Can Usage-Based Pricing Reduce Traffic Congestion?*
[Extended abstract; Please do not circulate]

Itai Ater[†] Adi Shany[‡] Brad Ross[§] Eray Turkel[¶] Shoshana Vasserman[‖]

September 2023

Abstract

This paper analyzes the effects of the largest field experiment to date that incentivizes drivers to limit driving during peak hours and congested areas via usage-based congestion pricing. The experiment monitored the driving behavior of 10,000 Israeli drivers who were recruited over the course of 2020. During the first six months of a driver's participation in the experiment, their driving behavior was monitored and recorded; afterwards, drivers received an annual budget and were charged for each kilometer driven during historically congested times in congested areas. Whatever remained in each driver's budget was paid to them a year later. We use comprehensive data on driving behavior of participating drivers over the course of 2020 and 2021 to evaluate how usage-based congestion pricing affects driving behavior. The staggering of driver recruitment over time facilitates identification of treatment effects via a difference-in-differences approach. We find that drivers decrease their congested driving behavior by approximately 10% or more of their pre-treatment medians across a myriad of outcomes designed to detect both extensive margin (i.e. whether to take a trip) and intensive margin (i.e. when to take a trip and via which route) responses. We also find that there is significant treatment effect heterogeneity across drivers that can be predicted by pre-treatment driving behavior. The most affected drivers tend to be those who contributed more to congestion, those who appear to have more flexibility in their driving choices, and those with easier access to public transit, but not those who are disproportionately socioeconomically advantaged or disadvantaged. Finally, we estimate how traffic density causally affects traffic speed on the largest highway in Israel using a year of high-frequency traffic sensor data. We use this structural relationship to translate our treatment effects on driving behavior up to 30% increases in traffic speeds during the most congested times of day.

*First and foremost, we thank Noa Bar-Nir, Fabian Barrett, Hadar Fuks, Shai Gilat, Andrew Meyers, Yotam Nir, Doron Zamir, and Michelle Zheng for their exemplary research assistance. We are also grateful to our colleagues and numerous seminar and conference participants for their helpful comments.

[†]Tel Aviv University, Coller School of Business and CEPR. Email: ater@post.tau.ac.il.

[‡]Tel Aviv University, Coller School of Business. Email: adishany@tauex.tau.ac.il.

[§]Stanford Graduate School of Business. Email: bradross@stanford.edu

[¶]Google. Email: eturkel@stanford.edu

[‖]Stanford Graduate School of Business and the NBER. Email: svass@stanford.edu.

1 Introduction

Traffic congestion is a notorious source of inefficiency, frustration and environmental damage. On a social level, traffic delays are associated with billions of dollars of waste (Chandler-Wilde, 2023) and substantial greenhouse gas emissions (Barth and Boriboonsomsin, 2008). On an individual level, traffic congestion is associated with serious physical health problems (Currie and Walker, 2011), and time spent in congestion is reported to be among the worst moments of the day (Kahneman et al., 2006). Given these costs, policymakers have sought effective means to reduce congestion, increasingly through policies that charge fees for driving in highly demanded areas. Zone-based entry fees ("Cordon pricing") have been in place in London since 2003, Stockholm since 2006, and Milan since 2008, and proposals for similar fee structures in Seattle, San Francisco, and New York City are currently progressing to the final stages.

However, there is little causal evidence that these programs work to meaningfully reduce congestion. A direct inspection of London congestion shows that road speeds initially increased following the implementation of the Cordon price policy, but have since rebounded. Existing experimental studies—which took place prior to the COVID-19 pandemic, and were relatively small in size and scope—have found that drivers make almost no significant changes in their rush-hour driving patterns in response to small monetary incentives (Kreindler, 2022; Martin and Thornton, 2017). These studies suggest that the price levels that transportation officials are willing to impose on congested roads are not large enough to meaningfully change the driving choices of heavy commuters, who already incur substantial costs from driving in the form of lost time and parking expenses. Still, it is unclear whether these economically insignificant effect estimates were due to short treatment durations, ineffective pricing, smaller sample sizes, and/or pre-pandemic norms around schedule flexibility and working from home.

This paper describes the results of what is, to our knowledge, the largest and longest-running usage-based congestion pricing experiment that has ever been implemented: a government-sponsored experiment called "Derech Erech"[1] that tracked the driving behavior of 10,000 drivers in Israel over the span of two years. As in many other countries, congestion is a first-order policy problem in Israel. The experiment recruited drivers from across commuting zones throughout Israel in a staggered fashion between 2020 and 2021. Each participating driver installed a monitoring device in their vehicle that could access their GPS coordinates over the course of their trips. During the first six months of participation, participants' driving behavior was monitored, but they did not receive any communication or monetary incentives. After the initial six-month period, drivers began to incur a time-and-location-specific fee for each kilometer driven out of a virtual budget, the remainder of which was paid out in cash upon drivers' annual vehicle registration renewals.

