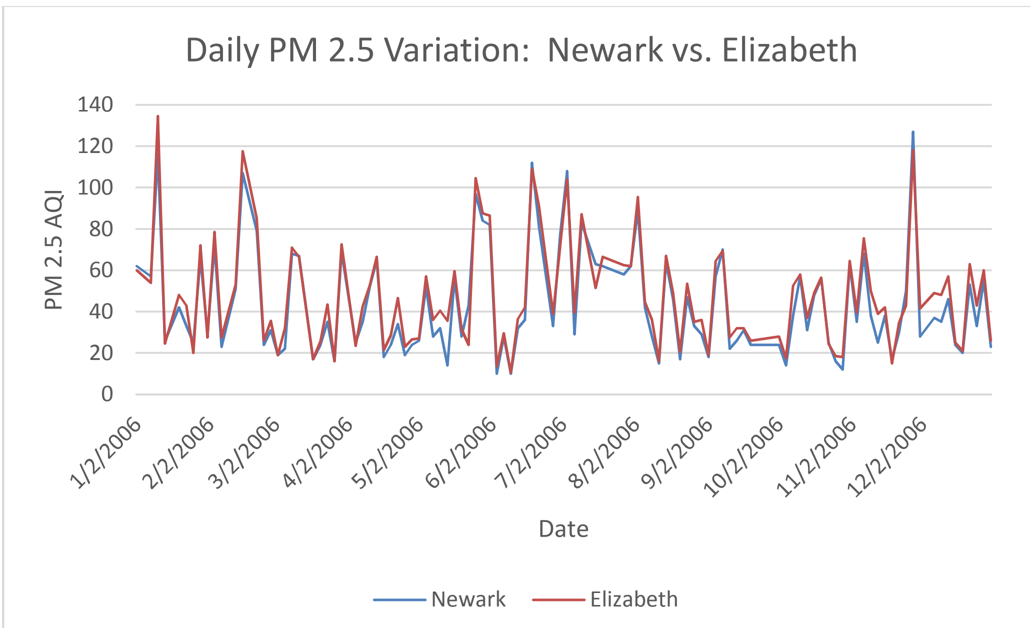

FIGURE I.II

Variation of daily PM 2.5 AQI in 2006 for Two Distant Cities: Camden and Newark.

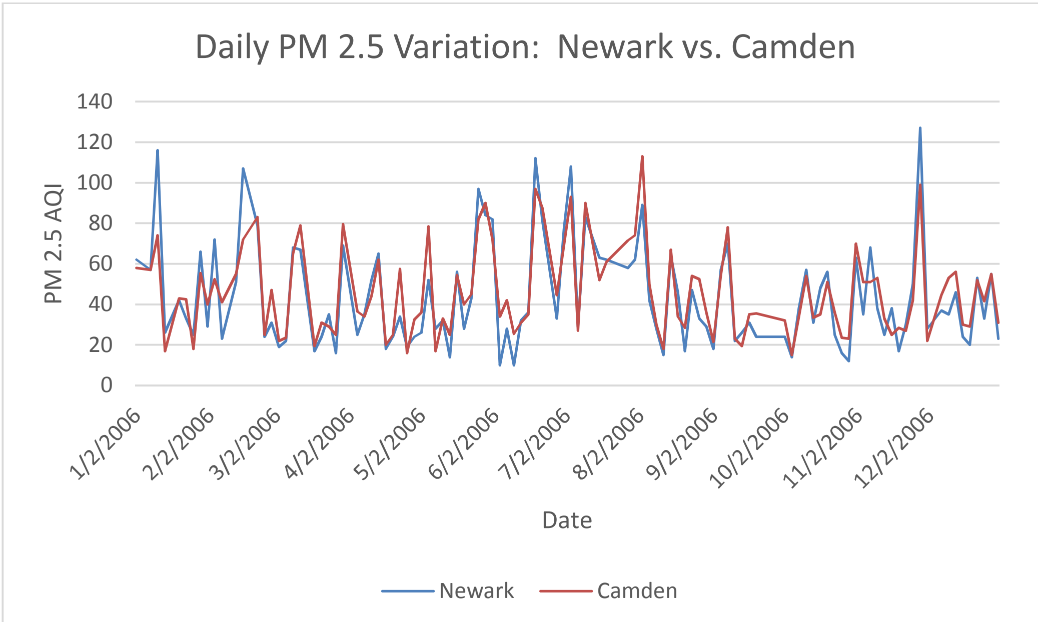

Notes: There were only 108 days out of the year where PM 2.5 was measured in all three cities. The remaining days were omitted.

The figures provide a sense of the regionality of the pollutant; concentrations at nearby cities are nearly identical across time, while concentrations at far-away cities differ substantially on many days. The monitoring sites are placed in urban areas and are intended to cover major population areas rather than the most polluted areas, thus the study sample should be representative of other urban areas nationwide.

Accident Data

The State of New Jersey Department of Transportation web site provides a database of accident reports filed by police in the state, found at (http://www.state.nj.us/transportation/refdata/accident/). New Jersey was chosen as the data source because it was the only state as of 2013 where accidents at the city-day level were readily available for the public to download. Its database is comprised of every accident report filed by the police, from which I was able to construct daily accident counts for each city each day. These counts provide a measure of daily accident frequency for a given city, with the assumption that the proportion of accidents reported is equal across days. One can be almost certain that not every accident is reported to the police, though all the serious ones likely are. There are also major issues with non-reporting in some cities, in which no accidents would be reported for several months despite a mean daily accident frequency comprising double digits. It is crucial that non-reporting be uncorrelated with PM 2.5 so that estimates are not biased. Although there is little reason to believe that pollution itself would reduce the rate of accident reporting, certain factors which are correlated with PM 2.5 concentrations such as rain and holidays might affect reporting, necessitating controls for them.

The likely omission of the most minor accidents also changes the way that estimation results should be interpreted. If the effect of pollution is the same for minor accidents as major ones, then it is reasonable to make inference about the effect of PM 2.5 on the accident rate as a whole and calculate damages accordingly based on aggregate figures. If one is to conservatively assume, however, that there is no effect on the very minor accidents for which there is no reporting, then one should only base damages on the expected increase in accidents predicted by model estimates.

Lastly, injuries and fatalities are also reported in the police data which allow one to observe the frequency of accident-related casualties as well. This can serve to provide an alternate measure of the effects of PM 2.5, and allows one to determine whether injurious accidents, which entail much greater economic costs, are disproportionately affected relative to reported accidents with no injuries.

Weather Data

To control for meteorological factors in the analysis, historic weather data was obtained from the Weather Underground website (http://www.wunderground.com/history). This resource archives weather data from airports and weather stations throughout the United States, providing historic daily weather for most major cities. Weather data were available for most cities with PM 2.5 monitors; for many cities, however, weather data needed to be obtained from adjacent or nearby cities. In these instances the accuracy of the weather data will be reduced, but there is not an obvious reason to believe that the daily measurement error in weather will be correlated with daily PM 2.5 levels. There are also many instances of data errors in which numerical entries are coded either as “N/A” or “T”. In these cases observations were dropped; approximately 5-10% of observations were dropped in this fashion. Once again, there is no reason to believe that the likelihood of these omissions are correlated with daily PM 2.5 variation. Lastly, a small number of temperature values are erroneously listed as zero, causing observations for that city-day to be dropped.

Traffic Count Data

Daily traffic count data were obtained from 41 key road segments in New Jersey to serve as a control for daily fluctuations in vehicle traffic; this control is important due to the high correlation of traffic with both accident frequency and on-road PM 2.5. Daily traffic counts, measured by induction loops in the roadways, were intended to serve as a proxy for traffic volume at nearby cities. These data were too raw, however, to use in the analysis. Many abnormal spikes and drops in vehicle counts were found at most monitoring sites due to the intermittent operation of detectors in one or more lanes. A separate statistical model would be needed to identify sensor outages and impute their counts. Given that daily variation from even complete sets of counts from fully operational detectors were uncorrelated with daily accident totals in nearby cities, the effort required to build such a model did not seem worthwhile.

Complete Dataset

Cities and days were matched across the three usable data sources, providing “city-day” observations containing PM 2.5 concentrations, accident totals, and weather data. Data were processed for 2006 – 2010, though future work will expand the scope to 2001-2013.

I.III EMPIRICAL STRATEGY

Because one is unable to observe individual drivers each day to collect data on their PM 2.5 exposure and whether or not they were involved in a motor vehicle crash, analysis was done at the city level. Variation in daily AQI level from PM 2.5 measured in a given city was exploited to identify a possible causal effect on the daily car accident rate for that city. Increases in accidents on days in which ambient PM 2.5 concentrations were higher, and vice versa, would be indicative of a positive correlation between the two. Two key assumptions required to facilitate this empirical strategy are that variation in ambient PM 2.5 is strongly correlated with variation in driver exposure to the pollutant, and that the acute effects of PM 2.5 are experienced on same day as exposure – both of which are discussed earlier in the paper.

Although multiple cities were used in the analysis, city-specific fixed effects and interactions were added to regressions so that identifying variation is found across time rather than across cities. The use of multiple cities served only to increase observations and gain statistical power in the regression described later.

Exogeneity of PM 2.5

To estimate the effect of PM 2.5 on car accidents, exogenous variation in PM 2.5 levels is needed. Due to the regional nature of the pollutant, the formation and accumulation of PM 2.5 is largely influenced by meteorology (Li et al. 2013, Stanek et al. 2011). In fact, much of the PM 2.5 comes from sources outside of New Jersey such as coal power plants in Pennsylvania; as a result much of the variation in New Jersey PM 2.5 concentrations is driven by wind (NJDEP 2009). The most important meteorological factors determining PM 2.5 concentrations are temperature, wind, relative humidity (Li et al. 2013), precipitation (Tai et al. 2009), and cloud cover (Wu 2014). Many of these same factors such as precipitation, wind, and cloud cover also play a role in accident frequency (U.S. DOT FHA 2014), and therefore must be conditioned-on for meteorology-driven PM 2.5 variation to be plausibly exogenous.

Important stationary sources of PM 2.5 include residential fuel combustion, electricity generation industrial processes, agriculture, and waste disposal (NJDEP 2009). Although these factors are not directly linked to car accidents, correlation might exist between them as a result of monthly and day-of-week patterns inherent to human economic activity; these patterns must therefore be controlled for as well. Stationary sources are more prevalent in some cities than others, and because accident rates also vary among cities due to variations in traffic management and infrastructure (Greibe 2003); these correlations provide an additional reason why city fixed effects are essential to the regression.

Lastly, automobile emissions along with dust from roads are also important sources of PM 2.5. However, temporal variation in traffic volumes explains a very small fraction of variation in even on-road concentrations (Li et al. 2013), thus it is very unlikely that any relationship found between PM 2.5 levels and accidents would be driven by reverse causality in which congestion from car accidents increases ambient PM 2.5 concentration reported by monitors. Nonetheless, it is still important to control for daily variation in vehicle counts, even if only a weak predictor of PM 2.5, given how important of a predictor they are for accident rates. Although other covariates such as temperature, precipitation, and daily and monthly patterns should control for a sizeable portion of traffic variation, the lack of a reliable measure of daily traffic variation is a serious limitation.

Other On-road Air Pollutants

One limitation of the empirical strategy is that on-road PM 2.5 levels are potentially correlated with other on-road pollutants that accumulate under the same traffic and atmospheric conditions. One category of these pollutants are ultrafine particles (UFP), which are a subset of PM 2.5 restricted to particles less than .1 micrometers. These are even more likely to adversely impact cognition than the larger particles which make up the vast majority of PM 2.5 mass. UFP and other major on-road pollutants are much less spatially homogenous than PM 2.5, making them prohibitively difficult to measure on a large scale and control for in the analysis. PM 2.5 might therefore serve as a proxy for these pollutants, and any potential effects they have on accident frequency might be falsely attributed to PM 2.5 in regressions.

Impacts through Visibility vs. Cognition Channels

Effects of PM 2.5 through two possible channels might be identified using this estimation strategy. One is the theorized cognitive effect, while another possibility is decreased visibility. The U.S. EPA cites PM 2.5 as a major cause of reduced visibility in the United States (http://www.epa.gov/pm/basic.html), which is also identified as a contributing factor to accidents (U.S.

DOT FHA 2014). The visibility channel is of less interest, because the effects of reduced visibility on accident frequency are obvious to drivers. Furthermore, it will be challenging to identify the contribution of PM 2.5 to visibility, since PM 2.5 accumulation and other visual obscurations such as fog could be affected by the same meteorological factors such as temperature and humidity (Li et al. 2013, SUNY Suffolk). Without controlling for visibility, these visual obscurations will bias estimates of PM 2.5's effect by causing additional accidents during the same atmospheric conditions that increase PM 2.5.

By controlling for visibility (which can be measured directly) in the regression, the point estimate for PM 2.5 will reflect its effect through non-visibility channels, (presumably cognition), so long as there are no uncontrolled factors influencing visibility that are correlated with the frequency of accidents.

Model Specification

Car accidents in a city-day can be thought of as a "Poisson process," i.e. one in which a few cases are generated per period of time (Nussbaum et al. 2007). Although the observed city-day accident counts in this study resemble a Poisson distribution, their variance is substantially higher than the mean, indicating over-dispersion. Therefore the negative binomial distribution was used for modeling these data because it contains an extra parameter to capture variance. This specification was also favored by Greibe (2003) and Eisenberg (2004) to model traffic collisions. Annual population was used as an exposure variable, which modifies the likelihood function so that the dependent variable is an accident rate rather than the total number accidents. This simultaneously served as a control for population growth, as well as prevented larger cities with higher traffic volumes (and consequently higher accident totals) from dominating the estimation of the PM 2.5 parameter.

The following model of accident frequency was estimated to determine the impact of PM 2.5:

$$Y_{it} = G_{it} * e^{\beta * PM_{it} + \gamma * X'_{it} + w_i + R_t + \varepsilon_{it} + \vartheta_{it}}$$

Variable names are defined below:

Y_{it} - Number of accidents in city i on day t

G_{it} - Population in city i on day t (used as the exposure variable to make the dependent variable a rate)

PM_{it} - PM 2.5 concentration in city i on day. In some regressions, the air quality index (AQI) of PM 2.5 is used instead. The variable is also stratified by AQI warning levels defined by the EPA.

X'_{it} - Matrix of city-day controls

- Month of year dummies for coastal cities
- Month of year dummies for cities near NYC
- Mean wind speed
- Quadratic function of temperature
- Functions of mean and minimum visibility
- Cloud cover.

w_i – City fixed effects

R_t – City invariant temporal effects.

- Month of year
- Day of week
- Linear trend

ε_{it} – Error term in the regression.

ϑ_{it} – Parameter estimated in the negative binomial model to correct for over-dispersion.

Dataset Restrictions

City-days with missing or erroneous weather data were dropped, although future work might impute these weather values to keep the observations.

City-months with unusual stretches of zero-accident days, which are discussed previously in the data section, were also dropped.

Cities with less than 40,000 people were excluded from the analysis due to the disproportionate number of zero-accident days in these cities. The accident totals fit a negative binomial distribution much better after doing so. A potential downside to dropping smaller cities is that these cities will have fewer sources of local PM pollution, which may not be perfectly correlated with ambient levels. As a result, data from the monitors in smaller cities are a better proxy for driver exposure within the city – making estimation more precise for those cities. Therefore some alternate specifications include cities with as little as 20,000 people, and use a zero-inflated negative binomial to control for the greater share of zero-accident days.

National and state holidays and holiday weekends were also omitted because factors such as drunk driving, unusual traffic volumes, and long-distance trips might affect the accident rate. At the same time these holidays might also be correlated with PM 2.5 concentrations due to less people working on the holidays (producing less industrial PM) and different home-energy usage patterns.

Lastly, observations were restricted to days with no precipitation. An extensive literature review, summarized in Appendix A.I, outlines the substantial difficulties in properly controlling for precipitation. Precipitation is strongly correlated with both accidents and PM 2.5 concentrations, and the relationship with accidents is highly complex and non-linear – especially in the absence of traffic count controls. Thus, observation-days with precipitation were dropped to prevent omitted variable bias from precipitation effects. If the true effect of PM 2.5 on car accidents is different on days with precipitation than without it, then the findings from this analysis will either overstate or understate the effect.

Instrumental Variables approach

To address the empirical challenges presented by the lack of controls for daily vehicle traffic and other local on-road pollutants, an alternate specification is used in which PM 2.5 is instrumented for. The instrument exploits variation in the direction of wind originating from Pennsylvania locations with coal-fired power plants. Emissions from coal-fired plants in Pennsylvania are a significant contributor of PM 2.5 in New Jersey (Bhattacharjee et al. 1999), and one would expect the direction of the wind to affect the extent to which these emissions affect New Jersey air quality.

The instrument is a function of the prior-day direction of the wind in two Pennsylvania locations near the New Jersey border with coal power plants (Northampton County and Philadelphia County). The more the prior-day wind-direction points towards the center of New Jersey, the higher the value that the instrument takes; this value does not vary from city to city, only from day to day. A visual inspection determined that wind from Northampton County blew most directly into the populous areas of New

Jersey from an angle of 250 degrees, and from Philadelphia County from an angle of 220 degrees. These vectors are illustrated on the map shown in Figure I.III.

FIGURE I.III

Vectors indicating the most direct wind paths from coal power plants in Philadelphia to populous New Jersey cities.

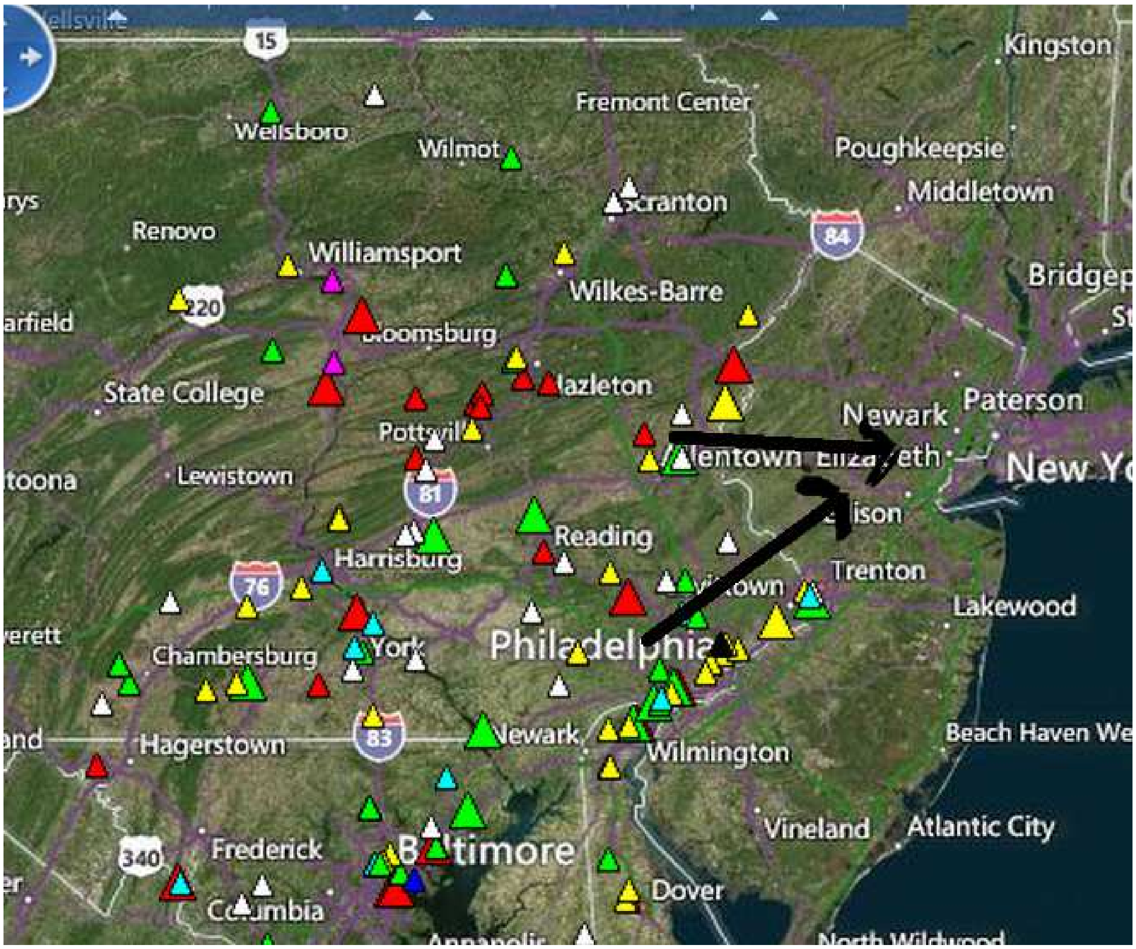

Notes: Coal-fired power plants are represented by red triangles. There is a high capacity coal power plant located below the "h" in "Philadelphia" whose icon is hidden behind a green triangle.

The absolute value of the difference between the actual wind direction and the "ideal" wind direction was computed each day for both counties, and then subtracted from 90. If the result was less than 0 for either city, it was recoded as a zero. The instrument was then calculated as the sum of the direction-score for both cities, and therefore can take values from 0 to 180.

The exogeneity of this instrument is contingent upon prior-day wind-direction only affecting accidents in New Jersey through the deposition of air pollution from Pennsylvania coal plants. While it is true that day-before wind direction won't directly cause accidents, it is correlated with some other factors that are also correlated with car accidents, necessitating the use of controls. These controls consisted of month of the year dummy variables, a cubic function of temperature, and cloud cover.

Conditioning on these variables, the instrument is a plausibly exogenous source of variation in air pollution concentrations in New Jersey. It addresses the prior empirical concerns, because unlike PM 2.5 levels, the instrument is uncorrelated with both traffic volumes and concentrations of local on-road pollutants.

One caveat to this approach is that other pollutants besides PM 2.5, such as sulfur dioxide, are likely carried from coal burning power plants by wind. Without including them in the regression, one cannot be certain that estimated effects of PM 2.5 are not due to the other pollutants. Unlike UFP, these wind-blown pollutants are measurable on a regional scale and widely available. Future work will include these pollutants in the analysis.

The same dataset restrictions from the previous methodology were also used for the IV regression. In addition, city-days with wind speed of greater than 15 mph were dropped to strengthen the instrument, because PM 2.5 concentrations are always low when strong winds are present.

I.IV RESULTS

This section contains regression results for the various specifications described in the methodology section. Table I.I.A and I.I.B define each of the variables listed in the results tables.

TABLE I.I.A

Variable Definitions for the Main Regression, Part 1

Variable	**Definition**	**Description**
Accidents	Number of accidents in city i on day j	Continuous
Injuries	Number of accident injuries in city i on day j	Continuous
AQI	The AQI from PM 2.5 in city i on day j	Continuous
lowAQI	AQI in city i on day j for AQI levels < 50	Continuous = AQI if AQI < 50 = 0 if AQI >= 50
highAQI	Indicator for whether AQI exceeds 50	Binary = 1 if AQI >=50 = 0 if AQI < 50
aqi50	AQI, for AQI levels >= 50	Continuous = AQI – 50 if AQI >= 50 = 0 if AQI < 50
aqi100	AQI, for AQI levels >= 100	Continuous = AQI – 100 if AQI >= 100 = 0 if AQI < 100
City2	City indicator	Factor variable
Month	Month-of-year indicator	Factor Variable
Coast#Month	Interaction of coastal city indicator and month-of-year indicator	Factor Interaction Variable
NYC#Month	Interaction of city-near-NYC indicator and month-of-year indicator	Factor Interaction Variable
Day2	Day-of-week indicator	Factor Variable

TABLE I.I.B

Variable Definitions for the Main Regression, Part 2

trend	Linear temporal trend term	Continuous
meanwind	Average wind speed in city i on day j	Continuous
meantemp	Average temperature in city i on day j	Continuous
temp2	Average temperature squared	Continuous
temp3	Average temperature cubed	Continuous
meanvisibility	Average visibility in city i on day j	Continuous
minvisibilility	Minimum visibility on city i on day j	Continuous
meanvisone	Indicator for whether mean visibility = 1.	Dummy
minvisone	Indicator for whether minimum visibility = 1.	Dummy
cloudcover	Cloud cover in city i on day j	Categorical variable, 1-8
precipitation	Total precipitation in city i on day j	Continuous
meanhumid	Average humidity in city i on day j	Continuous
sumw	Instrument for AQI, comprised of wind direction on day j.	Continuous

Main Regression

The main regression described in the methodology was unable to precisely identify an effect of PM 2.5 on car accidents. The results for the covariates are shown in Table I.II, and the results for the full regression are shown in Appendix A.II.

TABLE I.II

Results from the Main Regression

Negative binomial regression			Number of obs = 4727
			LR chi2(60) = 1704
Dispersion = mean			Prob > chi2 = 0
Log likelihood = -13411			Pseudo R2 = .0598
Accidents	Coef.	Std. Err.	P-value
Lowaqi	0.0004	0.0006	0.5060
Highaqi	0.0096	0.0263	0.7150
aqi50	-0.0005	0.0008	0.5180
aqi100	0.0002	0.0038	0.9490
ln(population)	1.0000	(exposure)	
/lnalpha	-3.3697	0.0721	
Alpha	0.0344	0.0025	

Replacing visibility variables with a variable for humidity (to make sure fog is still controlled for) had very little effect on the magnitude or precision of the estimates for the PM 2.5 variables.

Including smaller cities in the regression also had little impact on AQI-related estimates. The population threshold for cities was lowered from 40,000 to 20,000, which added Atlantic City, Camden, Pennsauken, Rahway, and Galloway to the dataset. Another regression was then performed using a zero-inflated negative binomial regression to adjust for the increased proportion of zero-accident days in smaller cities. These results, provided in Table I.III, still show no effect from ambient PM 2.5 on the frequency of vehicle collisons.

TABLE I.III

AQI-related Regression Results with 20,000 Population Cutoff and Zero-Inflation

Zero-inflated negative binomial regression			Number of obs = 6455
			Nonzero obs = 6328
			Zero obs = 127
Inflation model = logit			LR chi2 (65) = 2294
Log likelihood -16964			Prob > chi2 = 0
Accidents	Coef.	Std. Err.	P-value
Lowaqi	0.0003	0.0006	0.6460
Highaqi	-0.0011	0.0240	0.9630
aqi50	-0.0000	0.0007	0.9760
aqi100	-0.0004	0.0034	0.9010

Notes: Only estimates related to the effects of PM 2.5 are shown.

Ruling out Data Issues

Because other predictors for accidents were statistically significant and had the expected signs, there is little that suggests accident-data issues were the reason why a significant effect for PM 2.5 could not be found.

Furthermore, regressions of PM 2.5 concentrations on meteorological variables, shown in Table I.IV, yielded significant estimates that were predicted by theory. Therefore, it is unlikely that significant issues exist with the ambient PM 2.5 data either.

TABLE I.IV

Regression of PM 2.5 concentrations on meteorological variables

Aqi	Coef.	Std. Err.	P-value
Precipitation	-0.0186	0.0100	0.0630
Meanwind	-1.3344	0.0514	0.0000
Meantemp	0.2459	0.0105	0.0000
Meanhumid	0.0611	0.0844	0.0490
Cloudcover	-0.1659	1.1381	0.0000

Potential Reasons for Lack of Positive and Significant Findings

One possible reason for the lack of precise estimates for the PM 2.5 coefficient could be that the effect is very small. Given the substantial noise in the occurrence and reporting of accidents, it may be too difficult to estimate a very small effect without more data. Another possibility is that driver exposure to cognition-affecting pollutants is not correlated strongly enough with ambient levels of PM 2.5 detected by monitors. One possible way to address this in the future is to use a model to predict on-road pollutant exposure using ambient levels and meteorology. Alternately, ambient PM 2.5 levels themselves might negatively affect the decision to drive. To test this, more reliable traffic data is needed. Lastly, variation in ambient PM 2.5 might still not be exogenous despite controls, and might possibly even be anti-correlated with other causes of accidents. This issue is best addressed by using an instrument for PM 2.5. The following section provides results from the instrumental variable regression described in the empirical strategy.

Results from Instrumental Variable Approach

Table I.V on the following page shows the results from regressing PM 2.5 level on the instrument and controls, in order to demonstrate the relevance of the instrument. The analysis has only been completed for data from the year 2006, but future work will include later years.