Our analysis is composed of three parts. First, we consider the direct effects of the congestion pricing experiment: to what extent did treated individuals reduce their driving in congested times and places, and how? To examine these questions, we apply an event study approach in the spirit of Borusyak et al. (2021) to estimate average treatment effects across a number of driving behaviors over the course of 16 weeks after individuals are treated. Next, we ask whether and how treatment

[1] Derech Erech is a play on words in Hebrew, as "Derech" means both "route" and "direction/way", and "Erech" means "valuable" or "to value").

effects co-vary with observable driver characteristics. This line of questioning is informative about two important dimensions along which congestion pricing policies like the one we study should be evaluated. The first dimension is the extent to which a congestion pricing scheme is effective at reducing the driving behavior that contributes the most to congestion. The second dimension is the inefficiency or regressivity of a congestion pricing scheme: does it impose high costs on individuals who are highly sensitive to prices but do not have enough flexibility in their schedules to avoid being charged? To address these questions, we apply the Sorted Effects method developed in Chernozhukov et al. (2018) to compare the drivers with the smallest and largest estimated treatment effects across covariates measuring their pre-treatment driving behavior, socioeconomic status, and the accessibility of public transit alternatives to trips they take by car. Finally, we consider the welfare implications of this congestion pricing policy in equilibrium were it to be applied to all Israeli drivers. To do so, we estimate a structural model of the demand by drivers in our sample for different routes and travel times for their commuting trips, along with a nonparametric statistical model of how traffic speeds change with the number of cars on the main highway in Israel. We then simulate equilibrium flows along the highway if a all drivers nationally were to be subjected to congestion pricing and evaluate the changes in road speeds, commute times and driver utility.

To identify the causal determinants of driving choices in our panel, we rely on two types of variation: variation in when drivers were recruited to the experiment, and variation in drivers' baseline commuting behavior that is nationally representative of drivers who would be impacted by proposed congestion pricing policies. Our dynamic treatment effect estimates are predicated on the staggered recruitment of individuals: an average of about 200 individuals were recruited per week over the course of 50 weeks. As such, on a typical week in which a driver switches from "monitoring" to "treated" status, there is a segment of other drivers are still in "monitoring" who can serve as a control group to separate individual-specific idiosyncrasies in driving behavior and time-varying aggregate conditions from the effects of our congestion pricing treatment. In keeping with Borusyak et al. (2021), our treatment effect estimator models each individual's baseline driving behavior using their "monitoring" period. It then infers the change in behavior due to the onset of pricing by comparing the observed "treated" behavior against a predicted baseline based on the same driver's previous ("monitoring") driving, and the concurrent driving of other drivers still in their "monitoring" phase.

Our dynamic treatment effect analysis suggests that drivers take some time to fully adjust to the congestion pricing scheme, as the treatment effects take several weeks to stabilize. However, after 10 weeks of operating under the pricing treatment, drivers substantially decrease their driving in congested areas and times. Aggregating across time, we find that drivers decrease a variety of measures of congested driving by an average of 10% or more relative to pre-treatment median outcomes, both on the extensive (e.g. number of trips) and the intensive (e.g. price per Km driven) margins.

Going beyond averages to understand the heterogeneity in effects, we focus on three broad dimensions: flexibility, socioeconomic status, and baseline demand for driving on congested roads. A well-designed congestion pricing intervention needs to take all of these factors into account. If a policy doesn't affect enough flexible drivers or flexible trips, it will raise revenue but fail to reduce congestion. Similarly, if prices are not high enough, drivers will not find it worthwhile to reduce their congestion-

inducing driving behavior. On the flip side, if we mostly affect price-sensitive drivers who have lower incomes, we may reduce congestion but our intervention will have a regressive effect. Moreover, if we only affect the behavior of flexible and price-sensitive drivers but if they are not the main contributors to congestion, our policy will again be ineffective in reducing congestion.

Using the Sorted Effects method to test for differences along these dimensions between drivers who were most and least responsive to the congestion pricing treatment, we generally find support for the potential efficacy of a congestion pricing policy in our setting. Across different characterizations of flexibility and baseline demand, the most responsive drivers in our sample exhibited more flexibility in their arrival and departure times and took more expensive cross-city and metro-oriented trips during their pre-treatment "monitoring" period. Treatment effects also correlated with the availability of alternatives to driving. Considering an index characterizing the convenience of the best public transit alternative to each vehicle trip taken by car during each driver's "monitoring" period, we find that the most responsive drivers also tended to have better access to relevant public transit options. We did not find significant differences in treatment effects based on indicators of socioeconomic status based on individuals' home neighborhoods and vehicle characteristics. While this last exercise is limited by the availability of precise individual data due to privacy restrictions, our results are suggestive that the possibility of disproportionate regressive impacts may not be a major concern.

Having established that our usage-based congestion pricing scheme meaningfully decreases driving during the most congested times and in the most congested areas, we next assess whether our treatment effects on driving behavior translate into meaningful impacts on road speeds. To do so, we first estimate how reducing the number of drivers on each road affects the speed at which cars can travel in congested areas. Focusing on driving along the largest highway in Israel, the Ayalon, we use a year of high-frequency, contemporaneous measurements of traffic flows to estimate the structural relationship between traffic densities and speeds nonparametrically. Importantly, we instrument for traffic density using exogenous variation in sensor distances from geolocated car accidents in space and time since traffic density and speed are determined simultaneously in equilibrium due to drivers' endogenous route choices.

We use our estimate of the congestion function to assess how the treatment effects from our experiment would impact highway speeds if they were to be applied to the entire population of Israeli drivers, proportionally to their area of residence. Applying our average treatment effect estimates directly—under a ceteris paribus assumption in which drivers do not additionally react to the change in road speeds—we find that the 7% decrease in highway trips implied by our ATT estimates would lead to an increase in traffic speeds of up to 30% during the most congested times of day. In ongoing work, we are working to relax the ceteris paribus assumption to account for individuals' equilibrium adjustments to their driving choices. To do so, we estimate a model of each individual's demand for taking their most frequent trips at different departure times (or not at all) given the prices and travel durations that the individual would have expected from each option under their treatment status and the contemporaneous traffic conditions. We then simulate the counterfactual demand for each stretch of highway across time in equilibrium, where the shares of drivers expected to use each stretch at a given time due to their travel decisions takes into account the impact of the aggregate travel flows on

the road speeds that they encounter.

Our paper offers several contributions to the literature: First, by studying a long-term experiment at a national scale, we are able to quantify the average impact of a congestion pricing mechanism that resembles the types of policies that are under consideration. The long duration of the experiment enables us to characterize drivers' responses over a longer period than previous studies. This point is particularly important since we find considerably larger effects on driving behavior when we extend the treatment time window, allowing time for participants to learn to adjust.

Second, observing driver-level incentives and driving responses allows us to identify the effect of congestion pricing on different subsamples of drivers. Understanding which populations are more sensitive to treatment and which are less could be instrumental for informing the design of maximally effective congestion pricing policies, as well as for assessing the scope for adverse distributional impacts. Third, we flexibly estimate the structural relationship between traffic density and speed while accounting for the endogeneity of these two quantities, which allows us to more accurately predict the counterfactual distribution of traffic speeds under our congestion pricing scheme.

In future work, we intend to extend our equilibrium framework to account not only for travel on major highways, but for travel throughout the Israeli road network writ large. To do this, we will combine data from Google Maps, public transit GPS records, and trips taken within our experiment to estimate a model of travel speeds along Israeli roads throughout our experiment. Using this panel of road speeds, we will scale up our travel demand model to account for all frequent trips (not just those that use the highway), estimate a non-highway congestion function, and simulate counterfactuals that target congestion in different areas. Finally, we will use our framework to examine the implications of a congestion pricing scheme not only in raw time savings but also in terms of shifts in consumer surplus across the population of drivers and to compare existing policy proposals against a hypothetical first best allocation of drivers to trips.

References

Barth, M. and K. Boriboonsomsin (2008): "Real-world carbon dioxide impacts of traffic congestion," *Transportation research record*, 2058, 163–171.

Borusyak, K., X. Jaravel, and J. Spiess (2021): "Revisiting event study designs: Robust and efficient estimation," *arXiv preprint arXiv:2108.12419*.

Chandler-Wilde, H. (2023): "These Are the World's Most Congested Cities," *Bloomberg*.

Chernozhukov, V., I. Fernández-Val, and Y. Luo (2018): "The sorted effects method: discovering heterogeneous effects beyond their averages," *Econometrica*, 86, 1911–1938.

Currie, J. and R. Walker (2011): "Traffic congestion and infant health: Evidence from E-ZPass," *American Economic Journal: Applied Economics*, 3, 65–90.

Kahneman, D., A. B. Krueger, D. Schkade, N. Schwarz, and A. A. Stone (2006): "Would you be happier if you were richer?" *science*, 312, 1908–1910.

Kreindler, G. (2022): "Peak-hour road congestion pricing: Experimental evidence and equilibrium implications," *Working paper*.

Martin, L. A. and S. Thornton (2017): "To Drive or Not to Drive: A Field Experiment in Road Pricing," Tech. rep., Working Paper.