TABLE I.V

First-stage regression of PM 2.5 AQI on its instrument (sumw) and other controls

Source	SS	df	MS	
Model	157267	16	9829	Number of obs = 544
Residual	166641	527	316	F(16, 527) = 31
				Prob > F = 0
				R-squared = .4855
				Adj R-squared = .4699
Total	323908	543	596	Root MSE = 17.7820

Aqi	Coef.	Std. Err	P -value
Sumw	0.1219	0.0155	0.0000
Month			
2	20.8208	4.8957	0.0000
3	0.4975	5.0282	0.9210
4	-28.7305	5.5168	0.0000
5	-27.6487	5.7366	0.0000
6	-40.1609	6.7438	0.0000
7	-49.1710	7.1755	0.0000
8	-60.9498	7.1658	0.0000
9	-42.2220	6.4215	0.0000
10	-24.7842	5.2444	0.0000
11	4.1512	5.1481	0.4200
12	-1.6441	4.7966	0.7320
Cloudcover	0.7090	0.3712	0.0570
Meantemp	-1.2457	1.7213	0.4700
temp2	0.0359	0.0329	0.2750
temp3	-0.0001	0.0002	0.6080
_cons	24.6604	28.6577	0.3900

The first-stage results show that the wind-direction instrument, denoted in the regression table as *sumw*, is a very strong predictor of ambient PM 2.5 in New Jersey cities. When the instrument attains its maximum value, the model predicts a 21.6 point increase in the AQI from PM 2.5, which is roughly one standard deviation. The F-statistic for the null hypothesis that the coefficients on the instruments are jointly zero in the first stage regression is 61.9 (7.87^2), indicating that the instrument is sufficiently strong.

The results for the instrumental variable regression are shown in Table I.VI. An instrumental variable Poisson regression was used, as opposed to negative binomial, due to the lack of a negative binomial IV estimate routine in Stata. Future work should determine how this changes affects the estimated standard errors.

TABLE I.VI

Regression results from instrumental variable specification

			Number of obs =
Number of parameters = 18		544	
Number of moments = 17			
Initial weight matrix: Unadjusted			
GMM weight matrix: Robust			
Accidents	Coef.	Std. Err	P - value
Aqi	0.0054	0.0025	0.0290
Month			
2	-0.0430	0.2025	0.8320
3	0.4933	0.1961	0.0120
4	0.5791	0.2222	0.0090
5	0.6301	0.2402	0.0090
6	0.8096	0.2577	0.0020
7	0.7596	0.2780	0.0060
8	0.8146	0.2985	0.0060
9	0.6889	0.2498	0.0060
10	0.5442	0.2192	0.0130
11	0.4523	0.2054	0.0280
12	0.3839	0.1999	0.0550
Cloudcover	-0.0131	0.0079	0.0940
Meantemp	-0.0970	0.0361	0.0070
temp2	0.0016	0.0007	0.0140
temp3	0.0000	0.0000	0.0120
_cons	-7.9559	0.6226	0.0000
ln(popula~n)	1 (exposure)		
Instrumented: aqi			
Instruments: 2.month 3.month 4.month 5.month 6.month 7.month 8.month 9.month 10.month 11.month 12.month cloudcover meantemp temp2 temp3 sumw			

The estimate for the effect of PM 2.5 AQI on car accident rates is significant at the 5% level. The coefficient on the AQI variable is .0054, which is interpreted by taking the exponent and subtracting 1 to give the percent increase in the accident rate per unit increase of PM 2.5 AQI. Exp(.0054) is approximately 1.0054, implying that a one unit increase in PM 2.5 AQI increases the rate of car accidents by .54 percent. This also implies that a one-standard deviation in PM 2.5 AQI increases the car accident rate by 13.2 percent. Even though this estimate likely includes effects from other pollutants carried from coal-fueled power plants emissions, it still seems implausibly high. Even if valid, one must be cautious generalizing the results due to the previously described data restrictions. For example, until further analysis is done, assumptions must be made regarding whether the results will hold for smaller cities, very windy days, and days with precipitation. There is no indication, however, that any possible effects of PM 2.5 on accidents would be diminished in such conditions.

Further Checks

Table I.VII shows that when visibility is controlled for in the IV regression, the estimates for the effect of PM 2.5 on accidents do not change. This suggests that the estimated effects are occurring through the cognition channel, rather than the visibility channel.

TABLE I.VII

Estimate of impact of AQI on accidents when visibility is controlled for in the IV specification

Accidents	Coef.	Std. Err	P - value
Aqi	0.0060	0.0034	0.0800

Notes: Only estimates related to the effects of PM 2.5 on accidents are shown.

When the daily total of car accident injuries is used as the dependent variable instead of accidents, there is an even stronger effect estimated for PM 2.5. Table I.VIII shows the results of this alternate model. The result suggests that PM 2.5 concentrations tend to affect serious accidents more strongly. This would make sense if the causal effect was due to slowed reaction time, which would be more detrimental at higher speeds.

TABLE I.VIII

IV specification results using injuries as the dependent variable instead of accidents

Injuries	Coef.	Std. Err	P - Value
Aqi	0.0156	0.0072	0.0310

Notes: Only estimates related to the effects of PM 2.5 on car-accident injuries are shown.

Estimated Benefits of Pollution Abatement through Avoided Car Accidents

If the regression results are taken at face value, one can directly calculate the number of accidents that can be prevented through pollution abatement. The model predicts the percent decrease in a city's accident rate that results from from a given decrease in the city's ambient PM 2.5 concentration, and this decrease can be multiplied by the city's population. Once the decrease in accidents is computed, it can be multiplied by the average cost of an accident to obtain the total benefit each day of PM 2.5 abatement. That approach is not appropriate for this study, however, because it cannot yet be ascertained whether the estimated effect of PM 2.5 using the IV regression is in fact due to PM 2.5 or another pollutant emitted from coal-fueled power plants. Future work will include these other pollutants in the analysis so that the specific contributions of each pollutant can be assessed. This will then facilitate the calculation of accident avoidance benefits from abatement of each pollutant.

I.V CONCLUSION

The main regression in the analysis did not provide evidence that PM 2.5 pollution has causal effect on the frequency of car accidents, while the instrumental variable regression provided evidence that PM 2.5 affects the frequency of both car accidents and car accident injuries. There are multiple possible reasons for this discrepancy. One possibility is that the endogeneity of PM 2.5 variation in the main regression biases estimates downwards and prevents the finding of a significant effect. Another possibility is that the estimate of PM 2.5's effect in the instrumental variables regression is also picking up the effect of other pollutants emitted from coal-fired power plants on car accidents. In either case, the IV regression shows evidence that at least some form of regional air pollution is increasing the frequency of car accidents through its effect on driver cognition. Further work will be needed to explore which pollutants are causing the effects. In spite of this limitation, the analysis does provide evidence that the inhalation of pollutants from coal power plant emissions increases the likelihood of being involved in a vehicle collision. The damages resulting from the increased propensity to be in a crash might justify further regulation of coal power plant emissions.

Due to the unexpectedly large magnitude of the estimated effects in the IV regression, further robustness checks should also be performed. For example, the IV regression is currently only performed for one year of data, so the inclusion other years would be beneficial. Furthermore, modifying the IV analysis so that the instrument attains its maximum value at a set of arbitrary wind directions would be a useful placebo test. Lastly, repeating the analysis for a different state would be instructive; a positive finding in two different states would be much stronger evidence that there is a causal effect of air pollution on accidents.

CHAPTER II

LEARNING EQUILIBRIA IN MULTI-STATE TRAFFIC NETWORKS

II.I INTRODUCTION

Planners assume that users of congestible networks such as roads and telecom infrastructure will eventually reach equilibrium through learning. Empirical studies confirm that users at least tend toward equilibrium behavior over time. This assumption is critical to the analysis and management of these networks, given that performance and reliability are generally improved once equilibrium is reached. In these cases, social welfare depends on the speed and extent to which equilibria are reached. For example, roads are often priced to create a new user equilibrium that corresponds more closely to the social optimum. Lengthy delays in reaching this new equilibrium would undermine the policy by postponing the intended welfare benefits.

Equilibrium learning is constrained by human cognitive abilities and is also influenced by characteristics of the network. Factors that affect the ability of users to perceive and recall differences in outcomes from various choices will significantly affect learning speed (Arthur 1991, Koster et al. 2015). Empirical studies have shown that in relatively simple networks, convergence to near-equilibrium outcomes is rapid (Selten et al. 2007, Iida et al. 1992). This work explores how human learning is affected in a more complex learning environment.

Route-choice equilibrium learning was studied in a traffic network made more complex by the existence of several discrete states, each corresponding to distinct system equilibria. Transitions between discrete states are analogous to many real-world phenomena, such as single or multi-lane blockages on freeways or servers taken online and offline in telecom networks. The presence of multiple states complicates learning because experiences gained in one state of the network are not necessarily relevant to explicit or implicit beliefs about the network in other states. If subjects in multi-state networks do not exercise discretion to avoid incorporating experiences from irrelevant states, they will form inaccurate state-specific knowledge and beliefs that will delay equilibrium convergence. While automated agents, such as networked applications making routing decisions, use look-up tables to avoid these errors when learning in the presence of multiple states (Watkins and Dayan 1989), cognitive limitations might preclude human-decision makers from doing the same.

This work studies whether human learning is impaired when networks are made more complex by the presence of multiple states. It answers whether users possess the cognitive capacity to perform "state-specific" learning in networks, where users only make decisions based on past experiences from relevant states. To address this question, a binary route-choice traffic experiment was conducted to determine whether humans commit cognitive errors learning equilibria in multi-state networks. In this experiment, a group of subjects repeatedly traversed a congestible network and chose one of two parallel routes. The network transitioned between several distinct states from trial to trial. Each state corresponded to a different incident scenario consisting of discrete capacity loss of varying degrees on one of the routes.

Route-choice equilibrium learning by subjects was analyzed and compared between two types of experimental treatments: one with a simple two-state network, and one with a more complex five-state network. Simple reinforcement-learning models were fit to the experimental data to determine how subjects incorporate travel time outcomes from past trials into their route choice decisions. The data revealed that subjects in the five-state network committed widespread learning errors, while those in the two-state did not. Subjects in the two-state network near-exclusively incorporated information from relevant states into learning, resembling the previously described error-free learning algorithms of automated applications sending data over telecom networks. Subjects in the five-state network, however, often based route choice decisions on travel times experienced in different states, which are irrelevant. This suggests that the cognitive demands imposed by additional states were detrimental to learning.

Next, simulations were used to determine the effect that learning errors in the five-state network had on long-term convergence to equilibria. A simulation based on subject behavior in the five-state network was compared to a simulation based on subject behavior in the two-state network (both simulations used a five-state network). The results showed that impaired subject learning in the five-state network significantly delayed convergence to equilibria, and sometimes precluded equilibrium from being reached at all.

This finding adds to the behavioral economics literature by identifying a novel interaction between limited human cognition and the learning environment. In a network with multiple, discrete states, memory limitations led subjects to incorporate information from irrelevant states into route choice decisions.

This work as also contributes to the travel behavior/transportation literature by being the first to show how route-choice equilibrium learning among multiple human users is delayed/precluded by added complexity in traffic networks. This finding is applicable to real-world traffic scenarios. Numerous individuals face repeated route-choice decisions in their freeway commutes, with non-recurring single or multi-lane closures and blockages that create discrete state changes. It is likely that equilibrium convergence is significant delayed in these instances due to learning difficulty. In such scenarios, it may be beneficial for planners to implement measures that address the cognitive demands of learning over multiple states.

It is also the first study of equilibrium learning in traffic networks that employs a video-game-like platform that allows subjects to control virtual vehicles and simultaneously traverse a 2-D roadway in real-time. This significantly increases the realism of the platform compared to previous studies where subjects make hypothetical route choices in the absence of a simulated environment.

II.II LITERATURE REVIEW

There is long tradition in the transportation literature of behavioral experiments to study route and departure time choices, especially in response to information. A thorough a review of these works can be found in Mahmassani (2009). The scope of this review is limited to works that focus on route-choice. Furthermore, only the aspects of route-choice learning most relevant to the work of this paper are discussed in section. These consist of studies that identify cognitive factors affecting route-choice learning, studies that test convergence to route-choice user equilibria, and studies that model route-choice learning.

Cognitive factors affecting learning

Many prior studies have identified human cognitive ability to be a limiting factor in route-choice learning. For example, Vreeswijk et al. (2011) find that drivers have difficulty perceiving changes in traffic intensity, and Vreeswijk et al. (2014) find that drivers overestimate travel times on routes they don’t choose. Memory limitations have been explicitly shown to impact driver learning in the studies of Bogers et al. (2007), Ben-Elia & Shiftan (2009), and Pu et al. (2009), which show strong recency-bias in the use of information to make route-choice decisions. Arentze and Timmermans (2003) and Koster et al. (2015) simulate such memory limitations to show that they delay learning and reduce driver utility. Iida et al. (1992) experimentally confirm the adverse impacts of learning impairments on system performance, and find that performance improves when subjects are provided with a table of past travel times to assist memory.

Network complexity affecting learning

Ben-Elia and Shiftan (2009) conducted a route-choice experiment where subjects chose between alternatives whose travel times exhibited varying degrees of stochasticity. They show that increased stochasticity increases the time it takes for subjects to determine the faster route. This finding is independent of human cognition, because even a computer would learn less efficiently from a noisier signal. Bogers et al. (2007) show that when a network undergoes a state change, there is a transition period required for drivers to learn the new equilibrium. This too is independent of human cognition, because a learning period would be required for a machine as well. Therefore, while both studies show that network complexity can delay learning, there is no evidence that limited human cognition is a contributing factor.

This study adds to the literature by identifying a novel interaction between limited cognition and the learning environment. In a network with multiple, discrete states, memory limitations lead subjects to incorporate information from irrelevant states into route choice decisions. This is detrimental to equilibrium learning; Watkins and Dayan (1989) show analytically that it is instead optimal for experiences in specific networks states to only influence future decisions made in those same states (Arentze and Timmermans (2003) find a similar result using simulations). This study empirically confirms that using information from irrelevant states delays equilibrium learning, and unlike other previously identified cognitive limitations, has the potential to preclude equilibrium convergence entirely. This distinction is important because there is evidence from laboratory experiments that in simple networks, aggregate system performance can quickly approximate equilibrium[1] despite the possibility of delays caused by the cognitive limitations identified in previous studies. For example, subjects in the route-choice experiments of Selten et al. (2007) and Iida et al. (1992) readily oscillate[2] about the aggregate equilibrium route split. The studies, which consisted of binary route choice games in simple

[1] Evidence from real-world networks is less conclusive. Watling et al. (2012) found a rapid aggregate switch away from a route with reduced capacity, but an increase in usage of the facility that includes the affected route. Zhu et al. (2010) found it took about a month for a new equilibrium to be reached after a bridge collapsed.
[2] Oscillations are consistent with mixed strategy equilibrium theorized by Arthur (1991) and Monderer & Shapley (1996). In the route choice studies of Mahmassani (1990), however, only pure strategy equilibria are considered, and the author concludes that equilibrium is not reached due to persistent route switching.

congestible road networks to test whether drivers could coordinate on a payoff-maximizing equilibrium route split, provide evidence that convergence delays from cognitive limitations are not particularly significant in simple networks. In complex networks used in this study, however, impaired multi-state learning may result in a lack of equilibrium convergence altogether.

Testing equilibrium learning and convergence

Field studies of equilibrium learning in traffic networks are rare, because researchers lack the ability to intervene in real-world traffic networks and face challenges in accurately observing driver behavior. A study by Watling et al. (2012) exploits a natural experiment created by the planned reduction of capacity in part of a traffic network. Aggregate driver behavior was estimated through an algorithm based on the partial matching of license plates. The authors found that during the six days following capacity reduction on a route, aggregate usage shifted to other routes as predicted by equilibrium theory. However, usage of the facility that included that route paradoxically increased. It is unclear whether the second result would persist if controls for aggregate traffic volumes were included. Extreme case where driver learns new route. A different study by Zhu et al. (2010) exploited a natural experiment resulting from a bridge collapse to study driver behavior in response to the capacity shock. Aggregate traffic counts measured by loop detectors revealed that it took approximately a month for drivers to reach a new equilibrium traffic pattern. The bridge collapse represented an extreme case because each user had to find a completely new route to take; this likely contributed to the extended time required for equilibrium to be learned.

Driver behavior is more easily studied in laboratory experiments that give researchers full control over network conditions and perfect observation of users. The study of Selten et al. entailed a binary route choice game in a simple congestible road network to test whether drivers could coordinate on a travel-time-minimizing equilibrium route split. The author linked individual travel times to aggregate decision route-choice decisions, which precluded the existence of a dominant strategy (either route could be faster depending on the decisions of other agents). The experimental design required that agents iteratively adjust their behavior through "best-response" dynamics (Crawford 1995), which in theory results in equilibrium convergence (Monderer and Shapley 1996). Iida et al. (1992) conducted a similar experiment but had subjects guess the faster route, which then became their de facto choice. This reduced the emphasis on coordination, because it ruled out mixed strategies where subjects choose a route a fraction of the time but are indifferent between the two. Mahmassani and Stephan (1988) studied equilibrium convergence under a broader scope; allowing subjects to provide both route-choice and departure time decisions which were then supplied as inputs to a traffic simulation. Both route-choice and departure time decisions are critical components travel behavior analysis, however in the incident-management context of this study, route-choice decisions are made in real-time and independently of prior departure-time decisions. Therefore, this experimental framework of this study is restricted to route-choice decisions, similar to Selten et al.

The use of human-subject experiments allowed the authors to obtain better insights into actual route-choice behavior than a simulation. Furthermore, the experimental design made use of endogenous traffic (and thus payouts) rather than stochastic independent travel times, enhancing the realism of the route-choice decisions being made. This framework is adopted for this study to analyze how route choice learning is affected when network complexity is increased. This study utilized the same basic experimental design as Selten et al., with some major differences. The first difference is the introduction of multiple states in the network to test whether they adversely affect learning. More complex networks have also been used in the route choice experiments of Ben-Elia & Shiftan (2009) and Bogers et al. (2007), where stochastic travel times were drawn from distributions that periodically shift throughout the experiment. These works lack endogenous traffic however, and therefore only study single-user equilibria rather than interactive multi-user equilibria. The second way my experiment differs is in the utilization of a video-game-like platform that allows subjects to control virtual vehicles and simultaneously traverse a 2-D roadway in real-time. This significantly increases the realism of the platform compared to the aforementioned studies where subjects make hypothetical route choices in the absence of a simulated environment. This enhanced realism is largely incidental to this study (it was implemented to address other research questions), but should still increase the externality validity of the study to real-life route choice scenarios on freeways.

Analysis of route-choice learning

Reinforcement learning models have been used in many of studies of route-choice behavior to relate agents' choices their future decisions. In these models, choice outcomes influence an agent's overall route preferences rather than his/her beliefs about specific route attributes such as travel time (Chorus and Dellaert 2010). This makes these models useful for studies with human subjects, since the researcher will typically only observe an agent's choices and not his/her beliefs (an exception is the study of Iida et al., where subjects were asked to predict travel times on each route before choosing). They are especially practical in route-choice experiments with endogenous traffic, where travel times depend on aggregate decision-making rather than independent draws from an exogenous distribution. This is because it far simpler to model a comprehensive preference for each route, rather than a complex set of agent beliefs pertaining to the behavior of other drivers in addition to route characteristics.

Selten et al. and Pu et al. estimate reinforcement learning models from their route-choice data from human-subjects experiments, and Arentze and Timmermans use a reinforcement learning framework for simulation experiments. Selten et al. test a simple one parameter reinforcement learning model described in Arthur (1991) and shown by Erev and Roth (1998) to fit and predict human behavior in numerous experimental congestion games. Agents in this model have an initial propensity to choose each strategy based on implicit beliefs regarding payouts along with other unobserved random preferences. In the case of traffic networks, strategies correspond to routes, and payouts correspond to the inverse of travel time. The probability of an agent choosing a certain route is increasing in his/her propensity for that route relative to other routes. Propensities are updated each time a route is chosen, in a manner depending on the outcome of the choice. The propensity increase is inversely proportional to travel time, so that a very good payout (low travel time) makes an agent more likely to repeat their choice than a bad outcome. Choosing a route will always increase its propensity, which mimics the effect of inertia in choice behavior.

Selten et al. fit their data well with a modification of this model, and although they find that a small share of subjects likely behave according to a different set of heuristics, the reinforcement learning model captures the behavior of most subjects. Pu et al. also find that a reinforcement learning model is the best fit for the behavior of some, but not all, subjects (the authors do not specific the shares).

Another variation of the reinforcement learning model is introduced by Arentze and Timmerman (2003). In this model, stochasticity in route choice results from agents' desire to seek information about alternate routes, rather than unobserved preference parameters. This modeling framework is intuitively appealing and provides extra parameters to govern the frequency of route exploration. For certain parameter configurations, this model will be identical to the random-preference reinforcement learning models above. Because this feature of the model has been not empirically tested using experiments with human subjects or endogenous traffic, it is unclear whether the extra parameters are justified. The authors also introduced a state-based framework for reinforcement learning, where information gleaned during one state of the network will only be applied to other route choice decisions made during that same state. This framework is ideal for modelling state-specific learning, which Watkins and Dayan proved was optimal.

In each of these studies using reinforcement learning models, route-choice propensities are updated based on a linear function of prior payouts / travel times, though alternate functional forms can also be used. Selten et al. weight all past travel times equally, while the other two studies (Arentze and Timmermans, Pu et al.) parameterize the weighting of each past travel time. It is unclear from the latter two studies how much explanatory power is gained from parameterizing the relative weights of past travel times. Other classes of models described in the next section, which model beliefs rather than propensities, use Bayesian updating for instead. There is no compelling reason why Bayesian updating is advantageous for reinforcement learning models, however.

More theoretically explicit models outside the domain of reinforcement learning are used in the experimental study of Iida et al. and in simulation studies by Nakayama et al. (2001), Horowitz (1984), and Mahmassani and Chen (2003). In these dynamic models, subjects' specific beliefs about route attributes such as speed and reliability evolve through experience. Subjects choose a route based on either heuristics or utility maximization based on those attributes. Cantarella and Cascetta (1995) provide the theoretical underpinning of the sufficient and necessary conditions for drivers' route-choice behavior to user equilibria in these models. Models based on beliefs are less suitable for empirical work, because unlike route choice propensities, subjects' beliefs cannot be directly estimated from observed route choices. Furthermore, belief-based models are more deterministic than reinforcement learning models, which is inconsistent with homogenous mixed strategy equilibria theorized by Arthur and Monderer and Shapley. For example, the fully deterministic model of Nakayama precludes a mixed strategy equilibrium entirely, and in simulations was outperformed by models with no learning at all. The models used by Iida et al., Horowitz (1984), and Mahmassani and Chen introduce an element of stochasticity to his model through random errors in the perception of experienced travel times, consistent with the finding of Vreeswijk, Do, et al. (2011). The stochasticity of these models facilitates a mixed strategy equilibrium in theory, however for most values of parameters tested in the studies of Horowitz and also Mahmassani and Chen, the amount of stochasticity was not sufficient to produce enough route-switching for equilibrium to be learned. Instead, a pure-strategy steady-state was reached that did not coincide with the user equilibrium.

A study by Bogers et al. (2007) tested a hybrid model that incorporates both reinforced preferences and also explicit travel time beliefs into route-decisions. Travel time beliefs were updated through a separately estimated Markov model, since they can't be observed directly. The model fit human subject behavior well in a network with exogenous travel times, was not tested against behavior in a network with endogenous travel times.

Any of the models discussed in this section are suitable for estimating the relevant learning errors from this experiment's data. To assess whether subjects incorporate irrelevant information into route-choice decisions, one need only observe that travel times from a particular network-state influence route-choice probabilities in different states. Each of the aforementioned learning models can identify such a relationship. A reinforcement model with random unobserved preferences is advantageous, however, due the practicality, parsimony, and empirical validation of that framework. The model used by Erev and Roth exemplifies these attributes by fitting a wide array of learning environments using only one free parameter. I adopt this model to analyze my own experimental data, with the caveat that some of its theoretical predictions are contradicted by the empirical findings of other studies. For example, although subjects in the experiment of Selten et al. achieved aggregate route splits that approximated equilibrium, their behavior was not fully consistent with the homogenous mixed strategy equilibrium predicted by the reinforcement learning model. I modify this model to accommodate the state-based learning framework introduced by Arentze and Timmermans in order to test whether learning is impaired when multiple network states are introduced.

II.III METHODOLOGY

Experiment design

A route-choice experiment was used to study driver learning behavior. Human subjects used a driving simulator to control vehicles on a virtual road network, where they chose between one of two congestible routes. A learning model could then be estimated from the generated route-choice data to determine what information subjects used to inform their decisions.

The experimental sessions analyzed for this study are a subset of those conducted for a larger project that studied driver responses to traffic information. The project's final report (Kong et al. 2017) contains a comprehensive description of the driving simulator used for the experiment, while a higher-level overview is provided below.

Each experimental trial consisted of 39 subjects simultaneously traversing the road network from a shared starting area to an end zone. The road network consisted a three-lane main route, and a two-lane alternate route with a more circuitous path. Subjects were incentivized to reach the end-zone as quickly as possible by a subtractive payment scheme that reduced earnings monotonically during their time of travel.

Congestion developed endogenously on the network through the interaction of the subject-controlled vehicles sharing each route's limited capacity. Throughput on the main path was often impacted by exogenous lane blockages of varying severity analogous to unforeseen traffic incidents. The incident severity levels were implemented as discrete reductions in route capacity, each corresponding to a unique network state with a distinct aggregate Nash equilibrium route-split. The network transitioned to a different state/incident-scenario after each trial.

Subjects had a single opportunity to switch to the alternate route upstream of the incident area, before there was visual evidence of congestion. Therefore, subject decisions were informed only by their prior knowledge (experience) of the network, the actions of other subjects in their limited field of view, and any information provided through variable messages which mimicked variable message sign (VMS) installations used on freeways. The messages informed subjects which incident-state the network was in but did not provide any other information.

The multi-user interactive implementation of the driving simulator was critical to the experiment's design. Each subject's state-specific route-choice outcome depended on the behavior of other users, which made state-specific learning non-trivial. Each state-specific equilibrium route-split required that some share of users chose each route; there was never a dominant pure route-choice strategy that minimized travel time. This feature was essential for ensuring that subjects must coordinate, even if only implicitly through best-response learning, to achieve equilibrium over time.

Two different treatments were used; one where the network transitioned between only two states, and another where the network had five states. This design exploits the fact that state-transitions impose a cognitive burden requiring subjects to learn separate route-choice equilibria. The two-state network comprised a simple learning environment with a lower cognitive burden, while the five-state network imposed a more significant cognitive burden. Thus, equilibrium learning impairment observed to a greater degree in the five-state treatment than the two-state treatment would support the hypothesis that state transitions impose a cognitive demand on subjects that induces learning errors.

In sessions with the two-state treatment, subjects encountered a pre-randomized sequence of incidents each trial that blocked either one lane or three lanes. In sessions with the five-state treatment, subjects encountered an incident sequence comprised of five discrete pre-randomized scenarios consisting of:

1. no blockage at all
2. one lane blocked
3. two lanes blocked
4. three lanes blocked with eventual intermittent one-lane capacity
5. three lanes blocked with eventual less-frequent intermittent one-lane capacity

Each session began with a presentation of instructions explaining the traffic network, driving mechanics, and the earnings scheme. Subjects then completed 3 unpaid practice trials that further familiarized subjects with these topics. Next, subjects completed 20 consecutive paid trials. Lastly, subject filled out a questionnaire that solicited demographics and past driving experience.

Human subjects were recruited through random selection from the population of UC Irvine undergraduate and graduate students from the Experimental Social Science Laboratory (ESSL) subject pool with IRB approval under HS #2011-8378.

Theoretical Model

Three models were compared to test the hypothesis that in networks with more states, subjects are more prone to erroneously incorporate irrelevant information into route-choice decisions. The first model represented "error-free" learning, where subjects only base route-choice decisions on travel times experienced during the same state. This reflects the finding from Watkins and Dayan (1989) that learning in a multi-state network ideally occurs independently for each state; travel times experienced during one state should not influence future decisions made in other states.

The second model reflected cognitive learning errors where users incorporate irrelevant information into route choice decisions. Rather than base decisions on experiences that occurred during the same state, users only update propensities based on the travel time that occurred most recently, regardless of state, because it was the easiest result to recall.

The third model was a hybrid of the first two, where subjects made route choice decisions based on both the most recent experience that occurred during the same state and the travel time experienced immediately prior.

The models were based on the one-parameter simple reinforcement learning model from Erev and Roth (1998). In this model, subjects playing games have a unique propensity to choose each possible strategy, and strategies with higher propensities have a higher probability of being chosen. Propensities are defined at the subject and strategy level, so that the propensity for subject n to employ strategy j is denoted as q_{nj} (Erev and Roth 1998).

When subjects choose a strategy, they receive payoff x which results in reinforcement $R(x)$ (Erev and Roth 1998). They then update their propensity to choose that strategy again based on the reinforcement they receive, such that $q_{nj}(t+1) = q_{nj}(t) + R(x)$ if strategy j was chosen, and $q_{nj}(t)$ otherwise (Erev and Roth 1998). The better the choice performs, the greater the reinforcement term $R(x)$ is, and in turn the more the subject's propensity for that choice increases.

When applying the model to the context of this experiment, roadway travel can be thought of as an anti-coordination game where drivers interact strategically through their route-choices. Drivers choose strategy j from the set of possible route choices, and they benefit from choosing differently than others. The set of n players is comprised of the 39 subjects that participated in each session. The set of j strategies consist of either route "A" (the main route) and route "B" (the alternate route). The payoff x is a subject's earnings from a given trial; the reinforcement function $R(x)$ is explicitly defined later in the section.

The presence of multiple network states necessitated modifications to the model. It was apparent from the data that the initial share of diverting subjects varied significantly for each state, indicating that route-choice propensities varied by state. As a result, route-choice propensities were defined not only at the subject and route level, but at the state-level as well. (q_{Ans}, q_{Bns}) defines a set of user-state propensities for routes "A" and "B" in a network with states $s = 1, 2, \ldots, S$ and users $n = 1,2, \ldots, N$.

State-specific route choice probabilities are derived from the state-specific propensities. The probability of choosing a route in a given state is equal to the quotient of the propensity to choose that route in that state and the sum of all route propensities in that state. Thus, the probability P of a given subject choosing route A during a given state would be denoted as:

$$P_{Ans} = \frac{q_{Ans}}{q_{Ans}+q_{Bns}}$$

A second modification to the model implements "state-specific" propensity updating in both the error-free learning model and the hybrid learning model. This modification specifies that when a route is chosen in a certain state, the route-choice propensity for only that specific state is updated. This change is implemented through the reinforcement term $R(x)$, and is where the three models being tested differ from one another. Recall that $R(x)$ determines how past outcomes affect route-choice propensities. In these models, it is specified as a linear function of a subjects' earnings each trial, x. This implies that $R(x)$ is also a linear function of travel time, TT, because the subtractive payment scheme used to determine x is linear in TT. Specifically, $R(x)$ was defined as the difference between the maximum travel time possible in a trial, TT_{max}, and the travel time experienced by the subject, TT_n . The maximum travel time serves as a reference value, which Kahneman and Tversky (1979) describe as a cognitive anchor from which subjects perceive differences in outcomes. Erev and Roth hold this parameter fixed in the simplest form of their model, but acknowledge that the explanatory power of the model can often be significantly increased by freeing it. I tested an alternate specification that estimated a different reference point for each unique state of the network. The resulting gains in explanatory power were minimal and did not justify the use of extra parameters.

In the optimal learning model with state-specific updating, the reinforcement term is also state-specific. It is equal to the reference payout minus the travel time experienced by subject n choosing route j during state s. That is, $R_{jns}(x) = TT_{max} - TT_{jns}$. Every time a subject chooses a route in a particular state, $R(x)$ is added only to the route-choice propensity for that same state.

Without state-specific learning, the reinforcement term is only subject and route specific. That is, $R_{jn}(x) = TT_{max} - TT_{jn}$. This means that any time a route is chosen, regardless of the state, $R(x)$ is added to the route-choice propensity for all states. This type of updating will be referred to as "state-agnostic" from this point on.

Synthesizing these model components, the three learning models tested are described below:

MODEL 1 (optimal **state-specific** propensity updating only):

$q_{Ans}(t) = q_{Ans}(t-1) + [TT_{max} - TT_{Ans}(t-1)]$

$q_{Bns}(t) = q_{Bns}(t-1) + [TT_{max} - TT_{Bns}(t-1)]$

where

TT_{max} is the reference value

TT_{Ans} is the travel time experienced by user n on route A during state s.

In this model specification, subjects only update state-specific propensities based on previous travel times experienced during the same state. For example, if subject 1 chooses route A during state 1, q_{A11} is updated, but not any other $q_{Aj1,\ j\neq 1}$.

MODEL 2 (sub-optimal **state-agnostic** propensity updating only):

$$q_{Ans}(t) = q_{Ans}(t-1) + [TT_{max} - TT_{An}(t-1)]$$

$$q_{Bns}(t) = q_{Bns}(t-1) + [TT_{max} - TT_{Bn}(t-1)]$$

In this model specification, subjects update route-choice propensities for all states based on their travel time from the last trial, regardless of which state the last trial occurred in. For example, if subject 1 chooses route A, q_{Aj1} are updated for all states j. In this model specification, it is as if the existence of multiple states is completely ignored for updating.

MODEL 3 (sub-optimal **hybrid** model):

$$q_{Ans}(t) = q_{Ans}(t-1) + \alpha * [TT_{max} - TT_{Ans}(t-1)] + \beta * [TT_{max} - TT_{An}(t-1)]$$

$$q_{Bns}(t) = q_{Bns}(t-1) + \alpha * [TT_{max} - TT_{Bns}(t-1)] + \beta * [TT_{max} - TT_{An}(t-1)]$$

This model combines both "**state-specific"** and "**state-agnostic"** updating. α and β sum to 1 and are the respective contributions to relative contributions to propensity updating from "**state-specific**" and "**state-agnostic**" travel times, respectively.

Comparing the relative fit of these models for the two-state and five-state treatment data can test whether an increase in the number of network states contributes to human learning errors. If the error-free "state-specific" model is the best fit for the two-state data, while either the "indiscriminate" or "hybrid" model is the best fit for the five-state data, then the experimental hypothesis will be supported.

Estimation

Each route choice model was estimated from experimental data to determine which were the best fit for the two-state and five-state sessions. The covariates of interest in these models are the reinforcement terms $TT_{max} - TT_{jns}$ and $TT_{max} - TT_{jn}$, which describe state-specific and state-agnostic propensity updating, respectively.

Ideally, one could directly observe subjects' route-choice propensities to determine whether they update in response to state-specific choice outcomes, state-agnostic choice outcomes, or both. However, only subjects' actual route-choices were observed. Therefore, each model was re-written in terms of route-choice probabilities, shown below:

MODEL 1 (optimal **state-specific** propensity updating only):

$$P_{Ans}(t) = \frac{q_{Ans}(0) + \sum_{i=1}^{t}[TT_{max} - TT_{Ans}(i-1)]}{q_{Ans}(0) + \sum_{i=1}^{t}[TT_{max} - TT_{Ans}(i-1)] + q_{Bns}(0) + \sum_{i=1}^{t}[TT_{max} - TT_{Bns}(i-1)]}$$

$$P_{Bns}(t) = 1 - P_{Ans}(t)$$

where

$q_{jns}(0)$ denotes initial route-choice propensities

$\sum_{i=1}^{t}[TT_{max} - TT_{jns}(i-1)]$ is the state-specific travel-time reinforcement received by subjects over all prior trials

MODEL 2 (sub-optimal **state-agnostic** propensity updating only):

$$P_{Ans}(t) = \frac{q_{Ans}(0) + \sum_{i=1}^{t}[TT_{max} - TT_{An}(i-1)]}{q_{Ans}(0) + \sum_{i=1}^{t}[TT_{max} - TT_{An}(i-1)] + q_{Bns}(0) + \sum_{i=1}^{t}[TT_{max} - TT_{Bn}(i-1)]}$$

$$P_{Bns}(t) = 1 - P_{Ans}(t)$$

where

$q_{jns}(0)$ denotes initial route-choice propensities

$\sum_{i=1}^{t}[TT_{max} - TT_{jn}(i-1)]$ is the state-agnostic travel-time reinforcement received by subjects over all prior trials

MODEL 3 (sub-optimal **hybrid** model):

$$P_{Ans}(t) = \frac{q_{Ans}(0) + \sum_{i=1}^{t} \alpha * [TT_{max} - TT_{Ans}(i-1)] + \beta * [TT_{max} - TT_{An}(i-1)]}{q_{Ans}(0) + \sum_{i=1}^{t} \alpha * [TT_{max} - TT_{Ans}(i-1)] + \beta * [TT_{max} - TT_{An}(i-1)] + q_{Bns}(0) + \sum_{i=1}^{t} \alpha * [TT_{max} - TT_{Bns}(i-1)] + \beta * [TT_{max} - TT_{Bn}(i-1)]}$$

$$P_{Bns}(t) = 1 - P_{Ans}(t)$$

where

$q_{jns}(0)$ denotes initial route-choice propensities

$\sum_{i=1}^{t}[TT_{max} - TT_{jns}(i-1)]$ is the state-specific travel-time reinforcement received by subjects over all prior trials

$\sum_{i=1}^{t}[TT_{max} - TT_{jn}(i-1)]$ is the state-agnostic travel-time reinforcement received by subjects over all prior trials

These linear models of probability were then replaced with logit models to ensure that estimates were constrained to values that guaranteed positive route-choice probabilities. Furthermore, the expressions were simplified by replacing subjects' entire history of past travel times with a single lagged reinforcement term. The result is a model where the logit of route-choice probability is a linear function of the most recent relevant travel-time outcome. In the case of state-specific propensity updating, this outcome would be:

$TT_{max} - TT_{jns}$	if route j was chosen last time the network was in the same state, and
$-[TT_{max} - TT_{jns}]$	if route j was **not** chosen last time the network was in the same state

Thus, the probability of a subject choosing route "A" simplifies to:

$$logit(P_{Ans}(t)) = TT_{Bns}(t-1) - TT_{Ans}(t-1)$$

Although much simpler, this expression still captures the essence of traditional reinforcement learning models. The probability of a route being chosen in the future increases after the route provides favorable outcomes, and the probability decreases after alternatives perform well.

These models were estimated by logistic regression, where the dependent variable was a binary indicator of whether a subject diverted to the alternate route during a given trial. The covariates of interest, reinforcement terms comprised of travel times experienced by subjects during prior trials, are labeled and defined below.

Travel Time Last Round *(*$TT_{Bns}(t-1) - TT_{An}(t-1)$*)***:** This covariate is the travel time experienced by a subject in the previous trial. It is positive if the subject chose the alternate route ("B"), and negative if the subject chose the main route ("A"). It pertains to propensity updating on a "trial-to-trial" basis, independent of state, and is used in both the "state-agnostic" and "hybrid" models.

Travel Time Last State *(*$TT_{Bns}(t-1) - TT_{Ans}(t-1)$*)***:** This covariate is the travel time experienced by a subject during the last instance of a trial whose state matched the current one. It is positive if the subject chose the alternate route, and negative if the subject chose the main route. This covariate pertains to propensity updating on a state-specific basis; it is used in the "state-specific" and "hybrid" models.

The following controls were also included in the regression:

Diverted Last Round – Binary indicator of whether a subject diverted to the alternate route during the previous trial. This variable captures the contribution of the "inertia effect", or the inherent tendency to repeat a past choice, to propensity updating.

Diverted Last State – Binary indicator of whether a subject diverted to the alternate route during the last trial in which the same state as the current one occurred. This variable captures the "inertia effect" on a state-specific basis.

First State Occurrence – Binary indicator of whether it's the first time a given network state occurs during an experiment. In these instances, travel times for the most recent occurrence of the same state do not exist and are coded as zeros. Without this control the zeros would bias the covariates of interest.

No Prior Incident – Binary indicator of whether the prior trial did not have an incident (5-state treatment only). It is included because there is no congestion (and therefore no equilibrium or equilibrium learning) on trials without incidents. Trials which themselves had no incident are omitted for the same reason.

Scenario – Categorical variable capturing scenario-specific fixed effects. This variable captures the initial propensity to divert to the alternate route for each network state.

The exact specifications of the three models being estimated are shown below, with the covariates of interest in bold:

MODEL 1 (optimal "state-specific" learning only):
logit(*Probability of diverting*) = α + β****Travel Time Last State*** + η**controls*

MODEL 2 (sub-optimal "state-agnostic" learning only):
logit(*Probability of diverting*) = α + β****Travel Time Last Round*** + η**controls*

MODEL 3 (sub-optimal hybrid model):
logit(*Prob. of div.*) = α + β1****Travel Time Last Scenario*** + β2****Travel Time Last Round*** + η**controls*

To assess which learning model best describes subject behavior, the Bayesian information criterion (BIC) for each model was compared. The BIC is a modification of the likelihood function which includes a penalty term for adding model parameters; the penalty term is important for this application because the hybrid model benefits from the explanatory power of an extra covariate. The preferred model is the one with the lowest BIC. First, each model was estimated on data from the two-state treatment, and the BICs were compared to select a preferred model. Next, the process was repeated for data from the five-state treatment.

II.IV RESULTS

Two States

For the two-state network, the **state-specific** learning model specification had the lowest Bayesian Information Criterion. In this model, subjects only updated route-choice propensities based on travel times experienced during the last relevant state. This implies, for example, that subjects' travel times experienced during "major" incidents would not affect their probability of diverting during "minor" incidents. The fit of each model is shown in Table II.I.

TABLE II.I

Bayesian information criteria of learning models for the two-state treatment (in descending-order of fit)

Updating Model	**Bayesian Information Criterion**
State-specific (optimal)	760
Hybrid (sub-optimal)	766
State-agnostic (sub-optimal)	770

Estimation results from the hybrid learning model are provided in Table II.II to show the relative contributions of both state-specific and state-agnostic propensity-updating (shown in bold) to route-choice probability. The results indicate that subjects update route-choice propensities for each state based on travel times experienced during the last occurrence of the same state, but not based on travel times experienced during the most recent trial. This corroborates the finding based on the BIC indicating that the learning model with only state-specific updating fits the data best.

TABLE II.II

Logistic Regression Results for Hybrid Model in the 2 State Case

Logistic Regression		Number of obs = 608	
Log likelihood = -360.81		Pseudo R2 = 0.0879	
Subject Diverted	Coef.	Std. Err.	P-value
Travel Time Last State	**-0.505**	**0.158**	**0.001**
Travel Time Last Round	**0.000**	**0.142**	**0.999**
Diverted Last Round	-0.347	1.238	0.779
Diverted Last State	5.472	1.410	0.000
First State Occurrence	2.906	0.789	0.000
Scenario			
3	0.540	0.209	0.010
_cons	-3.327	0.659	0.000

The coefficient for **Travel Time Last State** is negative and significant (at the 1% level), indicating that route choices resulting in longer travel times are less likely to be repeated during the next occurrence of the same state.

The coefficient for **Travel Time Last Round** is 0 and not significant, indicating that travel times from the most recent trial have no effect on route choice in the following trial.

These results support the hypothesis that in simple networks with only two states, optimal state-specific updating can be achieved by human users. Experimental subjects were able learn exclusively from state-specific outcomes, disregarding irrelevant outcomes from prior trials with different states.

Five States

For the five-state network, the **hybrid** learning model specification had the lowest Bayesian Information Criterion. In this model subjects updated route-choice propensities based both on travel times experienced during the last relevant state, as well as travel times experienced during the most recent trial (regardless of state). This implies that subjects' travel times experienced during "major" incidents would affect their probability of diverting during "minor" incidents. The fit of each model is shown in Table II.III.

TABLE II.III

Bayesian information criteria of learning models for the five-state treatment (in descending-order of fit)

Updating Model	**BIC**
Hybrid (sub-optimal)	4547
State-specific (optimal)	4552
State-agnostic (sub-optimal)	4574

Estimation results from the hybrid learning model are provided in Table II.IV to show the relative contributions of both state-specific and state-agnostic propensity-updating (shown in bold). The results indicate that subjects update state-specific route choice propensities based both on travel times experienced during the last occurrence of the same state, and also travel times experienced during the most recent trial - regardless of the state. This corroborates the finding based on the BIC that the learning model with both state-specific and state-agnostic updating fits the data best.

TABLE II.IV

Logistic Regression Results for the Hybrid Model in the Five State Case

Logistic regression		Number of obs = 3536	
Log likelihood = -2233		Pseudo R2 = 0.0595	
Subject Diverted	Coef.	Std. Err.	P-value
Travel Time Last State	**-0.268**	**0.045**	**0.000**
Travel Time Last Round	**-0.137**	**0.037**	**0.000**
Diverted Last Round	1.483	0.334	0.000
Diverted Last State	3.284	0.439	0.000
First State Occurrence	1.549	0.232	0.000
Scenario			
2	0.345	0.119	0.004
3	0.677	0.109	0.000
4	0.672	0.116	0.000
No Prior Incident	0.469	0.112	0.000
_cons	-3.187	0.276	0.000

Unlike in the two-state case, the estimate of the coefficient for **Travel Time Last Round** is negative and significant for the five-state treatment data. This confirms that subjects in the five-state network make route-choice decisions based on travel times experienced in the trial immediately prior, regardless of the network state during that trial. A possible explanation for this behavior is that subjects struggle with the cognitive challenge of updating route-choice propensities for each state using past travel times exclusively from the same state. This explanation is consistent with the fact that subjects then also incorporate the last trial's travel times - which are easier to remember – possibly to compensate. This adjustment is counterproductive, however, and impairs equilibrium learning. Evidence presented in the next section demonstrates the detrimental effect of state-agnostic updating on equilibrium learning.

Evidence that State-agnostic Updating is Counterproductive

Although it has been proven that state-specific learning is theoretically optimal in a multi-state network (Watkins and Dayan 1989), it is still possible that outside of equilibrium, learning based on the most recent travel time can be rational behavior rather than a counterproductive response to memory limitations. This could be true in a circumstance where the average performance of the main route relative to the alternate route remains similar from trial to trial and does not undergo substantial shocks due to state changes. That is, if the inter-state variation in route choice performance does not exceed the intra-state variation, then incorporating information from the most recent trial could be beneficial. Furthermore, even if the outcome of the most route choice is on average uninformative to the next route choice decision, updating based on tit might still be rational if subjects believe that route performance is similar among similar states. The five-state treatment defines states less coarsely in terms of incident severity (none, minor, medium, major, severe) compared to the two-state treatment (minor, major). If subjects update based on the most recent travel time only when the current and previous states are similar, it reflects a plausible assumption that travel times from similar states are informative.

After testing for these possibilities using data from the five-state treatment, I find that neither occur. Table II.V shows that the relative performances of the two routes are not correlated between immediately consecutive trials. At the same time, this measure is highly correlated between sequential trials of the same state.

TABLE II.V

Correlation between Current and Prior-trial Route-choice Outcomes with Five States

Time Savings from Diverting during the ____ Round		**Correlation Coefficient**
Current	Previous	-.01
Current	Previous	.92

Furthermore, subjects did not exercise discretion when updating route-choice propensities based on the most recent travel time. Similarity between current and prior network states played no role in the extent to which the most recent travel time was used for propensity updating. The propensity to divert during minor-incidents, for example, is no less affected by major-incident travel times than by medium-incident travel times – even though the medium incident is more similar to a minor incident. This was tested by re-estimating the hybrid learning model on the five-state experimental data with the updating covariate **Travel Time Last Round** recoded as two separate variables: **Travel Time Last Round_Similar** when the most recent trial was in a similar state to the current trial, and **Travel Time Last Round_Different** when the most recent trial was in a substantially different state from the current trial. Table II.VI shows this estimation result.

TABLE II.VI

Regression Results for five-state treatment, separating the "prior-trial" updating covariate into two based the similarity of network state between the current and prior trial.

Subject Diverted	Coef.	Std. Err.	P>z
Travel Time Last State	-0.267	0.045	0.000
Travel Time Last Similar Round	**-0.120**	**0.037**	**0.001**
Travel Time Last Different Round	**-0.154**	**0.037**	**0.000**
Diverted Last Round	1.444	0.333	0.000
Diverted Last State	3.260	0.439	0.000
First State Occurence	1.553	0.233	0.000
Scenario			
2	0.324	0.120	0.007
3	0.655	0.110	0.000
4	0.650	0.116	0.000
No Prior Incident	0.440	0.113	0.000
_cons	-3.143	0.278	0.000

The magnitude of the coefficient for **Travel Time Last Similar Round** is not larger than that of **Travel Time Last Different Round**, implying that when subjects update route-choice probabilities based on outcomes from irrelevant states, they are not doing so because they believe similar states provide similar outcomes. Thus, the experimental evidence strongly suggests equilibrium learning based on travel times from route-choices immediately prior, disregarding state, is counterproductive and irrational. No plausible scenarios were found for updating based on the most recent travel time outcome to be rational behavior. This was reinforced by examining the correlation between subjects' monetary earnings during the experiment and the extent to which their route choices were influenced by the most recent travel-time experienced, regardless of state. The hybrid route-choice model was re-estimated for two groups of subjects: the highest earning quartile and the lowest earning quartile. Table II.VII shows these results.

TABLE II.VII

Marginal Effect of 10 Seconds of Travel Time on the Probability of Diverting for the Top and Bottom Quartiles of Subjects

Quartile of subjects	**Marginal effect for Travel Time Last Round**	**P > |z|**	**Marginal effect for Travel Time Last State**	**P > |z|**
Top performing	-.017	.331	-.048	.012
Worst performing	-.035	.018	-.048	.013

The marginal effect for the estimate of **Travel Time Last Round** is statistically significant for the lowest-earning quartile of subjects, but not for the highest-earning quartile. At the same time, there was no difference in the magnitude or precision of the marginal effect of **Travel Time Last State** between the two groups. In other words, the worst-performing subjects updated route-choice probabilities based on prior outcomes from irrelevant states, while the best-performing subjects did not. The supplementation of state-relevant information with recent information that ignores the network state was demonstrably counterproductive and detrimental to learning.

Simulation

The estimation of learning models from experimental data was able to demonstrate that subjects commit route-choice learning errors in networks with multiple states. The long-run implications of these errors on equilibrium convergence could not be observed, however, because the time horizon of the experiment was limited to 20 trials. Therefore, a long-run simulation of drivers' route choice behavior was performed. The simulated environment replicated that of the route choice experiment. 39 virtual subjects simultaneously “chose” between one of two alternate routes each trial. The simulated agents then received feedback in the form of a hypothetical travel time associated with their route choice. The travel times were based on a congestion function estimated from actual travel times experienced by experimental subjects. This congestion function maps the number of simulated agents choosing each route to the hypothetical travel-time feedback they receive.

The route-choices of the simulated agents were determined by the exact reinforcement learning model specified earlier in this paper. Recall that in the five-state case, the best-fit learning model was a hybrid model where subjects’ route choices were based on both state-specific and state-agnostic route choice outcomes. This model specifies route-choice probabilities for each agent as a function of initial route-choice propensities and cumulative travel-time feedback. Thus, each time a simulated agent “chose” a certain route, their probability of choosing that route again would update based on the hypothetical travel time feedback it received. Recall the expression for the probability that simulated agent n would choose route A during state s in period t:

$$P_{Ans}(t) = \frac{q_{Ans}(0)+\sum_{i=1}^{t}\alpha*[TT_{max}-TT_{Ans}(i-1)]+\beta*[TT_{max}-TT_{An}(i-1)]}{q_{Ans}(0)+\sum_{i=1}^{t}\alpha*[TT_{max}-TT_{Ans}(i-1)]+\beta*[TT_{max}-TT_{An}(i-1)]+q_{Bns}(0)+\sum_{i=1}^{t}\alpha*[TT_{max}-TT_{Bns}(i-1)]+\beta*[TT_{max}-TT_{Bn}(i-1)]}$$

where

$q_{jns}(0)$ denotes initial route-choice propensities

$\sum_{i=1}^{t}[TT_{max}-TT_{jns}(i-1)]$ is the state-specific travel-time reinforcement received by subjects over all prior trials

$\sum_{i=1}^{t}[TT_{max}-TT_{jn}(i-1)]$ is the state-agnostic travel-time reinforcement received by subjects over all prior trials

α and β sum to 1 and are the relative contributions to propensity updating from “state-specific” and “state-agnostic” travel times, respectively. Their values were derived from the best-fit learning model estimated from the human-subjects experiment, specifically the marginal effects of the covariates **Travel**

Time Last State and **Travel Time Last Round** on route choice probabilities. Because the means of the travel times in the simulation match those of the actual experiment, the ratio of α to β was set as the ratio of the estimated marginal effects. Table II.VIII below shows that the marginal effect for **Travel Time Last State** is approximately twice that of **Travel Time Last Round**.

TABLE II.VIII

Marginal Effects of 10 Seconds of Travel Time on the Probability of Diverting for both "State-specific" and "State-agnostic" Updating.

Covariate of interest	**Effect of 10 secs of travel time on probability of diverting**	**Standard Error**
Travel Time Last State	-.059	.010
Travel Time Last Round	-.030	.008

Thus, the value of α was fixed to be twice that of β. α and β could then be calculated from solving the two simultaneous equations:

1. $\frac{\alpha}{\beta} = 2$
2. $\alpha + \beta = 1$

The solution is $\alpha = .67;\ \beta = .33$

Next, the set of initial route-choice propensities assigned to all simulated agents was determined. This was done by first determining the ratio of initial state-specific propensities, $\frac{P_{Aj}(0)}{P_{Aj}(0) + P_{Bj}(0)}$ for states $j = 1,2,3,4,5$, which are equivalent to subjects' initial probability of diverting. Thus, they were set as the initial experimental diversion rate, averaged across all five-state sessions, for each state. Next, the magnitude of initial propensities was calibrated using a procedure from Erev and Roth (1998). The calibration is performed by iterating the simulation over a range of possible parameter values to find the magnitude of initial propensities that minimizes the sum of mean square differences between simulated and empirical diversion rates.

Figure II.I shows the fit of the simulation, averaged over 1000 iterations, to experimental data:

FIGURE II.I

Fit of simulation to experimental data for the five-state treatment

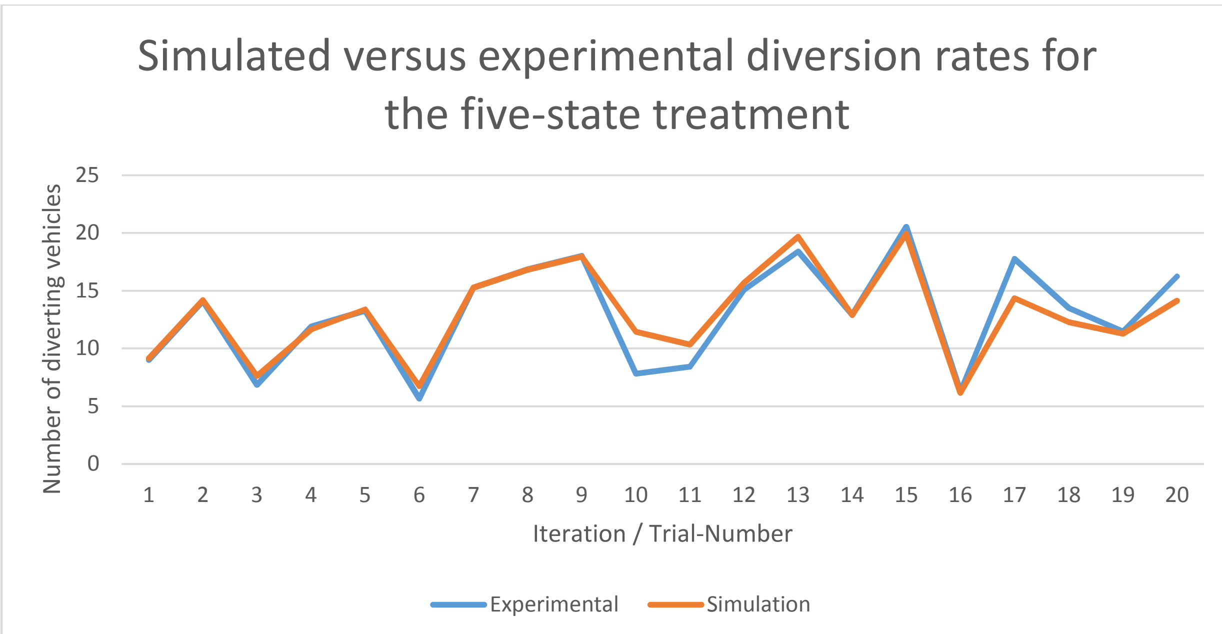

Based on visual inspection, the simulation using the empirically-determined model with estimated and calibrated parameters fit the experimental data well.

The horizon of the simulation was extended well-beyond the number of experimental trials to study longer-run implications of learning inefficiency on convergence to network equilibria. The order of incidents was preserved in the extended the simulation; the sequence was simply repeated several times.

The simulation results were then compared to that of a counterfactual simulation representing "error free" state-specific learning in a five-state network. The propensity-updating parameters for this simulation came from the best-fit learning model estimated for the two-state network treatment, where subjects learned from past travel times from relevant states only. It provided a prediction of equilibrium learning in a hypothetical case where subjects in a five-state network were not cognitively limited, and thus were able to update propensities exclusively on a state-specific basis. Figure II.II shows a comparison of equilibrium convergence between the empirical and counterfactual simulations for the four states with incidents.

FIGURE II.II

Convergence to Equilibrium Diversion Rates in Simulations Based on “Actual” Experimental Data and "Ideal" Counterfactual Behavior

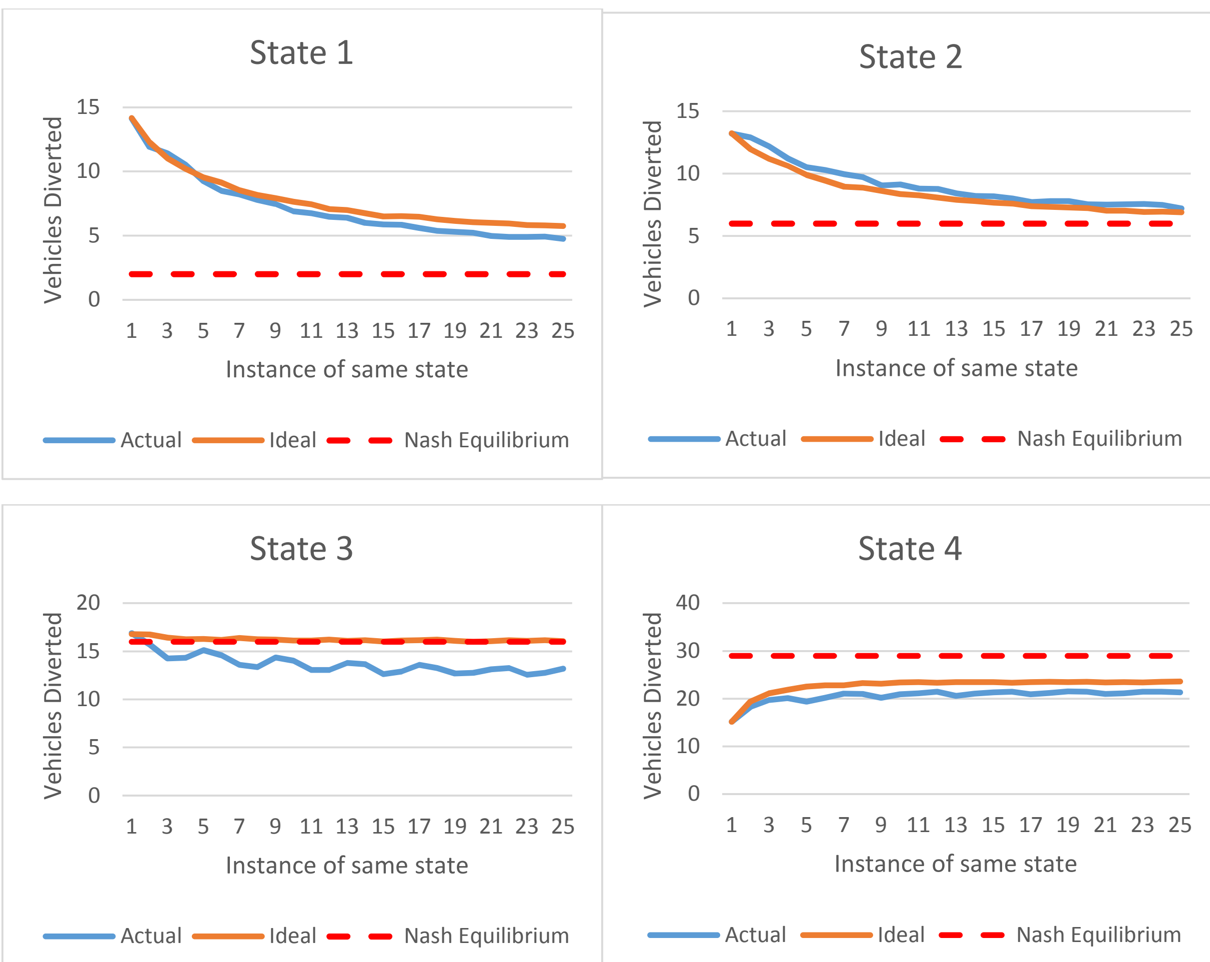

The equilibrium route split was approached more rapidly in the counterfactual simulation for three out of four states (recall that the no-incident state was omitted from the analysis). Furthermore, simulated agents always tended towards the Nash equilibrium route-split in the counterfactual simulation. In the empirically-derived simulation, however, divergence from the equilibrium occurred in one of the states (State 3, "Major" Incident). These results are evidence that learning based exclusively on state-specific propensity-updating approaches equilibrium more efficiently than learning that also incorporates the most recent state-agnostic outcome. The former was attainable when the network only had two states, but the not when the number of states was increased to five. This suggests that the presence of five states in the network imposed a cognitive burden on subjects which impaired their ability to learn.

Limitations

Although any learning model specification would have likely yielded estimates showing that subjects used irrelevant information in the five-state network, the choice to use a reinforcement learning model potentially influenced the simulation results significantly. The different types of models discussed in the literature review displayed substantial differences in long-run convergence behavior. Thus, it is possible that a different type of model without reinforcement learning would produce either greater or smaller differences between the empirical and counterfactual simulations. The differences between the two simulations would also be more substantial if some degree of memory decay was included in the model. In the current simulation the marginal effect of past outcomes on route-choice propensities is continually decreasing, so that after many iterations the differences between "state-specific" and "state-agnostic" updating became negligible.

Furthermore, simulated agents were modeled under the assumption that the subject behavior and learning observed during the 20-trial experiment would persist indefinitely. It is possible that even after that 3 practice trials and 20 paid trials, much of the behavior observed in the experiment could still be exploration rather than optimization (this tendency is discussed in Ben-Elia and Shiftan (2010)), and that learning behavior could change significantly after more time.

II.V CONCLUSION

Route-choice learning in a traffic network was shown to be adversely impacted by the presence of multiple network states. The aggregate route-choices of subjects participating in a traffic experiment converged towards the user equilibrium more slowly, or in some cases failed to converge at all, when the number of network states was increased. Statistical analyses of the experimental data revealed that while learning was seemingly unaffected in the presence of two unique states, the presence of five states caused subjects to incorporate irrelevant information into route-choice decisions. These findings suggest that additional network states increase the cognitive burden of optimal route-choice learning beyond the capabilities of many human subjects. In the network with several states, subjects instead based route-choice decisions in part on irrelevant travel-time outcomes from different network states. The fact that the irrelevant information was gleaned from the trial immediately prior, whose outcome is easiest to remember, suggests that memory played a role in learning inefficiency. That is, subjects might utilize the most recent (albeit irrelevant) travel times to compensate for difficulty in recalling travel times from the appropriate state.

The role that increased network states played in the occurrence of cognitive errors during equilibrium learning is a novel finding. Cognitive limitations have been shown by other researchers to delay learning; this study shows that multiple network states exacerbate the issue, causing subjects to make route-choice decisions based on irrelevant information that disregards the network state. A simulation based on these empirically determined learning errors found that over an extended time horizon, convergence to the equilibrium route-choice allocation is significantly delayed and, in some cases, non-existent.

Although this experimental setting examined state transitions specifically, it likely that learning is adversely affected in any network with operational complexity that challenges the memory and cognitive abilities of users. In light of this experiment's results, interventions on behalf of planners that address learning limitations would likely increase user welfare in operationally complex networks. For the specific case of multi-state networks, traveler information systems that aid memory (e.g. individualized tracking/feedback systems) could improve learning and ultimately quicken and/or facilitate convergence to user equilibria.

Future work could address the limitations of the simulation analysis. Strong untested behavioral assumptions were imposed that potentially influenced the simulation's findings. Future work could repeat the human subjects experiments over a longer time horizon to verify that behavioral patterns observed in the first 20 trials persist. Furthermore, the implications of using different classes of model for the simulation should be explored.

CHAPTER III

DO TRUTHFUL-BIDDING MECHANISMS IMPROVE THE ALLOCATION OF DRIVERS TO EXPRESS LANES?

III.I INTRODUCTION

Priced freeway lanes, which now exist in ten states and are becoming increasingly prevalent, can improve aggregate driver welfare by accounting for heterogeneity in their willingness to pay for time savings. (Small and Yan 2001, Arnott et al. 2002). Users who value time savings more highly can pay to access lanes with faster service, while those who value them less get longer travel times but avoid tolls (Small and Yan 2001). In practice, however, welfare-maximizing pricing is not implemented. Instead, express lanes are dynamically priced to maintain a minimum free-flow speed. This pricing regime often results in significant underutilization of the express lane, and may be less efficient than no pricing at all (Small and Yan 2001).

Welfare-maximizing tolls have been the subject of a wide body of theoretical work including studies by Small and Yan (2001), Verhoef and Small (1999), and Arnott et al. (2002). They consist of single price applied to one or more lanes that achieves an allocation of drivers between the toll and free lanes such that the sum of congestion costs across all users is minimized. A crucial assumption made in these works is that toll operators know the distribution of values that each user places on travel time savings. In practice, however, welfare-maximizing tolls are not implemented because transportation agencies lack the means to determine each drivers' willingness to pay[3]. Without accurate estimates for these preferences, the toll will likely be set either too high or too low and result in inefficient utilization of the fast lane(s). Moreover, achieving sufficient knowledge of these driver preferences is prohibitively difficult. Values of time can vary dramatically by person, trip purpose, time of day, day of week, and locale (Koppelman 2012, Wardman 2007, Parkany 1999, Sullivan 1998, Brownstone & Small 2005). Only limited survey and field data are available for determining driver values of time (VOT) for corridors served by toll lanes, and VOT estimates can vary significantly for a specific locale based on data collection and model specification (Koppelman 2012) Furthermore, VOT and corresponding parameter estimates used to forecast VOT vary from region to region (Koppelman 2012), necessitating a significant amount of effort to determine a VOT distribution anytime express lanes are introduced to a new area. Lastly, VOT has been shown to vary based on congestion levels (Koppelman 2012). This means that a typical VOT distribution may be invalid for express lane pricing during periods when non-recurring incidents or events result in local congestion; posing an additional challenge for incident management using lane pricing.

Using the Vickrey-Clarke-Groves (VCG) mechanism to determine tolls and lane assignments could obviate the need to estimate valuations of time savings by eliciting them directly from drivers. VCG mechanisms select socially optimal outcomes by incentivizing participants to truthfully reveal their preferences. Their design ensures that no matter which outcome is selected, no participant can increase their utility by providing non-truthful preferences. If the mechanism was used to manage a toll lane, it would be in each driver's best interest provide their truthful travel time preferences – likely transmitted remotely through connected vehicle technology which is becoming increasingly prevalent. These preferences could then be used to determine and assign a lane allocation for drivers that achieves essentially the same welfare benefits as second-best tolling.

---

[3] Furthermore, express lanes converted from High Occupancy Vehicle lanes are required to provide a minimum speed of 45 miles per hour.

The use of VCG mechanisms in practice is limited, in part because there is mixed evidence regarding the extent to which these mechanisms able to elicit truthful preferences from the average user. One notable application of the mechanism is the pricing and allocation ads through the Facebook application, but its use in real-world transportation settings is non-existent. This work explores the feasibility of Vickrey-Clarke-Groves mechanisms for allocating freeway lane capacity to drivers. The primary aim is to understand the extent to which actual human users will provide truthful travel time preferences in practice, which will in turn determine the extent to which the theoretical welfare benefits of implementing such a mechanism can be achieved.

Although VCG mechanisms incentivize truthful preference revelation, they have failed to achieve perfect revelation in every experimental study reviewed prior to this project. VCG mechanisms are difficult to grasp conceptually, especially the strategic dominance of truth-telling. Furthermore, most new processes require some degree of learning. Travelers and participants in other strategic endeavors in general have a strong tendency to experiment and explore (Ben-Elia and Shiftan 2009, Erev and Roth 1998), which would result in preference misrevelation - at least initially. This misrevelation could also persist, however, because sub-optimal equilibria are possible in many applications of the mechanism. Perfect truth-telling is not required to achieve near-optimal outcomes, however; the mechanism outperforms alternative allocations schemes in experiments by Healy (2006), Brenner and Morgan (1997), and McLaughlin & Friedman (2016), despite some degree of misrevelation by subjects. In other studies, however, misrevelation is so substantial that the mechanism fails to increase user welfare over that of other schemes (Attiyeh et al. 2000, Kawagoe and Mori 2001). Freeway capacity allocation shares some traits with applications in prior experiments where the VCG outperformed other alternatives, and other traits with applications where the VCG was inferior. This is discussed in greater detail in the literature review, however the upshot is that there is uncertainty regarding the potential empirical performance of the mechanism in the context of this study. Furthermore, this application is unique in that travel time savings are allocated rather than goods. It is possible that humans have a different understanding of their willingness to pay for a good or service as opposed to travel time. In addition, there are significant impediments to myopically learning truth-telling in the context of lane assignment. One is that the mechanism relies on congestion functions to predict average travel times for a lane as a function of volume, but deviations may occur due to idiosyncratic driver behavior. This random variation in travel times that will distort drivers' perceived relationship between the willingness to pay they provide and the travel times they experience. Another is that misrevelation resulting in sub-optimal outcomes can be an equilibrium, which is explained in greater depth in the literature review.

To study the workability of using a VCG mechanism to allocate freeway lane capacity, I developed a traffic experiment using an interactive, multi-user driving simulator where human subjects travel a freeway with an express lane. I attempt to allocate the experimental subjects between the express and general lanes using an optimal tolling scheme where users reveal their valuation of the road through a VCG mechanism before driving. I work from a framework where individuals have heterogenous preferences that are unknown to the regulator, preventing the optimal allocation of individuals across lanes. While real-world travel time preferences consist of many factors including value of time, value of reliability, and urgency costs, the experiment incentivizes subjects such that the only relevant travel time preference is value of time. This simplification greatly enhances the ease of observing the truthfulness/accuracy of revealed preferences, as well as calculating the efficiency of the resulting allocation.

This work is the first to empirically demonstrate that a VCG mechanism can elicit truthful preferences in setting where travel time is the allocated good. Despite the strong overall correlation between subjects' true and stated VOT, however, only 34% of bids were truthful, and the same degree of misrevelation persisted after 10 rounds. This work finds strong evidence that these deviations from truth-telling were due in part to limited subject understand of the mechanism, as well as stochasticity in travel time outcomes. Nevertheless, I show that given the amount of truth-telling observed, the mechanism may still dominate alternative pricing schemes in terms of driver welfare.

III.II LITERATURE REVIEW

Welfare benefits of second-best value pricing

The VCG mechanism studied for this project elicits VOTs so that drivers can be allocated to the express lane in a way that achieves the same outcome as an optimally-set second-best toll. This outcome minimizes travel delay cost subject to the constraint of one of more lanes being left unpriced. High VOT users use the fast lane and are charged, while the remaining users get a slower trip in the free lanes. When there is significant VOT heterogeneity, such that travel delay costs for some users are much higher than those of others, this pricing scheme can improve driver welfare over the case where no lanes are tolled (Small and Yan 2000, Verhoef and Small 1999, Arnott et al. 1991, Braid 1996). This pricing scheme also offers better performance than other currently implemented schemes such as profit maximization and keeping the fast lane in free flow (Small and Yan 2000).

Difficulty in estimating driver travel time preferences

Many studies have identified challenges associated with estimating travel time preferences needed to set efficient prices for managed lanes. The Second Strategic Highway Research Program (SHRP 2) produced the study "Improving Our Understanding of How Highway Congestion and Pricing Affect Travel Demand" (Koppelman 2012) that describes many of these issues. The study finds that value of time (VOT) can vary dramatically across income groups, vehicle occupancies, household composition, and trip purposes. Only limited survey and field data are available for determining a VOT distribution for a particular corridor, and VOT estimates can vary significantly for a specific locale based on data collection and model specification. It also finds that the distribution of driver VOTs vary by time of day, day of week, and by place. Lastly, it finds that VOT can vary based on congestion levels, which means that a typical VOT distribution may be invalid for lane pricing during periods when non-recurring incidents or events result in local congestion.

Wardman (2004) analyzed stated preference survey results and found that a travelers' VOT varied by income and trip purpose. A study by Parkany (1999) based on revealed-preference data for SR-91 users showed that commuters chose the toll lanes for a varying subset of days of the week, implying that drivers' VOT change from day to day. Sullivan (1998) made the same observation based on the same tollway. Brownstone and Small (2005) found that driver VOTs varied between California SR-91 and I-15 based on income and observed corridor distance, and that the median of the VOT distribution for each route's users were similar once the I-15 sample was reweighted to match characteristics of the SR-91 sample.

Travel-time preferences

The concept of value of time was first applied to a transportation setting by Johnson (1965), who defined it as the value of a traveler's leisure time. This type of travel-time preference is acknowledged by virtually all studies of road pricing, including Small and Yan (2001), Verhoef and Small (1999), Arnott et al. (2002), Koppelman (2012), Wardman (2007), Parkany (1999), and Sullivan (1998). The measure is a linear quantity reflecting drivers' willingness to pay for travel time savings.

Small, Winston, and Yan (2005) empirically demonstrated that drivers also value travel time reliability, which means having a route that provides consistent travel times. If arriving early or late is costly, then reliability is important to travelers (Noland and Small 1994). Braid (1996) argued these costs formed a "v-shaped" schedule delay cost function with welfare losses associated with drivers arriving both early (due to opportunity cost of leaving sooner) and late (less time spent at productive activity). He also cites empirical evidence of this functional form from Small (1982). Wardman (2004) also estimated that "late-time" is worth seven times more than typical "in-vehicle" time. The presence of these costs reflects non-linear components to travel time preferences, and also preferences regarding a traveler's departure time. Wardman (2004) explicitly estimated these preferences, and Arnott et al. (1991) analyzed these preferences in the context of efficient toll-setting (which differs from second-best value pricing).

More recent work has also identified urgency costs, defined as a willingness-to-pay for a discrete unit of time to avoid missing a deadline, as a critical component of willingness to pay for time savings (Bento, Roth, and Waxman 2014). These constitutes a unique component of travel preferences that do not scale with time.

This experimental study focuses solely on the value of time component of these preferences. Sullivan (1998) found that the need to be on time for a commitment (a key component of schedule delay and urgency costs) was not frequently cited among interviewed toll users as a reason for taking the toll road, and that value of time was more salient for these users. Starting from this simplified framework, the extent to which the mechanism reveals truthful VOT from users can be analyzed. Given the relative success of the mechanism, future work should incorporate and examine truthful elicitation of these other preferences as well.

Empirical values of time found in field studies

Values of time of actual freeway drivers have been estimated from several field studies. These studies provide a point of reference for what sensible VOTs look like for the general population, which could serve as an upper bound for comparison to the pool of student subjects participating in this study. The studies also identify demographic factors associated with VOT, which serve as a reference to see if those relationships hold when express lanes in this study's simulated road network are managed by a VCG mechanism. One survey study of California SR-91 found that the VOT of toll-users was 13-14 dollars per hour (Sullivan 1998), though this does not reveal much about the VOT distribution of freeway drivers overall (both express and free-lane users) the VOT of all freeway users - only that the median would be below this figure. Small, Winston, and Yan (2005) found from revealed preference surveys of SR-91 users that the VOT distribution of freeway drivers is a normal (bell-shaped) distribution with a mean of \$21/hour and a variance of \$10/hour, while stated preference surveys find a mean of \$12/hour with a variance of \$13/hour. These differ significantly, but both show sufficient driver heterogeneity to justify priced lanes. Based on revealed preference data from the SR91 and I-15 corridors in California, Brownstone and Small (2005) estimated the median value of time of commuters to be between \$20/hour and \$40/hr. Parkany (1999) identified some of the demographic factors that contribute to toll-lane usage. Increase age and being female were both associated with a higher probability of taking toll-lanes.

VCG mechanisms

Vickrey-Clarke-Groves mechanisms select optimal outcomes for a set of participants by eliciting truthful outcome preferences from each participant. The mechanism is generalized from the work of Vickrey (1961), who developed a sealed-bid second-price auction that incentivizes (as a dominant strategy) each bidder to bid their true value of the good. This framework was generalized to allocate public goods by Clarke (1971) and make optimal decisions by Groves (1973). Participants can never improve their utility by unilaterally misrepresenting their preferences, and thus in theory should always be truthful. This allows a planner to select that the socially optimal outcome that maximizes the sum or weighted sum of those preferences. In the case of freeway lane assignment, the set of possible outcomes would be any combination of drivers allocated between the toll and free lane. If travel time preferences are restricted to just consider value of time, the optimal outcome is the one which minimizes the sum of each traveler's travel time multiplied by their value of time.

VCG mechanisms in practice

Experimental tests of truth-telling and allocation efficiency with VCG mechanisms have yielded significantly varied results. In public good games with binary outcomes (a project is either undertaken or not), Attiyeh et al. (2000) and Kawagoe and Mori (2001) found there was substantial misrevelation and inefficiency. Attiyeh et al. find that near-truthful revelation occurs only 18% of the time, and that the efficient outcome (public good is/isn't provided) occurs only 70% of the time. The social welfare of participants was lower using the VCG mechanism than using a majority-rule voting scheme (though the welfare measure makes extremely conservative assumptions about how Clarke taxes are used, and may not appropriate). Kawago and Mori similarly find that truthful revelation occurs only 16% of the time, and that the correct binary outcome is achieved only 45% of the time. Second price auctions, which also creates a binary outcome for the user, are shown by Kagel (1995) to elicit inflated bids.

A recurring theme with binary outcome allocation is that there is a wide range of non-dominant bids that produce the same outcome as dominant bids; that is, subjects are rarely pivotal. This make it difficult for subjects to perceive a clear relationship between their bids and their payoffs and admits equilibria besides the truth-telling dominant strategy (Cason et al. 2003). As a result, the dominant strategy of truthful bidding is less apt to be learned. Higher levels of truth-telling are achieved when a payoff table is provided in the studies of Kawagoe & Mori and Cason et al.. The use of a payoff table is likely not feasible for real-time transportation applications, however.

VCG mechanisms elicit higher levels of truth-telling in experiments with continuous outcomes. These applications make the relationship between bids and outcome clearer and preclude the existence of non-truthtelling equilibria. When Cason et al. repeat their public-good experiment with a more continuous set of outcomes, truth-telling increases significantly and is learnable by more subjects. Healy conducted a public good game with more continuous outcomes but no truth-table, and finds that perfect truth-telling occurs 57% of the time, and that the outcome is not statistically different from the optimum 86% of the time. This was superior to voluntary contribution, Groves-Ledyard, Walker, and proportional tax mechanisms in terms of efficiency. McLaughlin and Friedman (2016) simulate a VCG auction with quasi-continuous outcomes (selling ad space), and find that the mechanism was 92% efficient and facilitated truth-fulness learning. The exact amount of misbidding was not explicitly stated. Lastly Brenner and Morgan (1997) run an experiment with a three-object combinatorial VCG, and find that 40% of bids matched the exact valuation of the bidder – achieving a 98% efficient allocation.

The literature shows that the mechanism fails in applications with a very small number of outcomes, and succeeds when there is more continuity. The lane-assignment application has elements of both regimes. From the subject's perspective, the most important outcome is binary - whether or not they assigned to the toll lane. However, there are also a more continuous range of outcomes for a subject depending on how many other subjects are assigned to the same lanes. For example, an outcome where 30 percent of subjects are assigned to the toll lane will result in a different toll and travel time than outcome where 25 percent of subjects are assigned. Given the wide difference in outcomes between the binary and continuous regimes in the literature, it is important to know which of these the application to freeways will resemble. Favorable outcomes from the literature find the mechanism to more efficient than all other pricing alternatives, while unfavorable outcomes find the mechanism to be dominated by alternatives. This study also adds to the literature through its novel application of a VCG mechanism to the allocation of time-savings. It is unclear whether the abstract nature of time as a good will impact the ability of users to provide truthful preferences. Lastly, this experiment is the first to include treatments where subjects are explicitly coached on the dominance of truth-telling as a revelation strategy, rather than be left to figure it out for themselves.

III.III METHODOLOGY

A laboratory experiment was designed to test whether the mechanism can elicit truthful VOTs for actual human subjects. The experiment platform consists of a real-time, interactive driving simulator implemented as a browser application. The simulator was adapted for this purpose from a prior version developed in collaboration with SiYuan Kong that was used to study driver responses to real-time traffic information by Kong et al. (2017).

The simulator supports up to 39 subjects who simultaneously travel a virtual freeway together. The freeway consisted of three lanes; eventually these lanes are separated so that one is tolled and two are left free. A bird's-eye view of the driving simulator is shown in Figure III.I.A, and a view from a single subject's perspective is shown in Figure III.I.B.

FIGURE III.I.A
Bird's-Eye View of the Driving Simulator

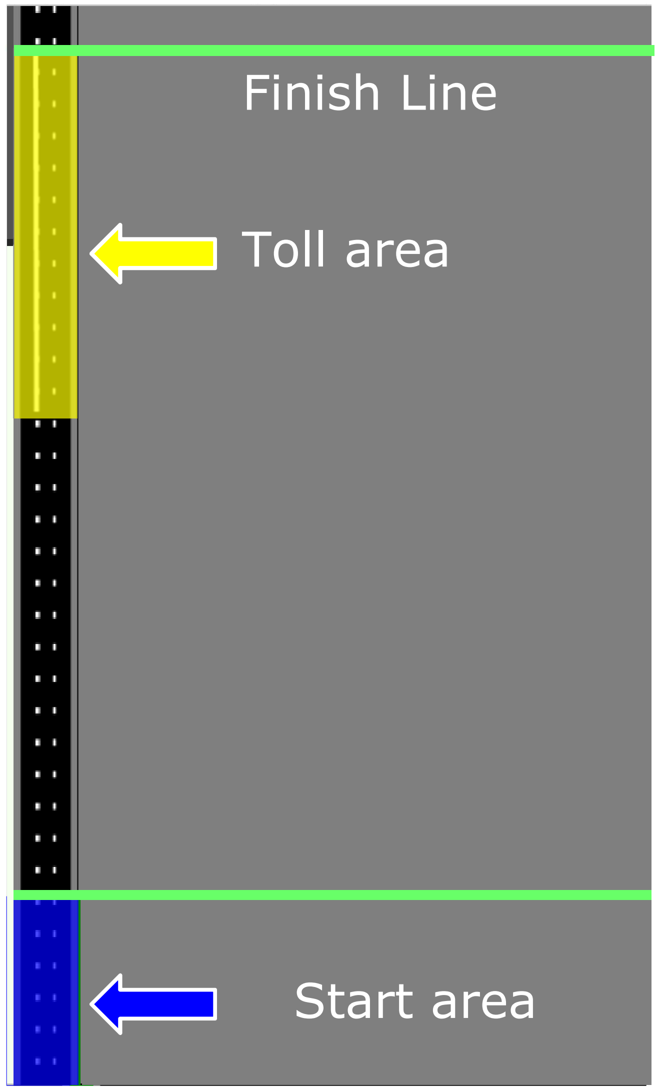

Notes: Subject-controlled vehicles are placed in the Start Area at the beginning of the trial. Subjects then travel from the Start Area towards the end of the freeway. Payment depreciation begins when they cross the Start Line, and ends when they cross the Finish Line. The freeway is three lanes wide. Subjects are free to switch between the three lanes until they reach the toll area, where the solid lane divider separates the left-most express lane from the two free lanes on the right.

FIGURE III.I.B
Subject-view of the Driving Simulator

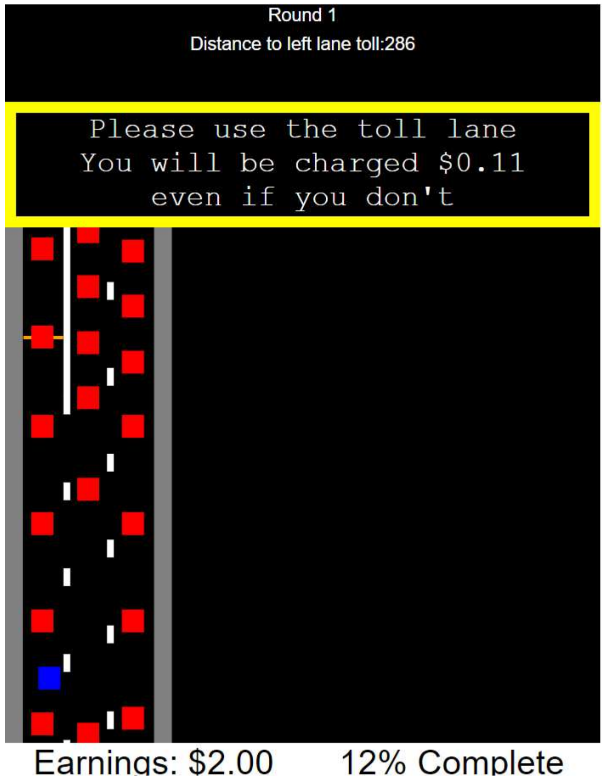

Notes: The subject-controlled vehicle is blue, all others are red. The subject is shown moving into the left-most lane of the three-lane freeway, which becomes the fast lane once the solid white lane divider is reached. The message within yellow border is informing the subject that he is assigned to the toll lane.

Subjects were given an endowment and instructed to perform a driving task which consisted of traveling the freeway for a fixed distance. They were incentivized to complete the task as quickly as possible through a linear subtractive payment scheme that deducted from their endowment for each second of travel. This payment depreciation rate served as the subjects' value of time. All subjects were dismissed from the experiment simultaneously, so that unless subjects experience disutility from simply using the simulator, the only consequence of increased travel time was greater earnings reduction. Thus, a subject incentivized by just money should value a second of time savings at the dollar amount lost per second due to the payment depreciation rate.

Before beginning each driving trial, all subjects revealed their value of time (VOT) using a slider. A VCG mechanism was then used to determine which lane each subject was assigned to and what tolls they paid (if any) based on the VOTs provided. Any toll that a subject incurred was deducted from the endowment they received at the beginning of each trial. Therefore, being assigned to the toll lane increases subject earnings by reducing the amount of payment depreciation that occurs, but reduces subject earnings by deducting the toll. Truthful revelation of VOT optimizes this tradeoff for subjects to maximize possible earnings, and thus provides both an incentive and reinforcement to bid truthfully.

To capture the dynamics of multi-driver mechanism use, subjects performed the bidding and driving phases together. This made users' tolls and travel times a function of not just their own bids, but also those of all other experiment participants.

Each session began with a presentation of instructions that explain the experiment and the VCG mechanism. Subjects then participated in two practice-trials where earnings did not count. The practice trials consisted of all 39 subjects simultaneously traveling the road network a single time together. They familiarized subjects with inputting their VOT, receiving lane assignment instructions, choosing the correct lane, and driving to the end goal. Afterwards subjects began the paid trials, where the same process was repeated13 times. Lastly, subjects completed a questionnaire once all the trials were completed. The survey solicited subject demographics, explanations for how bids were chosen, feedback about the mechanism, and other lifestyle information that may have affected how subjects bid.

To reflect the real-world phenomenon of VOT heterogeneity, subjects were randomly assigned a VOT that ranged from 10 cents to 3 dollars per second. These VOTs are equivalent to between 300 and 10000 dollars per hour, which are extremely high relative to real-world values. However, to be able to have repeated rounds over the course of an experimental session, the length of the driving task needed to be short (about a minute long, with about half that time spent driving after the lanes are separated into fast and slow lanes). The differentials between the two lanes were typically less than 10 seconds, and thus a rate of earnings decrease on the order of roughly a dollar per second was needed to make the truth-telling incentive salient.

Subjects who were given higher payment depreciation rates had a stronger incentive to take the toll lanes, since each second they spent driving was costlier. Although the experimental software assigned the VOTs to subjects, the mechanism only used self-reported VOTs for toll setting and lane allocation. This reflects what a real-world setting would be like, where operators cannot observe true user preferences. Subjects who bid high enough to take the toll lanes have the cost of the toll deducted from their earnings; this tradeoff is worth it for subjects assigned a high rate of earnings decrease, but not for those assigned a low one. All subjects left the experimental session at the same time (they waited for other subjects at the end of each round), so their only incentive was to keep as much of their initial earnings as possible.

The assigned payment depreciation rates are considered to be subjects' "true" VOTs, because they are salient given the absence of other incentives during each trial. Thus, assigning subjects a VOT allowed one to observe whether they provided their true VOT for the mechanism. It was also possible to observe the losses in earnings subjects incurred from providing non-truthful preferences, and ultimately compute a measure of user welfare. Lastly, this format facilitated the estimation of a model of how learning and demographics affect truthful revelation.

For half the sessions, each subject's VOT changed twice during the session. This tested whether subjects understood the mechanism well enough to change their bids accordingly, or whether they simply learned through trial and error. During this same group of sessions, the capacity of the free lanes was also changed once, which affected tolls and lane allocations. This tested whether subjects had the discipline to keep bidding their true VOT despite shifts in their travel times, tolls, and possibly lane assignments. For the remaining sessions, a subject's VOT changed only once and there was no change in capacity. This allowed for better observation of long-term learning. Lastly, the instructions were modified between sessions to test how varying the amount and type of information about the mechanism given to subjects affected the way that they bid. The quality of information varied for three distinct topics: how the mechanism determines lane assignment and tolls, the concept of value of time and how it relates to the experiment, and the strategic dominance of truth-telling. For each topic the explanation was either detailed or sparse.

In addition to assigned VOT sessions, a second type of treatment was used as a robustness check. For these sessions, a subject's earnings did not decrease with time. Instead, each subject left the experiment as soon his/her driving task was completed. Subjects in these sessions did not travel the freeway together; computer-controlled vehicles took the place of other subjects. This allowed subjects to complete the trials at their own pace instead of waiting on others at the end of each one. As with the prior treatment, those who bid high enough to be assigned to the priced lanes would have the toll deducted from their initial endowment. Thus, subjects balanced their innate desire to leave the session sooner (spending less time driving) against their desire to not spend money on tolls. Because each subject had already committed in advance to staying for the full duration of the session, the schedule delay costs and urgency costs were not relevant. Instead, only subjects' VOTs, which would be determined by the opportunity cost of whatever they could do if they left the session early, were relevant to their utility. In this case, a subjects VOT was however much money they were willing to spend to leave sooner. Subjects who were in more of a hurry to leave the session would find this tradeoff worth it, while those who were in less of a hurry would not. Although subjects did not travel the road together, in many sessions the mechanism used the bids from all subjects (even if they're not sharing the same roadway) to determine tolls and lane assignments. In a sense, each user's VOT was "assigned" to computer-controlled vehicles sharing the road with every other user. This preserves the dynamic aspect of bidding with the VCG mechanism.

The length of the network was significantly lengthened, and the lanes were made significantly more congestible, to increase the duration of the trials. This created much more noticeable differences between travel times in the toll lanes and free lanes, and as a result, was suitable for subjects to provide their VOT in dollars per hour rather dollars per second. A subject who provided a VOT of $0 / hour each round would take about 40 minutes in total to complete every round, while a subject who provided a VOT of $40 / hour (the maximum allowed) would take about 10 minutes to complete all rounds. Because their duration was extended, the number of trials in each session with the "innate" VOT treatment was reduced from 13 to 10.

This treatment avoids a confounding factor of the assigned VOT sessions, which is confusion over the very notion of being assigned a value of time. While VOT assignment of that type allowed for the observation of VOTs, there are not many real-world situations where people incur literal earnings losses as a linear function of travel delay. Thus, clarification is needed regarding whether deviations in bids from truthful VOT were in fact due to the inability of subjects to make the connection that their assigned depreciation rate is essentially their VOT (as opposed to conscious a decision to not reveal their true VOT). The "innate" VOT sessions do not have the issue; the salient VOT incentivizing subjects comes from their everyday preferences and opportunity costs. The truthfulness of bidding in these sessions cannot be directly observed, however. Instead, indirect measures can be used to examine whether truth-telling likely occurred. These measures include whether the mechanism elicits plausible VOTs that don't seem too low or too high, and whether elicited VOTs fluctuate significantly over time or remain stable.

Lastly, one other treatment was run that mimics conventional tolling used on present-day freeways. A single price was set for the toll-lane, and subjects chose whether to pay for the faster trip based on their VOT. This serves a robustness check to see if subjects behave rationally in the absence of a potentially confusing VCG mechanism. Supplementary Table B.I.I in Appendix B.I provides an overview of each session's treatments.

Mechanism implementation – outcome selection

In the context of freeway lane management, each possible outcome from the VCG represents a combination of drivers allocated along the fast lane and slow lanes. The value of each outcome to the driver is equal to their utility of the trip minus the cost their travel time. Trip utility is assumed to be equal for each user, and thus only the travel time costs are considered. Rather than necessitate that drivers provide their valuation of each possible outcome, one can take advantage of the fact that a driver's value of time is a linear quantity, enabling the social cost of each lane allocation to be calculated based on a single set of VOTs and congestion functions for the two sets of lanes.

Based on the values of time provided, the social cost resulting from every possible lane assignment is computed as the travel time for each set of lanes multiplied by the sum of the VOT of each user assigned to the lane. The selected outcome is the one that minimizes this cost.

For a set of a freeway users $i = 1{:}\,N$, the optimal allocation of users assigned to the fast lane can be defined as the x^* that minimizes

$$Time_{toll}(x) * \sum_{i=1}^{x} VOT_i \quad + \quad Time_{free}(N-x) * \sum_{i=x+1}^{N} VOT_i\,,$$

$$s.t.\ VOT(i) \geq VOT(i+1)$$

$Time_{toll}$ and $Time_{free}$ are the time it takes to complete the driving in the fast lanes and slow lanes, respectively, and are decreasing continuous functions of the number of users assigned to the lane. These congestion functions represent a simplifying assumption that each user of a lane exerts the same congestion externality on one another. In reality, delay cost calculations would be a dynamic problem that changes over time as more users enter and leave the freeway.

VOT_i is the value of time specified by each user i. To simplify the delay cost minimization computationally, the VOTs are sorted in descending order so that every allocation has the highest x VOT users assigned to the toll lanes.

Tolls are set using Clarke pivot, which charges users the marginal cost of their participation in the mechanism on other users. In the context of freeway lanes, the toll would be set as the externality that each driver imposes on other drivers due to the impact of his/her bid on the lane allocation. If the mechanism is used to allocate drivers spatially between lanes, rather than temporally to influence departure times, then it makes sense to take the presence of each driver on the road as given when computing externality costs. Externalities would arise from any bid above zero that changes the allocation. For example, any non-zero bid might cause an extra driver to be assigned to the fast lane, and furthermore, the marginal driver assigned to the toll lanes will likely have caused the next lowest bidder to be displaced to the slow lanes. The bidding externality is expressed as:

$Toll_j = F^*\left(VOT_{i!=j}, \boldsymbol{VOT_j}\right) - F^*\left(VOT_{i!=j}, \mathbf{0}\right)$, where

$Toll_j$ is the toll charged to user j,

$F^*\left(VOT_{i!=j}, \boldsymbol{VOT_j}\right)$ is the social cost of the optimal allocation given everyone's stated value of time,

$F^*\left(VOT_{i!=j}, \mathbf{0}\right)$ is the social cost of the optimal allocation given everyone's stated value of time **except** user j, whose value of time is fixed to 0.

The toll provides users with a dominant strategy incentive to provide their truthful value of time. Barring any errors in travel time estimates, users can never benefit from changing their bid, even ex-post. If a subject overstates his/her VOT and changes the lane-allocation compared to truth-telling, the externality cost (and thus toll he/she pays) will exceed his/her private benefit. If the user understates his/her VOT and changes the lane-allocation as a result, that user will have given up time savings worth more than the money he/she saves.

The mechanism as described above charges all drivers for using the freeway, even those in the slow lanes (provided they bid above \$0/hr and change the allocation in some way). Second-best pricing, however, described earlier as a regime where the "slow lanes" are always left un-tolled, is much more politically feasible and palatable to the public. The mechanism can therefore be modified so that only drivers assigned to the faster lane pay a toll. The toll is still set as the externality that drivers assigned to the toll lane impose on other drivers, which is made up of three components:

1. The congestion cost imposed on other drivers in the fast lane due to an extra user, minus the congestion benefit for drivers in the free lane due to one less user.
2. The delay cost imposed on any marginal drivers switched from the fast lane to the free lanes due to the VOT the user provided (compared to if that user had provided a VOT of $0/hour)
3. The congestion cost imposed on other drivers in the slow lane due to any marginal drivers switched from the toll lanes to the free lanes, minus the congestion benefit for other drivers in the fast-lane due to any marginal drivers switched from the toll lanes to the free lanes - resulting from the VOT the user provided (compared to if that user had provided a VOT of $0/hour).

In most cases, a driver assigned to the toll lane displaces a single other driver to the free lanes. As a result, components 1 and 3 cancel each other out, leaving the toll equal to component 2 - the delay cost for the marginal driver who was switched to the free lane. This is equivalent to a second-price auction, where each "winner" of fast lane access pays a price equal to the cost to the marginal "loser" of being switched to the slow lane. Given the similarity of the mechanism to an auction, the act of a user providing his/her VOT will often be referred to as "*bidding*" in this paper.

The "second-best" modification reduces the theoretical efficiency of the mechanism because drivers with a VOT below what would normally get them in the toll lane often have an incentive to slightly overstate their VOT. This is because such a driver will benefit if inflating their stated VOT raises the average free lane VOT just high enough so that another user is moved from the free lane to the toll lane – while at the same time keeping their bid low enough so that they are not switched to the toll lane. It is typically not a strong incentive because in most cases drivers cannot raise their VOT high to change the allocation without themselves becoming assigned to the toll lane. In addition, the risk of themselves being assigned to the toll lanes and overpaying for time savings serves as a deterrent to low VOT users from even attempting to inflate their bid to the "sweet spot" that provides a benefit. Simulation and experimental evidence from this project show that these perverse incentives have a very small effect on theoretical network efficiency, and an even smaller effect on actual experimental network efficiency. These results are shown in Appendix B.II.

Mechanism Implementation – Subject experience

Subjects began each round of the experiment by using a slide bar to select their value of time from a range of 0 to 4 dollar per hour. Based on the values of time provided, the VCG mechanism determined the optimal lane allocation and tolls subject to the "second-best" constraint of the slow lanes being un-tolled. As subjects approached the point where the freeway is separated into free and toll lanes, they were shown a message on their screen (representing an in-vehicle communication system) informing them which lane they were assigned to, along with what toll if any they are being charged. Subjects assigned to the toll lanes were assessed the toll no matter which lane they ultimately chose. This strengthens the incentive for high VOT users to comply with their lane assignment, and also strengthens the disincentive for low VOT users to modestly inflate their bid (by forcing them to pay the toll if they accidentally bid above the toll-lane threshold). Subjects assigned to free lanes were penalized their entire endowment for that round if they used the toll lanes, which provided a strong incentive to comply with the lane assignment.

Additional simulator details

The use of a driving simulator enhanced the realism of the experiment. For example, interactive driving dynamics often resulted in subjects needing to slow down to switch lanes, or being briefly impeded by other drivers ahead randomly reducing their speed. These instances resulted in travel time stochasticity that mimic real-world driving phenoma, which are important to capture due to their potential impact on the ability of subjects to perceive the relationship between their bids and their payouts.

Furthermore, there are perhaps factors associated with lane assignment that affect driver utility besides travel delay, such as frustration with traffic or satisfaction with travelling at high speed. The simulated driving makes it possible to see if an upward bias in VOT revelation occurs as a result. Most importantly, the driving simulator provides realistic cognitive stimulation during the "innate VOT" treatments that don't use a payment depreciation rate. This induces more realistic time preferences than alternative tasks that might occupy subjects' attention while they wait for simulated travel time to elapse.

The driving simulator was designed as 2-Dimensional top-down game implemented as a browser application using Node.js, JavaScript, and HTML5. Subjects saw a top-down view of the roadway where vehicles were represented as small colored squares - the driver's own vehicle was colored blue while all other vehicles were colored red. The driver's viewport constantly tracked his/her vehicle and presented a fixed window of visibility around it - the driver could see farther ahead than behind to simulate the forward-focused vision of real-world drivers. From top to bottom, the driver's screen contained the following elements: the secondary information area that displayed the current experiment round, the VMS display area, the driver's viewport, and the primary information area that displayed the driver's earnings and percent completion of their itinerary in real-time.

Using the keyboard Up Arrow or letter W, Left Arrow or letter A, and Right Arrow or letter D keys, drivers controlled their vehicles to accelerate or change lanes left and/or right. All vehicles accelerated at the same rate and quickly reached the same maximum speed. If a driver stopped accelerating, their vehicle decelerated at a constant rate until it reached the minimum speed. The minimum speed was designed to prevent a driver from completely blocking their lane, yet also be slow enough to prevent a driver who always traveled at the minimum speed from completing their itinerary before their entire endowment was expended. While cruising, vehicles were automatically guided to stay in the center of the nearest lane. A minimum following distance was enforced between cruising vehicles to allow space for lane changes to occur. If a subject's vehicle was obstructed by another vehicle when attempting to change lanes, the vehicle attempting the lane change was slowed down slightly to allow them to move in behind the obstructing vehicle. Drivers were informed that there were no rewards or penalties for colliding with other objects or vehicles. In addition to human controlled vehicles, computer controlled vehicles, which follow simple pre-defined control routines, were used to fill in the front of the driving platoon to create a sense of immersive traffic.

The basic engine of the simulator has elicited rational behavior in a prior study (see Kong et al. 2017), where subjects were able to reach a route-choice equilibrium between two alternate routes through reinforcement learning. This is evidence that the simulator itself is well-understood by subjects, and not a confounding factor in whether subjects choose to reveal truthful VOTs.

III.IV RESULTS

Truth-telling

Finding 1: Subjects on average tended to report their induced VOT when using the mechanism, but there was substantial mis-revelation.

The regression result in Table III.I shows that a subject's assigned VOT is a statistically significant predictor of what he/she will bid.

TABLE III.I
Regression of Revealed VOT on Assigned VOT

Revealed VOT	**Coef.**	**Std. Err**	**t**
Assigned VOT	0.549	0.038	14.39
_constant	0.668	0.042	15.79

Although there is a strong correlation between the value of time that a subject was assigned and the VOT that subject provided to the system, perfect truth-telling would result in an estimated coefficient of 1.000 for the assigned VOT and 0.000 for the constant. The estimates shown indicate that subjects overbid in aggregate (the coefficient for *constant* is greater than 0) and that bids are biased towards the mean (the coefficient for *Assigned VOT* is less than 1). A scatter plot of every revealed VOT vs. every assigned VOT is shown Supplementary Figure B.I.I in Appendix B.I.

Subjects' bids deviated from their assigned VOT by 60 cents on average. Subjects could bid from \$0.00 to \$4.00 per second, and the range of assigned VOTs was \$0.00 to \$2.50 per second, so magnitude of the deviation is relatively high. The average bid overall was 24 cents/second higher than the average assigned VoT.

The percentage of bids that were near-truthtelling was 34%. Bids "near truthtelling" were defined as those within 20 cents/second of a subject's assigned VOT; this cutoff was chosen because it corresponds to 10% of the range of possible VOTs.

There was substantial heterogeneity in misrevelation among subjects. This is shown in Figure III.II, where the average magnitude of misrevelation for each subject was placed in a bin of width \$0.25/second.

FIGURE III.II

Frequency of Average Magnitude of Misrevelation by Subjects

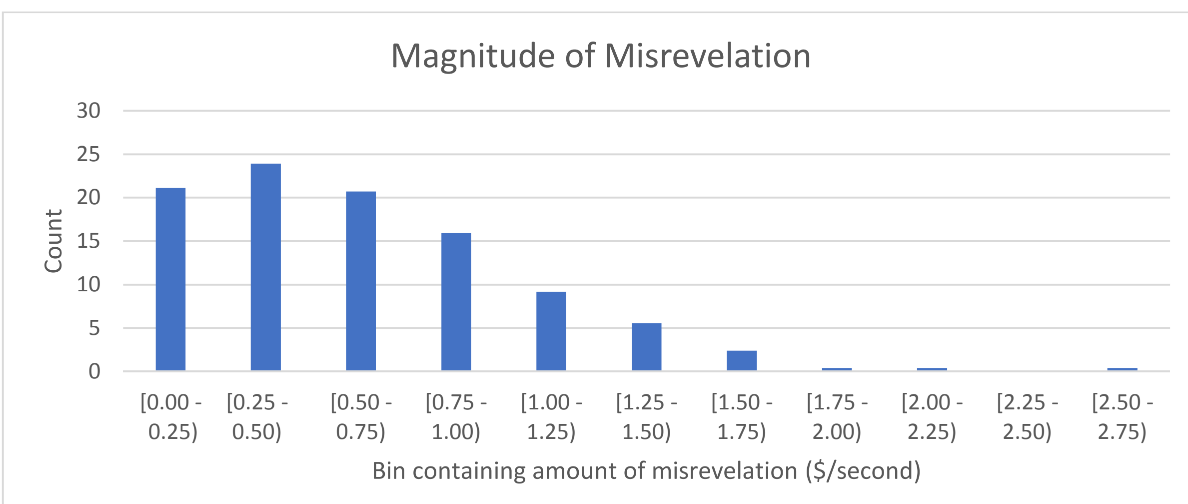

The distribution of average subject misbidding resembles a clipped normal distribution, suggesting that the ability/willingness of subjects to bid truthfully lies on a continuum. The mode is the bin representing misrevelation between 25 cents and 50 cents per second.

There was also heterogeneity among treatments in the share of near-truthful bids (within 20 cents/second of the assigned VOT, as defined previously). Depending on the experimental session, the percentage of near-truthful bids averaged over the first four driving rounds ranged from 28%-43%. The only treatment that varied during these rounds between sessions was the quality of information (pertaining to the three topics described in the methodology section) provided during the pre-trial instructions. Table III.II shows the share of near-truthful bids elicited during the first four rounds of each "assigned VOT" session, along with quality of mechanism information provided during the sessions.

TABLE III.II
Summary of Mechanism-Description Quality and Truth-Telling for each Assigned VOT Treatment

Session ID	**Quality of the pre-experiment explanation for the following topics:**			**Share of near-truthful bids. (first 4 rounds, no VOT shifts)**
	Tolls / Lane Assignment	*Truth-telling*	*VOT*	
1	Sparse	Sparse	Sparse	29%
2	**Detailed**	Sparse	Sparse	29%
3	Sparse	**Detailed**	Sparse	35%
4	**Detailed**	Sparse	Sparse	28%
5	**Detailed**	**Detailed**	**Detailed**	37%
6	**Detailed**	**Detailed**	**Detailed**	43%
7	**Detailed**	**Detailed**	**Detailed**	36%

The table shows that detailed explanations of the strategic dominance of *truth-telling* and the concept of *VOT* were associated with the highest fraction of truthful bids. The correlation between the effect of information quality for each topic and the share of near-truthful bids observed in a session was not statistically significant, however. These estimates are shown in Table III.III in the following section.

Factors affecting misrevelation

Finding 2: Mis-revelation was characterized by subjects employing a variety of coherent, yet non-dominant bidding strategies. Subjects with quantitative majors tended to bid closer to their induced VOT.

The magnitude of a subject's misbidding each round was regressed on treatment parameters and subject demographics to see how these factors affect misrevelation during the experiment. These results are shown in Table III.III.

TABLE III.III
Regression of Subject Misbidding on Experimental Treatment and Subject Demographics

Linear Regression		Number of obs = 1004	
		R-squared = 0.0629	
		(Std. Err. Adjusted for 251 cluster in ss)	
Magnitude of Misrevelation	**Coef.**	**Std. Err**	**t**
Good toll explanation	0.099	0.079	1.25
Good truth-telling explanation	-0.048	0.101	-0.47
Good VOT explanation	-0.158	0.117	-1.35
Assigned VOT	0.087*	0.047	1.84
Has a license	-0.081	0.072	-1.12
Has seen a variable message sign	-0.180**	0.084	-2.14
Subjects major is			
Somewhat Quantitative	-0.145**	0.064	-2.28
Very Quantitative	-0.104	0.073	-1.41
_constant	0.856	0.110	7.76

Notes: Standard errors were clustered at the subject level.

None of the estimated between-subjects treatment effects of information provision were statistically significant, likely because there were not enough sessions run to disentangle the three treatment variations. Estimated effects from other treatment parameters and demographic factors were significant, however.

The payment depreciation rate assigned to subjects was significantly correlated with misbidding; a \$1.00/second increase in a subject's VOT is associated with a \$.09 increase in misrevelation. A likely reason is that the lowest possible bid ($0.00/sec) is closer to the lowest assigned VOT ($0.10/sec) than the highest possible bid ($4.00/sec) is to the highest assigned VOT ($3.00/sec). Overbidding tends to be not be very costly to high VOT users (as opposed to underbidding), while underbidding tends to not be very costly to low VOT users (as opposed to overbidding). High VOT users have more "room" to overbid than low VOT users have for underbidding, which means high VOT users are able to misbid by greater magnitudes without receiving strong negative reinforcement.

Subjects who report to have encountered variable message signs on freeways had about $0.18/second less misrevelation on average than those who haven't. Possession of a driver's license is already controlled for, so this question could be a proxy for the subject driving on U.S. freeways. It is unclear why this would reduce subject misrevelation.

Lastly, subjects with a very quantitative major had about $0.15/second less misrevelation on average than subjects with a "soft" major. This is evidence that some element of human cognition plays a role in differences in subjects' understanding of the mechanism.

Next, a similar regression was run to estimate the impact of various self-reported bidding strategies on misrevelation. Table III.IV shows these regression results.

TABLE III.IV
Regression of Misrevelation Magnitude on Self-reported Non-dominant Bidding Strategies

Linear Regression		Number of obs = 251 R-squared = 0.1357	
Magnitude of Misrevelation	**Coef.**	**Std. Err**	**t**
N/A (incoherent response) (count = 37)	0.231***	0.077	3.01
Always avoided toll lane (count = 30)	0.319***	0.077	4.15
Always sought toll lane (count = 6)	0.022	0.139	0.16
Used trial and error (count = 37)	0.292***	0.075	3.88
Chose a constant or middle VOT (count = 12)	0.272*	0.153	1.77
Randomly chose bids (count = 28)	0.447***	0.125	3.57
Bid based on starting endowment (count = 9)	0.175	0.110	1.59
Bid a multiple of payment dep. Rate (count = 9)	-0.014	0.094	-0.15
_constant	0.428	0.039	11.10

Notes: The strategies were provided as free-response answers and then later categorized for analysis. The self-reported strategy chosen as the base level is the dominant strategy of subjects selecting their payment depreciation rate as their VOT.

There were many self-reported alternative bidding strategies significantly correlated with subject misbidding. These included toll avoidance, trial and error, picking a constant or middle VOT, or random bidding. Of these, subjects who chose their bids at random also misbid the greatest amount on average. While not a strategy itself, subjects who were unable to coherently articulate their strategy also misbid significantly more. These results show that misrevelation is often associated with subjects' conscious decision to employ specific and coherent (albeit non-dominant) bidding strategy, rather than confused/random behavior.

Learning

Finding 3: Subjects did not learn to bid closer to their true (induced) VOT over repeated trials, but did learn to place bids that reduce their likelihood of being assigned to utility-lowering lanes.

Under the baseline scenario of the mechanism where there are no exogenous changes to payment depreciation rates or traffic flow rates, subjects did not learn to provide their true VoT. Figure III.III shows the trend in subject misrevelation averaged over the first four rounds (post-practice) across all treatments.

FIGURE III.III

Magnitude of Misbidding vs. Round-number Averaged Across the First 4 Rounds of all Sessions

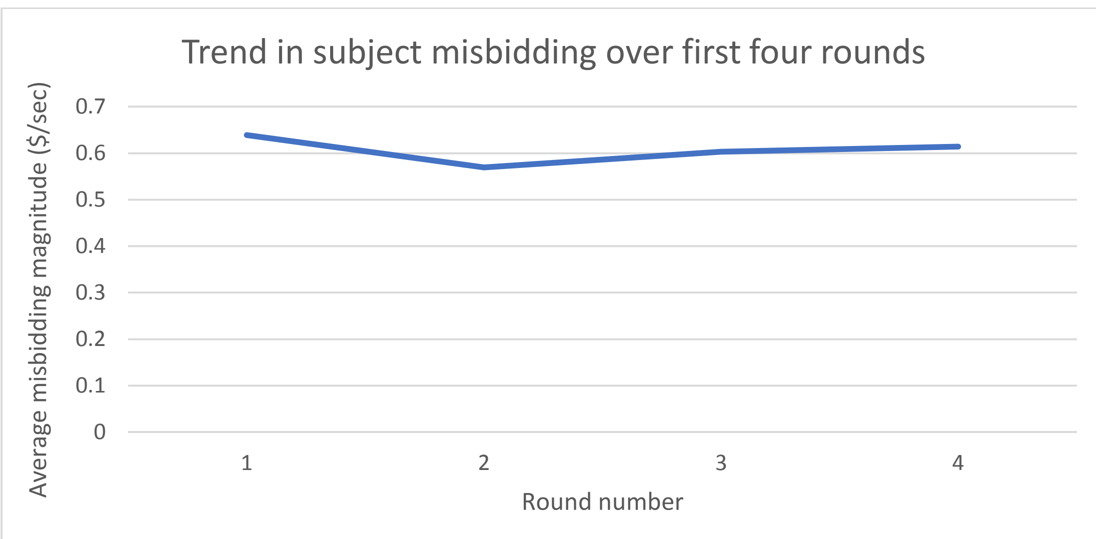

Notes: This graph is restricted to the first four rounds, because treatment variation involving changes to depreciation rates and traffic flow rates are introduced in some sessions beginning with round five. A graph showing subject learning over 13 sessions is shown in Supplementary Figure B.I.II in Appendix B.I. The magnitude of misrevelation does not decrease over time for that sample either.

The trend in misrevelation is essentially flat over the first four rounds. One likely contributing factor is that subjects do not always receive sufficient reinforcement to discourage misrevelation. Figure III.IV shows the relationship between subjects' misrevelation and earnings.

FIGURE III.IV

Relationship Between Misbidding and Earnings, Controlling for Assigned VoT

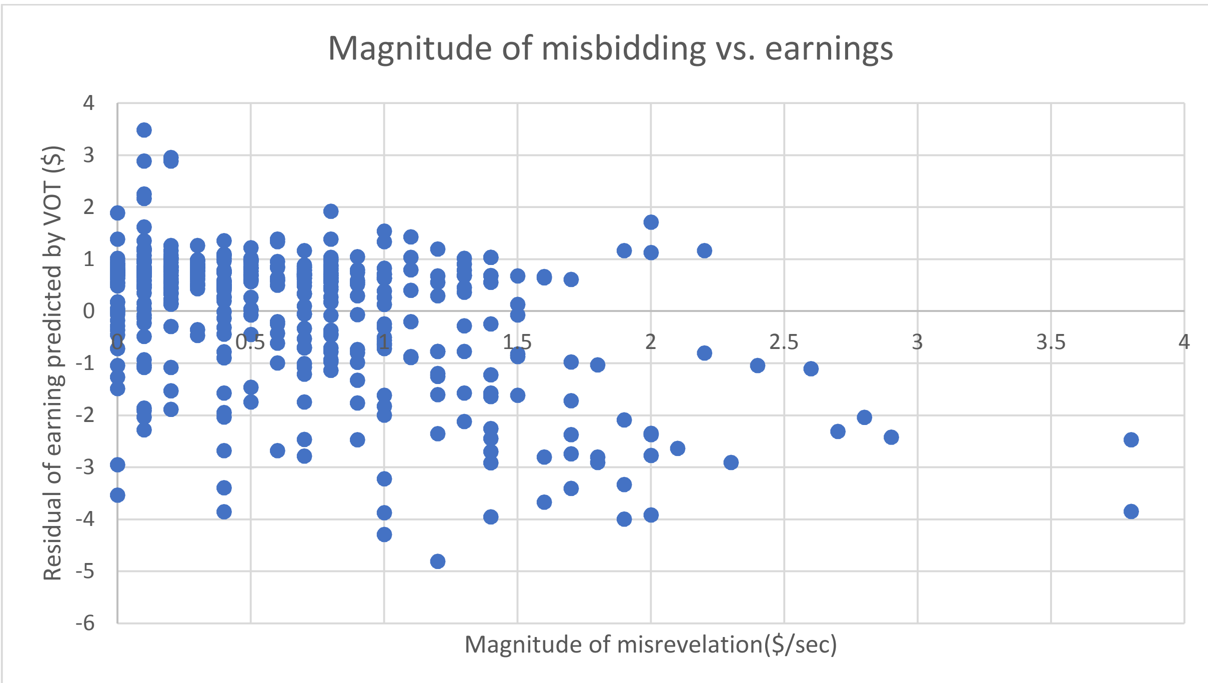

Notes: The graph controls for a subject's assigned VOT, because it has a strong effect on earnings. Subjects assigned higher VOTs earn more on average due to miscalibration that resulted in starting endowments being set as a multiple of depreciation rate that was slightly too high.

It is clear to see that increased misrevelation does not always result in a worse payout. There are a few reasons for this. One substantial source of payout variation comes from variable traffic dynamics. When drivers attempt to switch lanes to take the one they are assigned to, they can impede one another and cause temporary congestion delays. This stochastic process creates a noisy signal that causes random shocks to a subject's payout. An example of how this could impede the learning of truth-telling would be if a subject misbid by a smaller amount in a particular round but received a lower payout due to extra lane-changing congestion. That subject would erroneously conclude that bidding more truthfully lowered his/her earnings. It is possible that this effect is exaggerated by the experimental parameters, because driving times once the lanes are separated are relatively short relative to time spent changing lanes to get into one's assigned lane. More analysis needs to be performed to assess whether this issue could be as substantial in the real world.

Another important reason for the failure of subjects to learn truthful revelation is a lack of sufficient reinforcement. Misrevelation did not significantly impact users' earnings unless it caused them to be assigned to a different lane than the one they would've taken had they provided their truthful VOT. For example, if a subject with a high VoT - who would be ordinarily assigned to the toll lane - bids even higher, the allocation will only infrequently be changed (with one less subject being assigned to the toll lane). The converse is true for those with low VoTs (if one doesn't bid high enough to take the toll lane, the bid will only infrequently result in one less subject being put on the free lane). As a result, many non-truthful bids produce the same or near-same outcome for subjects as truthful revelation. Figure III.V shows how average earnings varied with the bids of two groups of subjects: those whose payment depreciation was either 60 or 70 cents per second, and those whose payment depreciation rate was 1.90 or 2.00 dollars per second.

FIGURE III.V

Average Payout Versus Bid for Subjects with a VOT of 60 or 70 Cents per Second and Subjects with a VOT of 1.90 or 2.00 Dollars per Second.

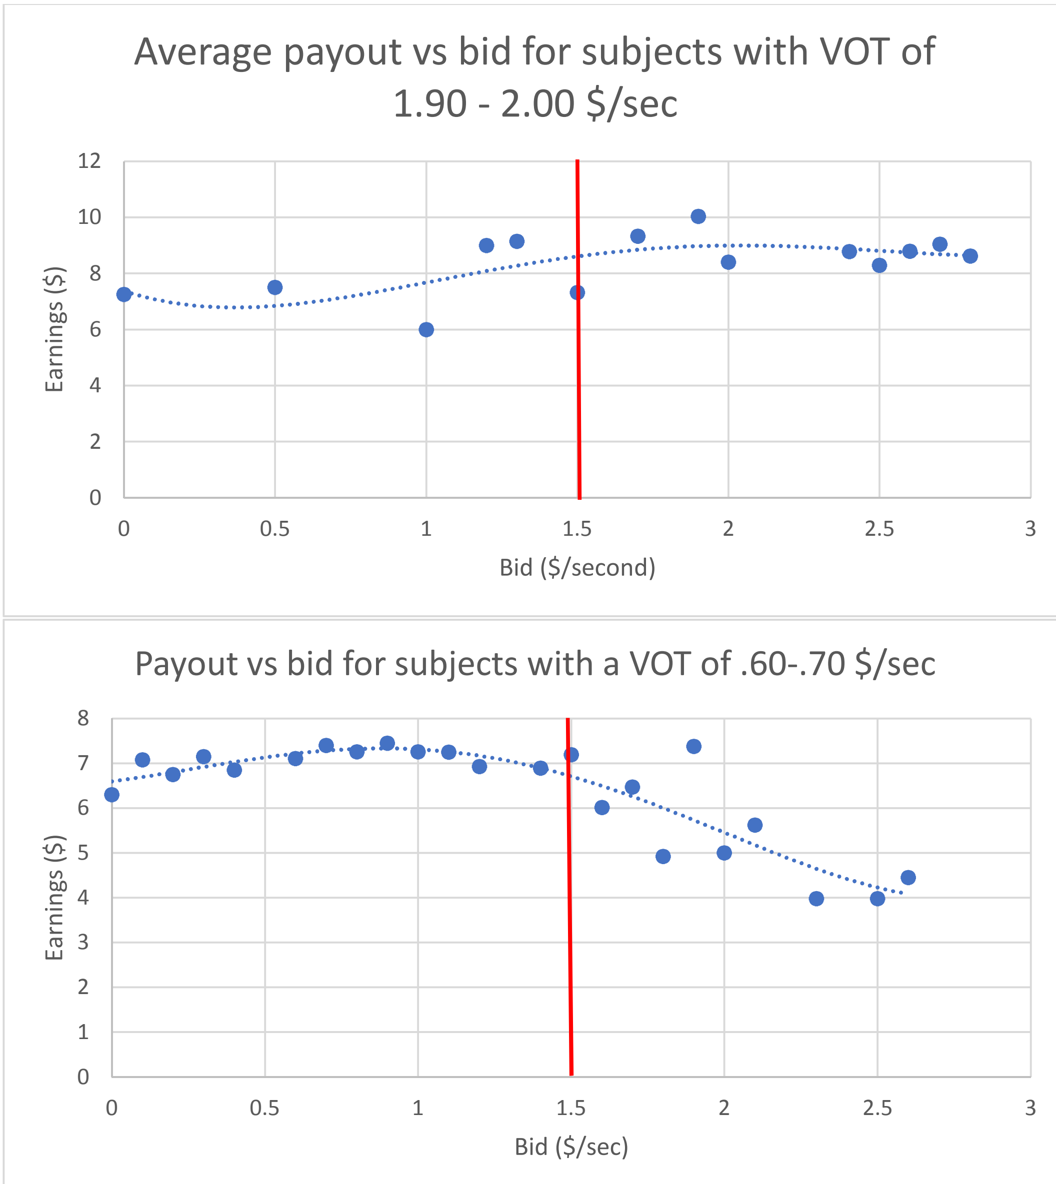

The graphs show that misbidding lowers earnings only slightly unless it causes subjects to be assigned to a different lane. Typically, the threshold for being assigned to the toll-lane was close to $1.50/second, so the two groups of subjects in Figure III.V represent payment depreciation rates on either side of the threshold. The graph suggests that subjects who don't understand the mechanism would not receive sufficient reinforcement to myopically learn that truth-telling is a dominant strategy; instead they would only learn that very high bids (for the low VOT group) or very low bids (for the high VOT group) strongly reduce their earnings. This is an attribute of freeway lane allocation that makes the use of a VCG less feasible in practice.

A learning model was estimated to conclusively demonstrate the effects of travel time stochasticity and weak reinforcement on the ability of subjects to learn turhtful revelation. The dependent variable of this model is the magnitude of the misrevelation of a subject's bid. The magnitude of the prior round's misrevelation is also included, so that the model estimates the effect of covariates on the **change** in the magnitude of a subject's misrevelation from round to round. In other words, this model shows whether certain factors cause subjects to bid further or closer to true their VoT, and by how much.

The model below describes a myopic learning model, where the magnitude of a subject's misbidding at time *t* is a function of inertia from misbidding at time *t-1*, feedback from changes in misbidding between periods *t-2* and *t-1,* feedback when changes in misbidding result in being assigned to a different lane, and feedback when changes in misbidding coincide with any random travel time shocks that experienced by a subject. Controls for assigned VOT, round, and experimental session are included as well. The model is written as:

$$
\begin{aligned}
Misbidding_{i,t} = {} & \alpha * Misbidding_{i,t-1} + \beta_1 * \left(Misbidding_{i,t-1} - Misbidding_{i,t-2}\right) + \\
& \beta_2 * LaneSwitched_{i,t-1} * \left(Misbidding_{i,t-1} - Misbidding_{i,t-2}\right) + \\
& \beta_3 * TimeShock_{i,t-1} * \left(Misbidding_{i,t-1} - Misbidding_{i,t-2}\right) + \\
& \delta_1 * AssignedVOT_i + \delta_2 * Session + \delta_3 * Round_t + constant + \varepsilon_{i,t}
\end{aligned}
$$

The dependent variable and covariates are described in more detail on the following page:

$\boldsymbol{Misbidding_{i,t}}$ – the difference between the revealed and assigned VOT provided by subject *i* at time *t*.

$\boldsymbol{Mischange_{i,t-1}}$ – Defined as $(Misbidding_{i,t-1} - Misbidding_{i,t-2})$, the change in misbidding between periods *t-2* and *t-1*. Any change in misbidding provides useful feedback for subjects, so intuitively the estimated coefficient should be negative.

$\boldsymbol{LaneSwitched_{i,t-1}}$ – Indicates that subject *i* was assigned to a different lane at time *t-1* than time *t-2*. When interacted with $\boldsymbol{Mischange_{i,t-1}}$, the covariate is called $\boldsymbol{MischangeLane_{i,t-1}}$. Any change in lane assignment gives useful feedback, so intuitively the estimated coefficient should be negative.

$\boldsymbol{TimeShock_{i,t-1}}$ – The difference in the stochastic time delay experienced by subject *i* between time *t-1* and time *t-2*. When interacted with $\boldsymbol{Mischange_{i,t-1}}$, the covariate is called $\boldsymbol{MischangeShock_{i,t-1}}$. Positive delays amplify feedback from increased misrevelation feedback, and negative delays amplify feedback from decreased misrevelation. Intuitively, the estimated coefficient should be negative.

This model is based on reinforcement learning models, where agents myopically adjust their behavior based on feedback. The model describes how subjects who do not grasp the dominant strategy of truthful revelation might still reach this outcome through myopic iterative best-responses. The feedback for subjects in this experiment comes from changes in their payout. This model separates subjects' payouts into several different covariates to isolate their specific contributions to learning. For example, if a subject bids further away from their true VOT than they did one round earlier, their payoff should decrease, which provides negative reinforcement for that deviation. One would expect misrevelation during the next round to decrease in response. This specific effect is captured by the covariate $\boldsymbol{Mischange_{i,t-1}}$.

The effect on a subject's payout from changes in misrevelation is amplified by a subject being switched to a different lane. Thus, the resulting adjustment that a subject makes to his/her misbidding should be larger in response to the stronger reinforcement. This effect is captured by the inclusion of

$\boldsymbol{MischangeLane_{i,t-1}}$. Lastly a time shock due to random lane-switching dynamics can either amplify or dampen the perceived effect on a subject's payout due to misbidding. If a subject increases misrevelation but benefits from travel time stochasticity, the apparent consequences will be reduced. The reverse is true, however, if the subject would have incurred extra delays due to travel time stochasticity. The role that this factor plays in reinforcement learning is captured by the variable $\boldsymbol{MischangeShock_{i,t-1}}$.

Observations from any experimental round that occurs after subject VOTs were reassigned were discarded, due to their confounding effects on misrevelation. Three of the seven experimental sessions where subjects were assigned VOTs were also omitted, because VOTs were reassigned after only four rounds. These would have only produced two usable rounds of data due to autoregression in the model. The remaining four sessions used for the estimation contain six rounds of data before VOTs were re-assigned. The estimation results are shown in Table III.V.

TABLE III.V
Regression of Subject Misbidding on Prior-round Outcomes

Linear Regression		Number of obs = 538 R-squared = 0.3412	
Misrevelation *Misbidding(t)*	**Coef.**	**Std. Err**	**t**
Prior round misrevelation *Misbidding(t-1)*	0.551***	0.038	14.67
Prior round change in misrevelation *Mischange(t-1)*	0.121	0.088	1.38
Prior round change in misrevelation; subject also assigned to different lane *MischangeLane(t-1)*	-0.285***	0.086	-3.31
Prior round change in misrevelation; interacted with lane delay shocks. *MischangeShock(t-1)*	-0.109***	0.033	-3.31
Payment depreciation rate *Assigned VOT*	0.060**	0.029	2.07
SessionID			
6	-0.009	0.049	-0.19
7	-0.055	0.050	-1.11
19	-0.032	0.052	-0.62
Round Number	0.017	0.016	1.07
_constant	0.172	0.089	1.94

The coefficient for **Mischange(t-1)** is not statistically significant and is a different sign than expected. This suggests that bidding closer or further to one's VOT does not provide subjects with sufficient reinforcement to learn truthtelling. However, the statistically significant negative estimate for **MischangeLane(t-1)** demonstrates that when a subject increases or decreases the magnitude of misbidding and is also switched to a different lane, the reinforcement is sufficient for that subject to reduce the magnitude of misbidding during the next round. These findings provide strong evidence that subjects only learn to decrease misrevelation when changes in their bidding result in being assigned to a different lane. Thus, there is a wide range of possible misrevelation that subjects will fail to learn to correct.

The model results suggest that subjects do not learn to explicitly reduce misrevelation, but instead learn to place bids that result in being assigned to utility-lowering lanes less frequently. Further evidence of this is provided in the Figure III.VI, which shows the average number of very low VOT subjects assigned to toll lane each round.

FIGURE III.VI

Average Number of Low-VOT (0.50 $/sec or less) Users Assigned to the Toll Lane each Round

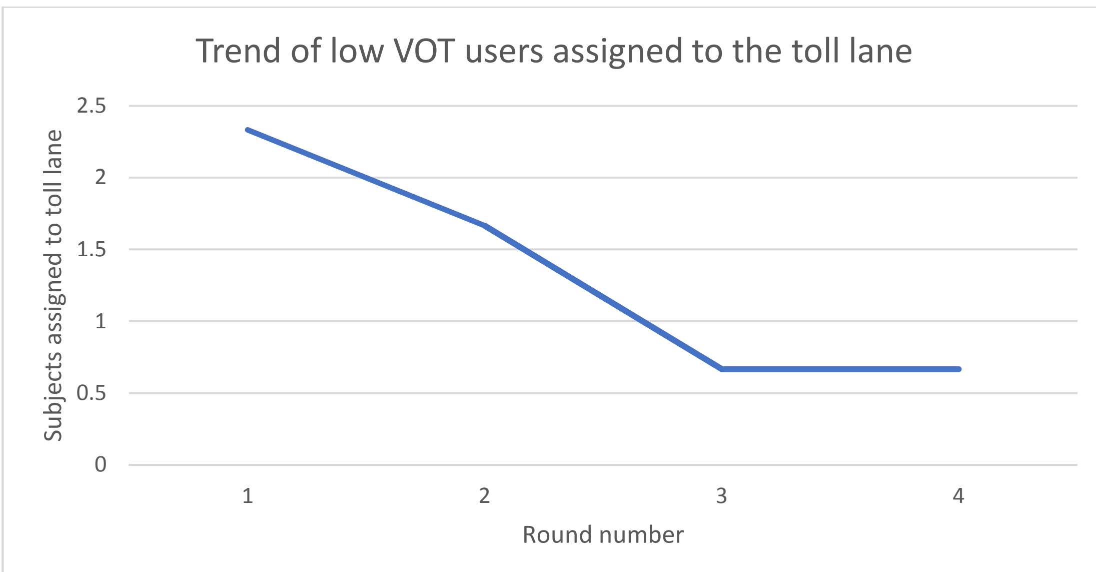

The graph shows that low-VOT users (the group for whom toll-lane assignment is costliest) learn to place bids low enough to not be assigned the toll lane. It seems paradoxical that costly lane mis-allocation can decrease while misrevelation overall does not. However, it is the case that misrevelation biased towards extreme bids increases, while misrevelation biased towards central bids decreases. The former occurrence has a minimal effect on lane allocation, while the latter reduces costly lane misallocation.

The statistically significant coefficient for **MischangeShock(t-1)** also indicates that travel time stochasticity affects learning pertaining to misrevelation. The time shocks either amplify or dampen reinforcement, depending on whether their direction coincides with the direction of changes in misbidding (though the variable was constructed so that any intuitive effect would be identified in the same negative direction). Although the estimate for this covariate provides strong evidence that travel time variation impacts learning, it is not apparent from the estimate whether the effect is beneficial or detrimental overall. To determine whether the presence of travel time stochasticity increases or decreases misbidding overall, each subject's average magnitude of misbidding over all rounds was regressed on that the total average magnitude of random delay shocks that each subject experienced. Identifying variation comes from the fact that each different subject experienced a different level of travel time stochasticity depending on their spatial positioning during each round. The regression also includes controls for assigned VOT and experimental session. The results, shown in Table III.VI, indicate that travel time shocks are significantly and positively correlated with misrevelation. For each second of travel time shocks experienced by a subject, they provide VOTs that are almost 5 additional cents per second further away on average from their true VOT.

TABLE III.VI
Regression of Subjects' Average Magnitude of Misrevelation on the Average Magnitude of Delay Shocks Experienced by Subjects.

Linear Regression		Number of obs = 137 R-squared = 0.0881	
Average Misrevelation (magnitude)	**Coef.**	**Std. Err**	**T**
Average Delay Shock (magnitude)	0.047*	0.027	1.73
Assigned VOT	0.522**	0.217	2.41
SessionID			
6	-0.253	0.369	-0.69
7	-0.638*	0.374	-1.71
19	-0.385	0.398	-0.97
_constant	2.057	0.328	6.28

Efficiency

Finding 4: The VCG mechanism achieved a more efficient lane allocation for drivers than all other studied alternatives. Furthermore, driver welfare improved over repeated trials.

Misrevelation results in efficiency losses from the mechanism, because the social welfare maximization will be performed based on inaccurate preferences. In every round of every experimental session, misrevelation resulted in a suboptimal allocation of users between the toll and free lanes. In every case, there were high VOT subjects assigned to the free lanes due to underbidding, and vice versa. In many cases, either too many or too few subjects were assigned to toll lanes as well. The calculation of welfare loss resulting from misallocation was simplified by assuming that the only relevant welfare measure was travel delay cost, computed as the sum of each driver's VOT multiplied by the amount of travel time spent reaching their destination. Another way welfare could be measured is to also consider toll expenditures. Misbidding and overpaying for time savings as a result has welfare implications, as does paying any toll at all. These welfare implications greatly depend on how toll revenue is redistributed, which is beyond the scope of this work. Future work studying how toll payments affect the welfare analysis of this mechanism would be beneficial.

System welfare was assessed in aggregate when computing mechanism efficiency loss from misrevelation, because a given user's misrevelation imposes externalities on other users. Take for example the case of an extremely high bid from a low VOT user that changes the allocation so that fewer users are assigned to the toll lane. This would impose negative externalities in two ways and a positive externality in one way. Negative externalities would result from marginal high VOT users being displaced to the free lanes and incurring higher delay costs, as well as from the extra users assigned to free lanes slowing down free lane users. A positive externality would come from increasing toll lane speed resulting from fewer toll lanes assignees. In this case, the negative externalities would outweigh the positive externality and lower system welfare. System welfare loss was therefore assessed by comparing the minimum aggregate travel delay cost possible with no misrevelation to the realized aggregate delay cost achieved by subjects based on their actual bids. System welfare loss is the difference between these two quantities, and efficiency loss is the quotient of this difference and the hypothetical minimum delay cost. Table III.VII shows the welfare loss associated with the VCG mechanism due to subjects misbidding.

TABLE III.VII
Mechanism Efficiency Loss by Treatment Type

Session ID	Quality of the pre-experiment explanation for the following topics:			Welfare loss as % of optimal (first 4 rounds, no VOT shifts)
	Tolls / Lane Assignment	Truth-telling	VOT	
2	Sparse	Sparse	Sparse	1.8 %
3	Detailed	Sparse	Sparse	2.4 %
4	Sparse	Detailed	Sparse	2.3 %
5	Detailed	Sparse	Sparse	2.4 %
6	Detailed	Detailed	Detailed	1.5 %
7	Detailed	Detailed	Detailed	1.3 %
19	Detailed	Detailed	Detailed	1.5 %

Depending on the treatment considered, subjects providing incorrect VOT increased the delay costs of drivers on the road anywhere from 1.3 % to 2.4 %. As is the case with misbidding, there is at least a weak correlation between the pre-experiment information provided to subjects and efficiency losses. Losses were lowest when detailed information about the mechanism, the importance of bidding truthfully, and the concept of VOT were provided.

These efficiency losses are put in better context by comparing them to travel delay costs in the absence of the VCG mechanism. With no pricing mechanism at all, travel delay costs increase by 3% provided that a user equilibrium is achieved where the speed across both sets of lanes is identical. This means that under the baseline scenario with not VOT or road capacity changes, the mechanism never performed so poorly that drivers would be better off with no pricing at all.

The VCG mechanism's performance can be put into further perspective by comparing it to other pricing mechanisms. Many scenarios were considered for comparison. One such scenario is a dynamic pricing regime where the toll lane is kept at a minimum speed that is 50% faster than the free lanes. Other scenarios included a second best tolling regime where planners knew the distribution of driver VOT to varying degrees of accuracy.

To calculate welfare losses under the minimum speed toll lane regime, an analytical model based on experimental parameters was used to solve for the equilibrium price and then compute driver welfare based on delay costs. For second-best tolling scenarios the same model was used, however VOTs of simulated freeway users were drawn at random from the underlying distribution used for the experiment, which was a uniform distribution ranging from \$0.10 to \$2.40 per second. For each these scenarios, the model assumes that drivers always make rational lane choices based on whether the toll is worth the time savings they will get in the toll lane.

Three second-best tolling scenarios were considered. In the first scenario, planners perfectly knew the underlying distribution of VOTs for the population of drivers, but did not know what the realized VOT sample would be at any given moment. The welfare-maximizing second-best toll was set based on the parameters of the underlying distribution, and a VOT for each of the simulated 39 drivers was drawn from that same distribution. If the sample produces a set of subject VOTs that is representative of the underlying distribution, then the ex-ante toll set by planners ends up being optimal, and there is no efficiency loss. If the sample of VOTs is non-representative of the underlying distribution, however, then the ex-ante toll will be suboptimal; a better toll could have been set had operators known the true values of VOTs in the sample. To determine the efficiency loss that results from this sampling error, the system welfare given the equilibrium distribution resulting from the ex-ante toll was compared to the system welfare resulting from the optimal toll that would have been set had the realized sample of VOTs been perfectly known. The sampling process was repeated 50 times, and the average of the corresponding efficiency losses was taken. It is worth noting that the efficiency losses calculated for this scenario are sensitive to the number of drivers whose travel times are affected by any one driver's lane assignment. This analysis assumes that each vehicle's assignment affects 38 others, simply because the experimental laboratory accommodated a maximum of 39 subjects per session. The higher the number of interdependent vehicles is, the lower efficiency losses from sampling will be by virtue of the increased sample size better matching the population. The actual amount of real-world travel time interdependency is determined by complex traffic dynamics that are beyond the scope of this project.

The remaining two second-best tolling scenarios assume that toll operators know the correct shape of the underlying VOT distribution but misestimated the population mean. In one scenario the true mean of user VOTs differs by 10% from what is estimated, and in the other the true mean differs by 25%. In these scenarios, the true mean is higher half the time and lower half the time. The welfare calculation otherwise followed the exact same procedure as the first second-best tolling scenario, the only difference being that the sampled VOTs were drawn from distributions with different means. Efficiency losses were higher in these scenarios, since it was almost guaranteed that a set of drivers' VOT will not match the expected value of the incorrect distribution. Obviously the further off the sample mean was from the true mean, the higher the resulting efficiency loss would be. If planners assume that the mean of the VOT distribution is lower than it really is, than the toll will likely be set too low – and vice versa. These two scenarios are not as sensitive to the assumptions made in this work regarding the interdependency of travel-times for vehicles sharing the road.

For each of the second-best tolling scenarios, a constraint was imposed that prevented the toll lane from having slower speeds than the free lanes. This constraint reflects the belief that even if the toll was set far too low, drivers would stop choosing the toll lane based on the (incorrectly) predicted travel times once it became clear that the toll lane was moving slower than the free lanes. However, toll lanes are occasionally observed to move to slower than adjacent free lanes, because drivers have limited opportunities to switch between the two sets of lanes. A second set of efficiency loss estimates were computed with this assumption relaxed so that toll lane speeds may be up to 12% slower than free-lane speeds. Table III.VIII shows the increase in welfare loss associated with various second-best tolling scenarios.

TABLE III.VIII
Second-best Tolling Efficiency Losses by Simulation Scenario

Misestimation of average VOT	Toll lane speed constrained (if no, toll lane can be 12% slower)	Welfare loss as a % of optimal
0%	Yes	1.4%
0%	No	1.7%
10%	Yes	2.0%
10%	No	2.3%
25%	Yes	4.0%
25%	No	5.2%

When the toll lane is constrained to be equal or faster to the free lane, efficiency loss from standard tolls ranges from 1.4% to 4.0%, depending on how well the distribution was known. This range increases to 1.7% to 5.2% when this constraint is relaxed by 12%.

The simulated welfare losses from second-best tolling assume that drivers always choose the lane that maximizes their utility given their VOT and estimated toll lane time savings. If the same assumption of driver rationality was applied to the welfare analysis for a VCG mechanism, welfare losses would be near-zero. Therefore, it is instructive to also compare actual human performance under a second-best tolling regime to welfare losses under a VCG in practice. To this end, an additional human-subjects experiment treatment was run that simulated a second-best toll. An optimal toll was set based on the VOT assigned to subjects, and the subjects chose whether to take the toll lane based on information provided to them about the toll lane price and estimated time savings. This simulates the decision environment drivers face in the real-world, since they must determine for themselves whether the toll lane is worth the price based on the amount of time it is expected to save them. Subjects in the experiment frequently chose the lane that gave them lower earnings, despite having the necessary information available to determine the optimal lane choice. These sub-optimal choices reduced aggregate driver welfare by 3%.

A ranking of each various lane pricing scenarios by their performance with respect to social welfare is shown in Table III.IX.

TABLE III.IX

Welfare loss under each pricing scenario in ascending order

Pricing regime	**Simulated vs. Experimental**	**Scenario details**	**Efficiency Loss**
VCG	Experimental	Best instructions provided	1.3%
Second best	Simulated	Correct VOT distribution known, Toll lane speed >= free lane speed	1.4%
Second best	Simulated	Correct VOT distribution known, Toll lane speed >= 0.88 * free lane speed	1.7%
Second best	Simulated	Mean of VOT dist. misestimated by 10%, Toll lane speed >= free lane speed	2.0%
Second best	Simulated	Mean of VOT dist. misestimated by 10%, Toll lane speed >= 0.88 * free lane speed	2.3%
VCG	Experimental	Worst instructions provided	2.4%
Second best	Experimental	Exact set of VOTs of known Subjects determine which lane is optimal	3.0%
Second best	Simulated	Mean of VOT dist. misestimated by 25%, Toll lane speed >= free lane speed	4.0%
Free-flow	Simulated	Toll lane speed >= 1.50 free lane speed	4.4%
Second best	Simulated	Mean of VOT dist. misestimated by 25%, Toll lane speed >= 0.88 * free lane speed	5.2%

The table shows that under the most favorable set of experimental conditions, the VCG mechanism achieved the most efficient lane allocation of all other alternatives. Under the worst conditions, however, the mechanism only outperformed a simulated second-best toll when the mean of the VOT distribution was misestimated by 25%. In experimental practice, however, the second-best toll performed consistently worse than all scenarios where the VCG mechanism used – even when the optimal second-best toll was set. The second-best toll produced experimental welfare losses because human subjects were not always able to calculate the utility-maximizing lane based on toll and time-savings information provided to them. The VCG does not incur welfare losses in this fashion because drivers are automatically assigned to the utility-maximizing lane, provided that they specified their travel time preferences truthfully. The tolling scenario where the toll lane was kept in free-flow produced greater welfare losses nearly all other alternatives.

The welfare results should be interpreted cautiously, because they are derived from an environment where value of time is the only relevant travel time preference. The results cannot be extended to an environment where other preferences such as reliability and urgency are relevant. Nevertheless, the findings from this work demonstrate that in a simplified environment with one relevant travel-time preference, the VCG mechanism can improve welfare over alternative pricing schemes.

Furthermore, the welfare loss under VCG pricing decreased over time – despite the fact that misrevelation did not. This results from subjects reducing the frequency of getting assigned to sub-optimal lanes given their VOT. Figure III.VII shows a downward trend in average welfare loss (meaning welfare is increasing) across all sessions for the first four rounds.

FIGURE III.VII

Average Welfare Loss with a VCG Mechanism for the First Four Treatment Rounds

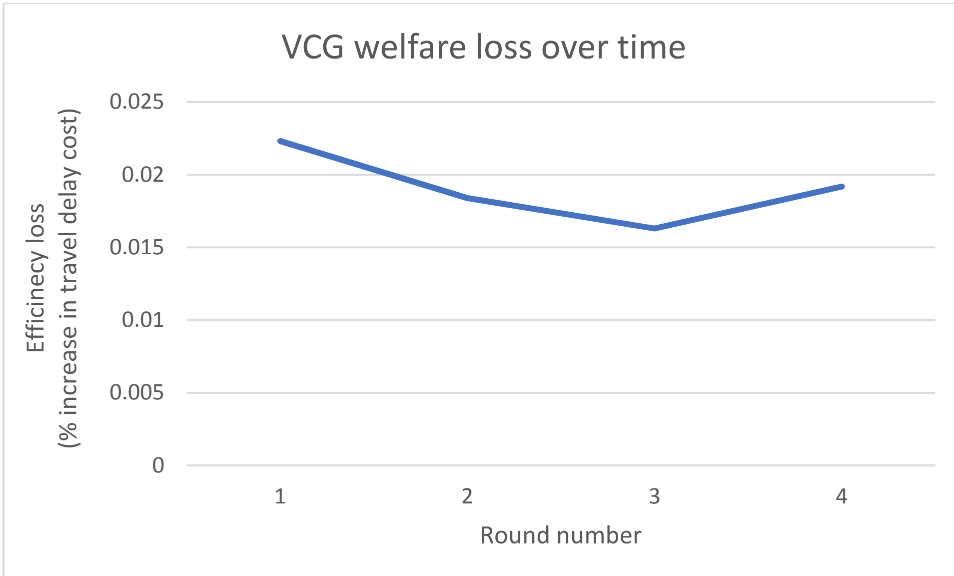

The downward trend in welfare loss is even greater in treatments where travel time stochasticity was intentionally reduced through design of the environment. Figure III.VIII shows welfare loss over time for scenarios where travel-times were made more deterministic through ease of lane-switching.

FIGURE III.VIII

Average VCG Welfare Loss for the First Four Rounds of Treatments with Less Travel-time Stochasticity

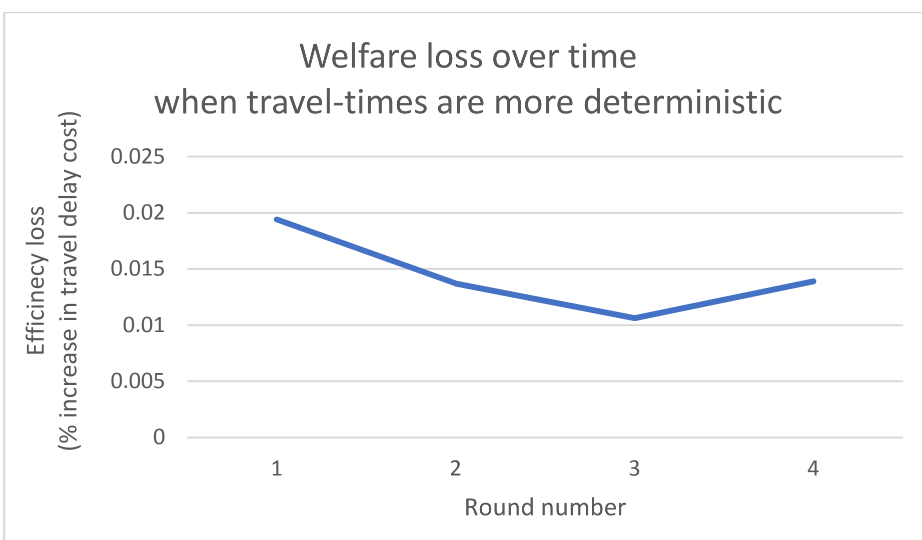

The faster rate of welfare improvement shown in this graph is evidence that learning is easier for subjects when more deterministic travel times strengthen the correlation between bidding truthfulness and driver utility. These results also suggest that the more experience that a population of drivers has with the mechanism, the more of the mechanism's theoretical benefits can be achieved.

Robustness check using "Innate" VOT treatment: Plausibility

Finding 5: The "innate" VOT incentive scheme elicited plausible bids from subjects.

In sessions where subjects were incentivized by leaving the experiment sooner (and thus their value of time based on real-world opportunity costs) rather than the subtractive payment scheme, the session-average VOT reported by subjects ranged from 8-12 dollars per hour. These values are understandably lower than those found from empirical studies of SR-91 in Orange County, which found that the average VOT of drivers was $12-13/hour using a stated preference methodology and $21/hour using a revealed preference methodology. The SR-91 study samples a representative group of actual toll-road users, while the sample for the experiment in this study was comprised exclusively of students attending University of California, Irvine. By comparison, these students likely have lower incomes on average, and more importantly, were participating in experimental sessions that typically pay $15/hour. Therefore, it is highly plausible that the true average VOT of the experimental subjects was below both that SR-91 users and the pay rate from ESSL experiments, which is consistent with the findings of this work.

Subjects' reported VOTs were regressed on demographic, attitudinal, and schedule related survey responses as a further check on their plausibility. These regression results are shown in Table III.X.

TABLE III.X
Regression of Elicited VOT on Various Survey Responses

Linear Regression		Number of obs = 166 R-squared = 0.2747	
Average VOT of a subject	**Coef.**	**Std. Err**	**t**
Gender	-1.632*	0.987	-1.65
Has a license	2.410*	1.278	1.89
Older than 23	5.005*	2.800	1.79
Prior toll history			
Has taken a toll lane	2.702*	1.390	1.94
Never seen toll lane	1.097	1.174	0.93
Boredom affected bidding	3.843***	0.978	3.93
Subject's post-experiment schedule			
Class or office hours	3.578**	1.386	2.58
Deadline to meet	9.136**	4.410	2.07
_constant	5.085	2.554	1.99

Notes: Experimental-session fixed effects were also included but not shown

The signs of most estimated coefficients match intuition or prior literature. A description of each variable and its estimated coefficient is provided on the following page.

Gender: Being male was significantly correlated with a VOT that was $1.63/hour lower on average. This is consistent with the finding from Parkany (1999) that female drivers are more likely to use toll lanes.

License: Subjects with a driver's license reported significantly higher VOTs, over $2/hour more. This matches intuition, since subjects with a license are more mobile and likely have a higher opportunity cost of experimental participation

Older than 23: Subjects older than the age of a typical undergrad reported significantly higher VOTs, over $5/hour higher. This matches the Parkany (1999) finding that age increased the probability of choosing the toll lane, and also matches intuition that older subjects would be more likely to have more responsibilities or a higher paying job.

Toll History: The baseline for this covariate is subjects who have encountered toll lanes before, but have never chosen to take one. Subjects who have taken toll lanes before has significantly higher VOTs on average by comparison. This agrees with intuition because subjects who have taken toll lanes before likely have a higher value of time.

Schedule: A subject's schedule had strong and significant effects on the VOT he or she chose. The baseline was subjects whose responses indicated that they were committed to staying on the premises of the experimental lab for the entire scheduled duration (e.g. had to wait on a ride home scheduled for the end of the time-period designated for the experiment). The baseline should be associated with a very low VOT, since there is very little opportunity cost. By comparison, subjects who didn't express this sentiment and said they had other matters on campus to attend to had an average VOT that was over $3/hour higher. Subjects who mentioned a deadline or something else important to do later that day had a VOT that was over $9/hour higher on average. These findings match intuition; higher opportunity costs resulted in higher stated VOTs.

Boredom: Subjects who reported being bored during the session stated values of time that were significantly higher, over $3/hour more on average. This finding also matches intuition.

Despite abundant evidence that VOTs elicited by the mechanism were plausible, there were several individual subjects that made seemingly irrational bids. Figure III.IX shows the frequency of every stated VOT, binned into groups of two, made during innate VOT treatments.

FIGURE III.IX

Complete Distribution of each VOT Elicited during "Innate VOT" Treatments

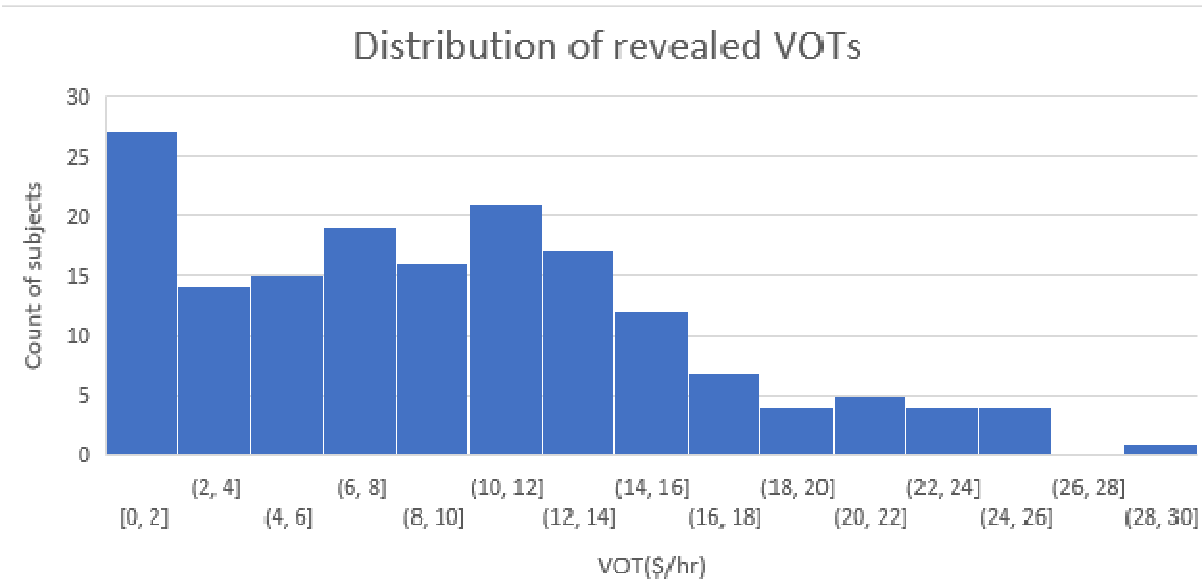

Many subjects bid well over $20/hour on average, which does not make economic sense given that they participate in experiments that pay $15/hour on average, and thus implicitly value their time at less than that rate. There are several reasons why seemingly irrationally bids may occur. Subjects may not understand the mechanism or their own value of time, which would reflect cognitive limitations. However, subjects may also derive such disutility from participating in this particular experiment, that they are willing to pay more than their typical value of time to leave. Given the strong correlation between boredom and VOT, this latter explanation is likely a contributing factor.

Robustness check using Innate VOT treatments: Volatility

Finding 6: The average volatility of an individual subjects' bidding across trials was identical between assigned VOT and innate VOT sessions. This is evidence that assigning subjects a VOT in the form a payment depreciation rate is a suitable treatment for studying mechanism performance.

The experiment is a controlled setting where one would not expect a subject's VOT to change much as it progressed. Subjects were not allowed to access their phones or the internet during the experiment, so there was no chance for a subject to become aware of an urgent event. Therefore, one would expect minimal fluctuation in the VOTs elicited by a subject if the mechanism was working as designed. This would indicate that subjects quickly grasp the mechanism and can immediately identify their own VOT. A subject whose reported VOT fluctuates significantly either does not understand the mechanism or does not have a strong sense of their own VOT, and in either case is experimenting with various values. A subject cannot be consistently bidding their true VOT if the value they give is fluctuating, and thus the fluctuations are indicative of welfare loss.

Examining evidence of misbidding in innate VOT treatments is an important robustness check on the findings from assigned VOT treatments. This is because assigned VOT treatments introduce an unnatural element to a subject's preferences; travelers are not typically incentivized by literal linear payment depreciation in their day to day lives. Therefore, the payment depreciation rate itself could be a source of subject confusion, in addition to the mechanism, and could be responsible for some or most of the misbidding observed.

If the volatility in bids is similar between the two types of treatments, however, it lends credibility to the results from the assigned VOT treatments, because it would show that substantial misrelevation is occurring in the innate VOT treatments as well. Table III.XI compares the average intra-subject standard deviation of bids between the two types of treatments. In other words, it shows a comparison of how much a subject's bids fluctuate in the two types of treatments.

TABLE III.XI
Comparison of Intra-subject Bid Volatility between Assigned and Innate VOT Treatments

Treatment type	**Average bid**	**Average standard deviation of bids**	**Normalized average standard deviation of bids**
Assigned VOT	1.17 $/sec	.42 $/sec	*.36*
Innate VOT	9.7 $/hr	3.48 $/hr	*.36*

Notes: Average standard deviation of bids shown in the right-most column is normalized by the mean of subjects' bids.

The normalized standard deviations of intra-subject bids are identical between the two types of treatments. This is evidence that misrevelation in the assigned VOT treatments reflects limitations of the mechanism itself rather than confusion over the premise of being assigned a VOT.

III.V CONCLUSION

Study participants often provided truthful values of time when a Vickrey-Clarke-Groves mechanism was used to allocate express-lane access, but frequently did not. The discrepancies between subjects' actual and reported VOT reduced the performance of the VCG mechanism compared to the theoretical optimum. However, the VCG mechanism was able to dominate alternatives in terms of social welfare.

The percent of truthful bids for any treatment never exceed 45%, nor fell below 25%. Truthful revelation was higher in this VCG application than binary-outcome public good games studied in the literature, but lower than applications encompassing a more continuous range of outcomes to select from. These results fit with trend of results found in the literature, since lane assignment encompasses one very significant binary outcome (a subject's lane assignment) and a minor, more continuous outcome (the speed of each set of lanes). This correspondence to the literature is strong evidence that time savings can be allocated similarly to other goods using truthful revelation mechanisms – a novel finding of this work. This work was also able to show that even with explicit "coaching," subjects still employ a variety of non-truthful bidding strategies. Many appeared to employ myopic reinforcement learning, and others appeared to simply choose at random. However, both truth-telling and system welfare were higher with the treatment where subjects received this pre-experiment coaching, suggesting that training can improve mechanism performance.

Subjects in this experiment did not learn to bid more truthfully in aggregate. This work identified two significant factors that impaired learning truthful revelation. One was the lack of salient incentives for truth-telling over ranges of VOTs that did not span the cutoff for toll vs. free lane assignment. High VOT subjects faced little consequences for bidding even higher, and the reverse was true for low VOT subjects. This incentive structure allowed for suboptimal equilibria to exist, which may allow misrevelation to persist. The second factor impairing learning was travel time stochasticity that results from traffic dynamics. This stochasticity distorts the relationship between a subject's bids and earning/utility, delaying myopic reinforcement learning. This phenomenon is inherent to the application of the mechanism for lane allocation, simply due to costly human errors in the driving task. Despite the lack of reductions in misrevelation over time, subjects do reduce the frequency of costly assignments to the lane that is sub-optimal given their VOT. As a result, the efficiency of the mechanism improved over time.

Although subject misrevelation resulted in welfare losses for the VCG mechanism compared to the theoretical optimum, the mechanism often improved driver welfare relative to alternatives. In treatments where subjects were provided with highly detailed explanations of the mechanism during pre-trial instructions, driver welfare was higher with the VCG than with alternative schemes such as second-best tolling and dynamic pricing that keeps the toll lane in free flow. In treatments with low-quality instructions, the VCG mechanism did not improve welfare relative to some second-best tolling schemes –

but the calculations for these latter schemes assumed implausibly accurate knowledge of user VOTs, given what field studies are capable of. The upshot is that given the parameters of this experiment, a well-implemented VCG mechanism performed better than alternatives in terms of social welfare.

Robustness checks using treatments that elicited subjects' "true" or "innate" VOT based on real-world opportunity costs, rather than one induced by a payment depreciation rate, validated much of the experiments findings. The VOTs elicited by the mechanism were plausible given the subject pool and typical laboratory payrate, and were influenced by demographic, attitudinal, and schedule-related factors in ways that matched the literature and intuition.

Just as the assigned VOT treatments elicited a substantial share of non-truthful bids, the innate VOT treatments elicited some VOTs that seemed implausibly high. Subjects revealed VOTs as high as $28/hour on average, almost double what the experimental laboratory typically pays per hour. Furthermore, there was significant volatility in elicited bids among subjects, which corresponded to the level of volatility observed in assigned-VOT sessions. This volatility confirms that subjects are not always revealing their true VOT, and thus the mechanism is limited in its ability to elicit truthful VOTs and in term select socially optimal outcomes.

Although the VCG mechanism was shown to achieve lane allocations with improved user welfare relative to other pricing schemes, the welfare results should be interpreted cautiously. A major limitation of this work is that the experimental environment was designed such that the only relevant travel time preference was linear value of time. In reality, reliability and urgency are important preferences that have significant welfare implications for express lane allocation. Eliciting several types of preferences simultaneously is a much more demanding for human users of a VCG mechanism, especially given the difficulty that many subjects had in truthfully providing only one preference. Furthermore, unlike VOT, nonlinear preferences such as urgency are logistically more complex to elicit. It is possible that once such preferences are taken into account, the VCG will be unable to improve driver welfare relative to other toll lane pricing schemes. For example, free-flow pricing of toll lanes is a more formidable benchmark when reliability is also considered rather than simply value of time. Nonetheless, this work demonstrates that VCG mechanisms can to elicit truthful travel time preferences. Based on the ability to improve social welfare when allocating drivers to freeway lanes based on value of time, it is worth exploring the mechanism further to ascertain whether benefits can also be achieved when additional, more complex preferences are considered.

One final limitation is that misrevelation and learning was studied on a short time horizon in the absence of influence from subjects' peers and the media. External sources of information could play a significant role in the revelation strategies of mechanism users, and their effects should be studied further.

BIBLIOGRAPHY

Adams, H. S., M. J. Nieuwenhuijsen, and R.N. Colvile, “Determinants of fine particle (PM2.5) personal exposure levels in transport microenvironments,” *Atmospheric Environment,* 35 (2001), 4557-4566.

Ailshire, Jennifer A., Philippa Clarke, “Fine particulate matter air pollution and cognitive function among older US adults,” *The Journals of Gerontology Series B,* 70 (2015), 322-328.

Air Quality Index (AQI) - A Guide to Air Quality and Your Health, Retrieved December 22, 2014, from http://airnow.gov/index.cfm?action=aqibasics.aqi

Arentze, Theo, and Harry Timmermans. “Modeling learning and adaptation processes in activity-travel choice,” *Transportation,* 30.1 (2003), 37-62.

Arnott, Richard, André De Palma, and Robin Lindsey, "Route choice with heterogeneous drivers and group-specific congestion costs," *Regional Science and Urban Economics,* 22 (1992), 71-102.

Arthur, Brian, “Designing Economic Agents that Act like Human Agents: A Behavioral Approach to Bounded Rationality,” *The American Economic Review*, 81 (1991), 353-359.

Attiyeh, Greg, Robert Franciosi, and R. Mark Isaac, “Experiments with the Pivot Process for Providing Public Goods,” *Public Choice*, 102 (2000), 93-112.

Ben-Elia, Eran, and Yoram Shiftan, “Which Road do I Take? A Learning-based Model of Route Choice with Real-time Information,” *Transportation Research Part A: Policy and Practice,* 44 (2010), 249-264.

Bento, Antonio, Kevin Roth, and Andrew Waxman, “The Value of Urgency: Evidence from Congestion Pricing Experiments,” Working Paper, (2014).

Bhattacharjee, H., M. Drescher, T. Good, and Z. Hartley, "Particulate Matter in New Jersey," (1999), Retrieved from https://www.princeton.edu/step/conferences-reports/reports/

Blincoe, Lawrence J., Ted R. Miller, Eduard Zaloshnja, and Bruce A. Lawrence, "The economic and societal impact of motor vehicle crashes*," Report No. DOT HS 812 013*, (2010).

Bogers, Enide, Michel Bierlaire, and Serge Hoogendoorn, "Modeling Learning in Route Choice," *Transportation Research Record: Journal of the Transportation Research Board,* (2007): 1-8.

Braid, Ralph M, "Peak-Load Pricing of a Transportation Route with an Unpriced substitute," *Journal of Urban Economics*, 40 (1996), 179-197.

Brenner, Donald R., and John Morgan, "The Vickrey-Clarke-Groves versus the Simultaneous Ascending Auction: An Experimental Approach," *SSRN Electronic Journal*, (1998).

Brownstone, David, and Kenneth A. Small, "Valuing Time and Reliability: Assessing the Evidence from Road Pricing Demonstrations," *Transportation Research Part A*, 39 (2005), 279–293.

California Department of Public Health, "Air Quality Data Query Results," (2012), Retrieved from http://www.ehib.org/page.jsp?page_key=130&year=2006&pmn=LTPM25AVG

Cantarella, Giulio E., and Ennio Cascetta, "Dynamic Processes and Equilibrium in Transportation Networks: Towards a Unifying Theory," *Transportation Science,* 29 (1995), 305-329.

Cason, Timothy N., Lata Gangadharan, and Charlotte Duke, "Market Power in Tradable Emission Markets: a Laboratory Testbed for Emission Trading in Port Phillip Bay, Victoria," *Ecological Economics*, 46 (2003), 469–491.

Chang, Tom, Joshua S. Graff Zivin, Tal Gross, and Matthew J. Neidell, "Particulate pollution and the productivity of pear packers," *NBER Working Paper*, (2014).

Chen, Roger, and Hani Mahmassani, "Travel Time Perception and Learning Mechanisms in Traffic Networks," *Transportation Research Record: Journal of the Transportation Research Board,* 1894 (2004), 209-21.

Chorus, Caspar, and Benedict Dellaert, “Travel Choice Inertia: The Joint Role of Risk Aversion and Learning,” *ERIM Report Series Research in Management, Erasmus Research Institute of Management,* (2010).

Clarke, Edward H, “Multipart Pricing of Public Goods,” *Public Choice*, 11 (1971), 17–33.

Crawford, Vincent P, "Adaptive Dynamics in Coordination Games," *Econometrica,* 63 (1995).

Eisenberg, Daniel, “The mixed effects of precipitation on traffic crashes,” *Accident Analysis and Prevention,* 36 (2004), 637–647.

Erev, Ido, and Alvin Roth, “Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria,” *The American Economic Review,* 88 (1998), 848-881.

Greenwald, R., J. Sarnat, S. Sarna, F. Yip, and T. Boehmer, “In-Cabin Air Pollution Exposure and Acute Respiratory Response in Healthy and Asthmatic Automobile Commuters,” *American Thoracic Society International Conference Abstracts*, (2012).

Greibe, Poul, “Accident prediction models for urban roads,” *Accident Analysis and Prevention*, 35 (2003), 273-285.

Groves, Theodore, “Incentives in Teams,” *Econometrica*, 41 (1973), 617-631.

Horowitz, Joel L, "The stability of stochastic equilibrium in a two-link transportation network," *Transportation Research Part B: Methodological,* 18.1 (1984), 13-28.

Iida, Yasunori, Takamasa Akiyama, and Takashi Uchida, "Experimental Analysis of Dynamic Route Choice Behavior," *Transportation Research Part B: Methodological,* 26 (1992), 17-32.

Johnson, Bruce, “Travel Time and The Price of Leisure,” *Economic Inquiry*, 4 (1966), 135-145.

Jung, Soyoung, Kitae Jang, Yoonjin Yoon, and Sanghyeok Kang, “ Contributing factors to vehicle to vehicle crash frequency and severity under rainfall,” *Journal of Safety Research*, 50 (2014), 1-10.

Kagel, John H., and Dan Levin, “Auctions: a Survey of Experimental Research,” *The Handbook of Experimental Economics*, 2 (2016).

Kawagoe, Toshiji, and Toru Mori, “Can the Pivotal mechanism induce truth-telling? An experimental study,” *Public Choice*, 108 (2001), 331-354.

Keay, Kevin, and Ian Simmonds, “Road accidents and rainfall in a large Australian city,” *Accident Analysis & Prevention,* 38 (2006), 445-454.

Kong, SiYuan, et al. “An Experimental Study of Route Choice and Incident Management Using a Real-Time 2D Driving Simulator with Variable Message Signs.” Transportation Research Board 96th Annual Meeting, (2017).

Koster, Paul, Stefanie Peer, and Thijs Dekker, "Memory, Expectation Formation and Scheduling Choices," *Economics of Transportation,* 4 (2015), 256-65.

Lavy, Victor, Avraham Evenstein, and Sefi Roth, “The Impact of Short Term Exposure to Ambient Air Pollution on Cognitive Performance and Human Capital Formation,” *NBER Working Paper*, (2014).

Li, Lianfa, Jun Wu, Neelakshi Hudda, Constantinos Sioutas, Scott A. Fruin, and Ralph J. Delfino, “Modeling the concentrations of on-road air pollutants in southern California,” *Environmental Science & Technology,* 47 (2013), 9291−9299.

Liu, Xiaozhen, and H. Christopher Frey, “Modeling of in-Vehicle Human Exposure to Ambient Fine Particulate Matter,” *Atmos. Environ*., 45 (2011), 4745–4752.

Mahmassani, Hani S., “Learning from interactive experiments: travel behavior and complex system dynamics,” *The expanding sphere of travel behavior research,* (2009), 131-158.

Mahmassani, Hani S., “Dynamic models of commuter behavior: Experimental investigation and application to the analysis of planned traffic disruptions,” *Transportation Research Part A,* 24 (2010), 465-484.

Mahmassani, Hani S., and D. G. Stephan, “Experimental investigation of route and departure time choice dynamics of urban commuters,” *Transportation Research Record Journal of the Transportation Research Board,* 1203 (1988).

McLaughlin, Kevin, and Daniel Friedman, “Online ad auctions: An experiment,” *ESI Working Paper*, (2016).

Monderer, Dov, and Lloyd S. Shapley, "Potential Games," *Games and Economic Behavior,* 14 (1996), 124-43.

Nakayama, Shoichiro, Ryuichi Kitamura, and Satoshi Fujii, "Drivers' Route Choice Rules and Network Behavior: Do Drivers Become Rational and Homogeneous Through Learning?," *Transportation Research Record: Journal of the Transportation Research Board,* 1752 (2001), 62-68.

National Highway Traffic Safety Administration, “Examination of Crash Contributing Factors Using National Crash Databases,” *DOT-VNTSC-NHTSA-02-07*, (2003).

NJDEP (New Jersey Department of Air Quality and Planning), “Fine Particulate Matter (PM2.5) State Implementation Plan, Chapter 2: Air Quality Ambient and Emission Inventory Data,” (2009).

Noland, Robert B., and Kenneth A. Small, “Travel-Time Uncertainty, Departure Time Choice, and the Cost of Morning Commutes,” *Transportation Research Record*, 1493 (1994), 150-158.

Nussbaum, E. Michael, Sherif Elsadat, and Ahmed H. Khago, “Best Practices in Analyzing Count Data: Poisson Regression,” *Best Practices in Quantitative Methods* (Chapter 21), (2007).

Parkany, Emily, “Traveler Responses to New Choices: Toll vs. Free Alternatives in a Congested Corridor,” PhD Dissertation, (1999).

Pu, Liang, Hong-Zhi Guan, Zi-Xu Lu, Shan Huang, and Jia Fang, "Individual and Population Learning in Route Choice Behavior of Drivers: Experiment and Model," *ICCTP 2009*, (2009).

Selten, R., T. Chmura, T. Pitz, S. Kube, and M. Schreckenberg. "Commuters Route Choice Behaviour,"*Games and Economic Behavior,* 58 (2007), 394-406.

Small, Kenneth A., Clifford Winston, and Jia Yan, “Uncovering the Distribution of Motorists' Preferences for Travel Time and Reliability,” *Econometrica*, 73 (2005), 1367-1382.

Small, Kenneth A., and Jia Yan. “The Value of ‘Value Pricing’ of Roads: Second-Best Pricing and Product Differentiation,” *Journal of Urban Economics*, 49 (2001), 310-336.

Small, Kenneth A., “The Scheduling of Consumer Activities: Work Trips,” *American Economic Review*, 72 (1982), 467-79.

Stanek, L. W., J. S. Brown, J. Stanek, J. Gift, and D. L. Costa, “Air pollution toxicology—a brief review of the role of the science in shaping the current understanding of air pollution health risks,” *Toxicol. Sci.,* 120 (2011), S8–S27.

Sullivan, E., “Evaluating the Impacts of the SR 91 Variable-Toll Express Lane Facility. Final Report,” *Transport Research International Documentation - TRID*, (1998).

SUNY Suffolk, “Lecture on Clouds, Fog, and Precipitation [Powerpoint Lecture Slides],” Retrieved from http://www2.sunysuffolk.edu/mandias/met101/Clouds_Fog_Precip.pdf

Tai, A. P. K., L. J. Mickley, D. J. Jacob, E. M. Leibensperger, L. Zhang, J. A. Fisher, and H. O. T. Pye, “Meteorological Modes of Variability for Fine Particulate Matter (PM2.5) Air Quality in the United States: Implications for PM2.5 Sensitivity to Climate Change,” *Atmospheric Chemistry and Physics*, 12 (2012), 3131–3145.

U.S. DOT FHA (US Department of Transportation, Federal Highway Administration). “How Do Weather Events Impact Roads?,” (2014), Retrieved from http://www.ops.fhwa.dot.gov/weather/q1_roadimpact.htm

U.S. EPA., “Understanding Particle Pollution,” (2004), Retrieved from http://www.epa.gov/airtrends/aqtrnd04/pmreport03/pmunderstand_2405.pdf

Verhoef, Erik, and Kenneth A. Small, “Product Differentiation on Roads,” *Journal of Transport Economics and Policy*, 38 (2004), 127-156.

Vickrey, William, “Counterspeculation, Auctions, And Competitive Sealed Tenders,” *The Journal of Finance*, 16 (1961), 8–37.

Vreeswijk, Jaap, Michael Do, Wilco Middag, Marieke Martens, Eric Van Berkum, and Bart Van Arem, "Travelers' ability to observe changes in traffic intensities and traffic light settings," *14th International IEEE Conference on Intelligent Transportation Systems,* (2011).

Vreeswijk, Jaap, Tom Thomas, Eric Van Berkum, and Bart Van Arem, "Perception bias in route choice," *Transportation,* 41.6 (2014), 1305-321.

Watkins, Christopher J. C. H., and Peter Dayan, "Q-learning," *Machine Learning,* 8 (1992), 279-92.

Watling, David, David Milne, and Stephen Clark, “Network Impacts of a Road Capacity Reduction: Empirical Analysis & Model Predictions,” *Transportation Research Part A: Policy and Practice,* 46 (2012), 167–189.

Wu, Jun, “Air Quality Modeling [Powerpoint Lecture Slides],” Retrieved by e-mail from author, (2014).

Zhu, Shanjian, David Levinson, Henry Liu, and Kathleen Harder, “The traffic and behavioral effects of the I-35W Mississippi River bridge collapse,” *Transportation Research Part A: Policy and Practice,* 44 (2010), 771-784.

APPENDICES

Chapter I: The Effect of Ambient Air Pollution on the Frequency of Motor Vehicle Collisions

Chapter III: Do Truthful-Bidding Mechanisms Improve the Allocation of Drivers to Express Lanes?

APPENDIX A

Chapter I Appendix

A.I In-depth Reasoning Behind Dropping Precipitation Observations

The difficulty in controlling for precipitation lies in the complexity of its relationship with both accidents and pollution. For example, precipitation, which itself is strongly correlated with PM 2.5 levels (Tai et al. 2009), affects accident rates differently depending on whether rain, snow, or ice is involved (Eisenberg 2004). The form that precipitation takes (rain, snow, or ice) is in turn affected by temperature, which is also strongly correlated with PM 2.5 (Li et al. 2013). Furthermore, both snow and rain both have major nonlinearities in their effects on accident rates, due in part to changes in driver caution as a response to challenging conditions, making them even more difficult to control for (Eisenberg 2004). The analysis becomes much more complicated without proper controls for vehicle counts, because both precipitation and temperature will also affect the number of vehicles on the road – introducing even more non-linearity into their effects on accidents. There will also likely be variation in precipitation's effect from city to city based on differences in local snow and ice management, and also based on differences in on-road water accumulation due to variation in road geometry (Jung et al. 2014). Precipitation's effects on accidents also vary significantly based on the amount of time transpired since the last precipitation (Keay and Simmonds 2005). Lastly, the effects of precipitation on accidents are also dependent on wind speed and humidity, which themselves affect the formation and melting of ice and snow (U.S. DOT FHA). Wind speed and humidity are also correlated with PM 2.5 concentrations, introducing another potential source of bias.

A.II Full Results from the Main Regression

TABLE A.II.A

Results from the Main Regression, Part 1

Negative binomial regression			Number of obs = 4727
			LR chi2(60) = 1704
Dispersion = mean			Prob > chi2 = 0
Log likelihood = -13411			Pseudo R2 = .0598
Accidents	Coef.	Std. Err.	P-value
Lowaqi	0.0004	0.0006	0.5060
Highaqi	0.0096	0.0263	0.7150
aqi50	-0.0005	0.0008	0.5180
aqi100	0.0002	0.0038	0.9490
city2			
Elizabeth	-0.3351	0.0244	0.0000
Jersey City	-0.3030	0.0486	0.0000
Newark	-0.3718	0.0277	0.0000
North Brunswick	-0.0896	0.0672	0.1820
Paterson	-0.1884	0.0297	0.0000
Toms River	0.1847	0.0522	0.0000
Trenton	-0.1445	0.0261	0.0000
Union City	-0.5205	0.0554	0.0000

TABLE A.II.B

Results from the Main Regression, Part 2

month			
2	0.0501	0.0351	0.1530
3	0.1742	0.0360	0.0000
4	0.1765	0.0394	0.0000
5	0.2224	0.0427	0.0000
6	0.2339	0.0469	0.0000
7	0.1546	0.0480	0.0010
8	0.1119	0.0481	0.0200
9	0.1833	0.0445	0.0000
10	0.1540	0.0401	0.0000
11	0.1275	0.0382	0.0010
12	0.1207	0.0359	0.0010
coast#month			
1 1	-0.1170	0.0639	0.0670
1 2	-0.1923	0.0654	0.0030
1 3	-0.4119	0.0682	0.0000
1 4	-0.2670	0.0678	0.0000
1 5	-0.1006	0.0653	0.1230
1 6	-0.1105	0.0653	0.0910
1 7	-0.0062	0.0635	0.9220
1 8	0.1004	0.0653	0.1240
1 9	-0.0890	0.0638	0.1630
1 10	-0.1160	0.0659	0.0780
1 11	-0.0894	-0.0704	0.2040
1 12	0.0000	(omitted)	

TABLE A.II.C

Results from the Main Regression, Part 3

nyc#month			
1 1	-0.0704	0.0612	0.2500
1 2	-0.0867	0.0596	0.1460
1 3	-0.1371	0.0594	0.0210
1 4	-0.0893	0.0580	0.1240
1 5	-0.0742	0.0588	0.2070
1 6	-0.0437	0.0616	0.4790
1 7	-0.0911	0.0573	0.1120
1 8	-0.0570	0.0579	0.3250
1 9	-0.0491	0.0570	0.3890
1 10	-0.0605	0.0591	0.3060
1 11	-0.0739	0.0610	0.2260
1 12	0.0000	(omitted)	
day2			
M	-0.1263	0.0183	0.0000
S	-0.3664	0.0192	0.0000
SA	-0.1463	0.0181	0.0000
TH	-0.1486	0.0184	0.0000
TU	-0.1688	0.0182	0.0000
W	-0.1758	0.0185	0.0000
Trend	0.0000	0.0000	0.5860
Meanwind	0.0043	0.0019	0.0250
Meantemp	-0.0092	0.0022	0.0000
temp2	0.0001	0.0000	0.0010
Meanvisibility	-0.0013	0.0078	0.8720
Minvisibility	-0.0040	0.0033	0.2260

TABLE A.II.D

Results from the Main Regression, Part 4

Minvisone	0.0005	0.0475	0.9920
Meanvisone	0.8389	0.2907	0.0040
Cloudcover	-0.0146	0.0027	0.0000
_cons	-8.6901	0.0903	0.0000
ln(population)	1.0000	(exposure)	
/lnalpha	-3.3697	0.0721	
Alpha	0.0344	0.0025	

Chapter III Appendix

B.I Additional Tables and Figures

SUPPLEMENTARY TABLE B.I.I

Summary of Relevant Experimental Sessions and Treatments

Session	Toll Type	Incentive Type*	Instruction Quality*	# VOT Changes	# Capacity Changes	Interactive driving*	Endogenous pricing*
1	VCG	Payment Dep.	S-S-S	2	1	Yes	Yes
2	VCG	Payment Dep.	D-S-S	2	1	Yes	Yes
3	VCG	Payment Dep.	S-D-S	2	1	Yes	Yes
4	VCG	Payment Dep.	D-S-S	1	0	Yes	Yes
5	VCG	Payment Dep.	D-D-D	1	0	Yes	Yes
6	VCG	Payment Dep.	D-D-D	1	0	Yes	Yes
7	VCG	Payment Dep.	D-D-D	1	0	Yes	Yes
8	VCG	Opp. Cost	N/A	N/A	0	No	No
9	VCG	Opp. Cost	N/A	N/A	0	No	No
10	VCG	Opp. Cost	N/A	N/A	0	No	No
11	VCG	Opp. Cost	N/A	N/A	0	No	No
12	VCG	Opp. Cost	N/A	N/A	0	No	No
13	VCG	Opp. Cost	N/A	N/A	0	No	No
14	VCG	Opp. Cost	N/A	N/A	0	No	Yes
15	VCG	Opp. Cost	N/A	N/A	0	No	Yes
16	Single	Payment Dep.	N/A	2	1	Yes	No
17	Single	Opp. Cost	N/A	N/A	0	Yes	No

Notes: Incentive types that are Payment Depreciation refer to assigned VOT sessions, while incentive types that are Opportunity Cost refer to "leave when finished" sessions.

Instruction quality refers to pre-trial explanations of VCG mechanisms, truth-telling, and value of time - respectively. "S" refers to a sparse explanation, and "D" refers to a detailed explanation.

Interactive driving refers to whether subjects traveled together on the same roadway.

Endogenous pricing refers to whether lane speed, pricing, and/or assignment were a function of other subjects' decisions.

SUPPLEMENTARY FIGURE B.I.I

Plot of Revealed Versus Assigned VOTs for each Subject-Round-Treatment

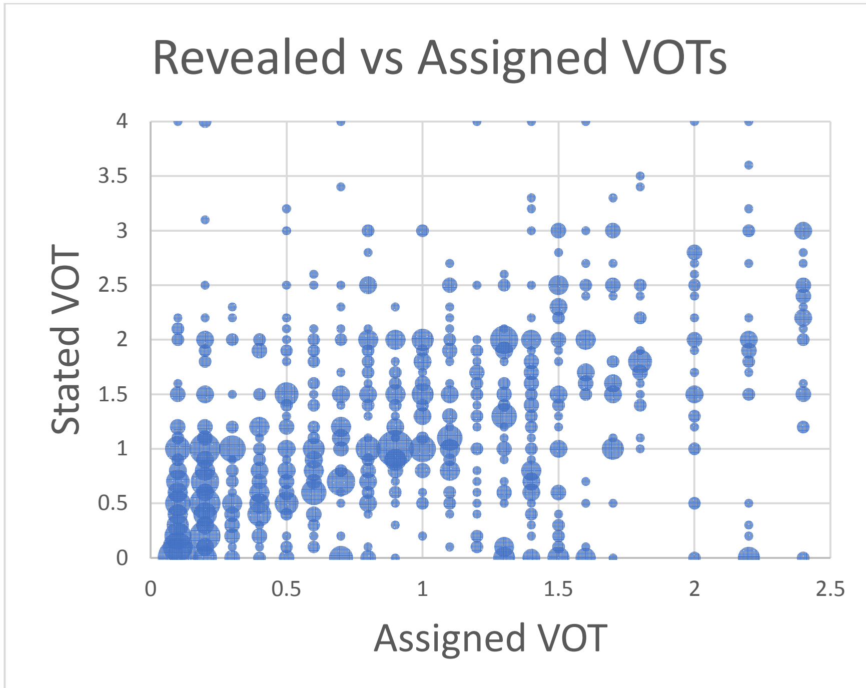

Notes: The figure above plots each VOT provided by experimental subjects along with the VOT assigned to that subject at the time of the bid. The data are from rounds occurring before the first time subject VOTs are re-assigned. Larger bubbles occur at coordinates where there is more density in bids.

SUPPLEMENTARY FIGURE B.I.II

Magnitude of Misbidding by Experimental Round Averaged across Sessions

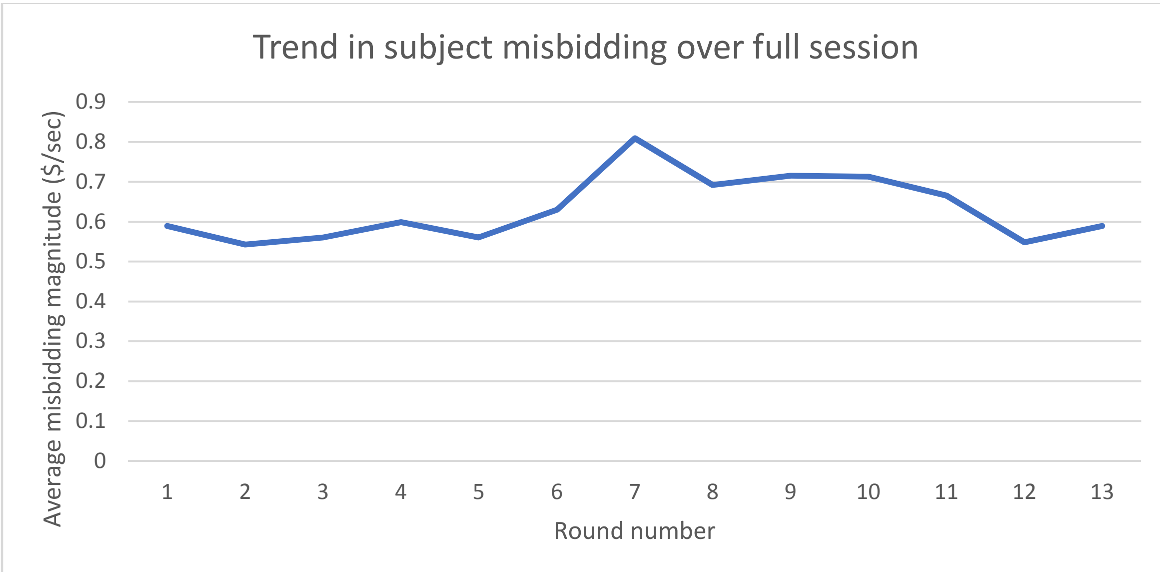

Notes: Only sessions that treated subjects with one or fewer changes to their assigned VOT were included.

B.II Welfare Costs of Second-Best VCG Mechanism Implementation

There were no significant welfare losses issues associated with implementing the VCG mechanism as a second-best solution where only the fast lane is tolled. The table below shows the efficiency losses incurred by restricting the slow lanes to be free.

TABLE B.II.I

Welfare Losses Incurred by Imposing a Price of $0 on the Slow Lanes.

Scenario	**VOT 1 Capacity 1**	**VOT 2 Capacity 1**	**VOT 3 Capacity 1**	**VOT 3 Capacity 2**
Incentivized Welfare Loss in Equilibrium	0%	1%	0%	0%
Worst Case Welfare Loss in Equilibrium	1%	1%	0.5%	0%
Experimental Welfare Loss	< .2%	< .2%	< .2%	0%

Notes: The row headings differentiate between three efficiency-loss metrics analyzed, each of which are described in further detail below the table. The column headings differentiate between combinations of three distinct VOT distributions tested in the analysis and two distinct free-lane road capacities tested.

1) Incentivized efficiency loss in equilibrium

Although in theory leaving the slow lanes unpriced could create slight incentives for low VOT to inflate their bids, these incentives were typically not present given experimental parameters. This was tested by simulating the experiment in the case where every subject truthfully provided a VOT that was equal to their payment depreciation rate. Next, instances were identified where a subject could increase their earnings by stating a VOT that was higher than their payment depreciation rate. Only one experimental scenario out of provided such an opportunity. Exploiting this opportunity increases the social cost of travel delay by 1%. Obviously in the other cases, there was a 0% increase since there was no opportunity. This constitutes what theoretical efficiency losses due to the second-best implementation would be in equilibrium. On average, these losses were .25%. To put this loss in perspective, it was equal to approximately 1/5th of the best-case efficiency loss when using a standard toll. Therefore in theory, the second-best implementation was not a deal-breaker for the VCG. Furthermore, for these losses to be manifested, however, subjects would need to identify and exploit these opportunities.

2) Worst case efficiency loss in equilibrium

Although in most cases subjects did not have an individual incentive to inflate their bid, there was also no disincentive for them to inflate their bid. Thus in the worst case of every low VOT increasing their bid up until the very point where they would be switched to the toll lane, it was possible for further efficiency losses to occur. This worst case, which corresponds to every free-lane user bidding the maximum possible VOT that still lets them use the free lane, was simulated. It was found that for every scenario but one, efficiency losses due to extra users on the toll lane would equal between .5% and 1%. On average, this loss was equal to .625%. To put this loss in perspective, it was equal to approximately 1/2th of the best-case efficiency loss when using a standard toll. Thus even in the worst-possible case, the second-best implementation of the VCG was still beneficial, but half the benefits were eroded. This worst case was almost impossible to be achieved in practice. There was virtually no way subjects would know exactly where the toll lane cutoff would be, and attempting to do so would almost certainly end up with them frequently being assigned to the toll lane (which would make them worse off).

3) Experimental efficiency loss

It was not possible to directly the social cost of using the second-best implementation of the VCG in practice, because no "first-best" counterfactual sessions were run for assigned-VOT treatments that used an implementation where all users (even those in the slow-lane) were charged. The best estimate from the available data was obtained by identifying the average bid-inflation of experimental low VOT users, and identifying what welfare losses this amount of bid-inflation would cause if all other subjects were to bid their true VOT. The average experimental bid inflation was about 15 cents per second, or 20% of the average free lane VOT. This amount of bid inflation never increased overall travel delay cost by more than .2% when simulated. Furthermore, overbidding was just as prevalent among high-VOT users who had no incentive to do so. Lastly, there were only a handful of instances where subjects who had an incentive to inflate their bid actually did. Therefore being constrained to the second-best implementation poses minor theoretical efficiency concerns for the VCG, and negligible practical ones. This shows that the mechanism was compatible with the desire of planners to leave one set of lanes unpriced.

Additional evidence of the practicality of second-best VCG mechanisms comes from assigned VOT treatments. The two figures below show each subject's average bid for two different sessions that differ only in whether or the first-best (tolls on both sets of lanes) or second-best (tolls on fast lane only) VCG mechanism is implemented.

FIGURE B.II.I

Distribution of Subjects' Average VOT Elicited from a Second-best VCG Mechanism

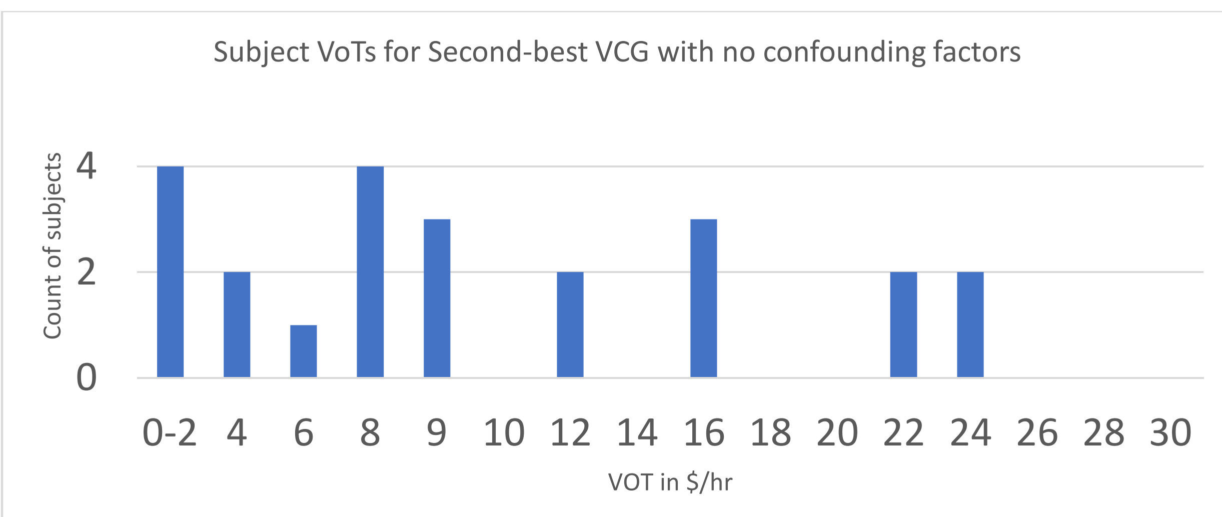

FIGURE B.II.II

Distribution of Subjects' Average VOT Elicited from a Second-best VCG Mechanism

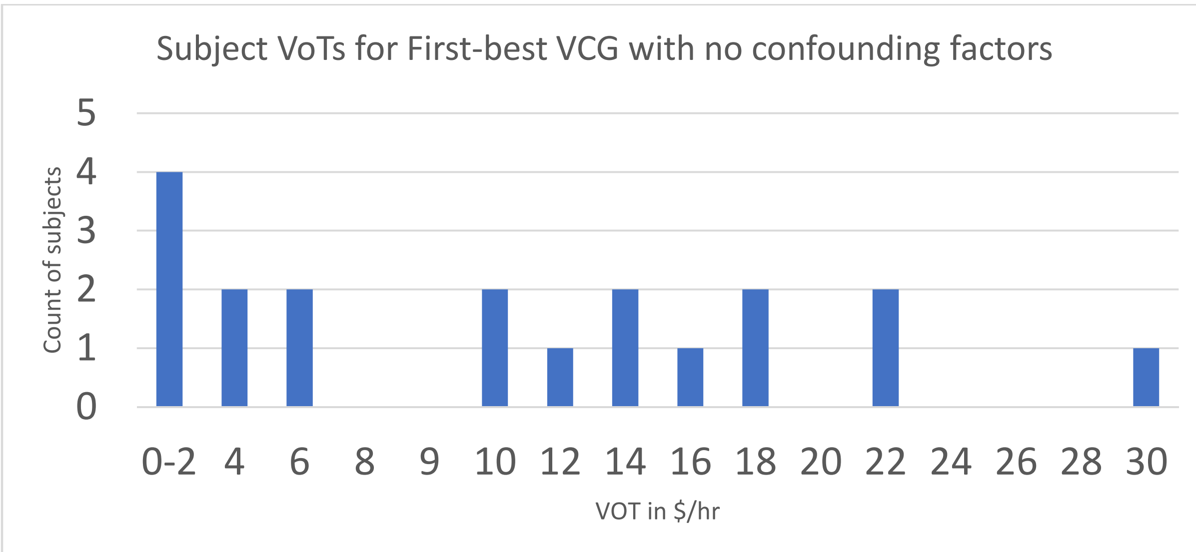

When comparing the two distributions, there are no obvious difference between the first-best and second-best implementation of the VCG mechanism. Both have roughly the same shape, and the same mean of $10/hour. This is additional evidence that a second-best implementation of a VCG mechanism does not lead to substantial deviations in user behavior, and therefore has a negligible impact on user welfare.