

NBER WORKING PAPER SERIES

WELFARE ESTIMATES OF SHIFTING PEAK TRAVEL

Robert W. Hahn
Robert D. Metcalfe
Eddy Tam

Working Paper 31629
<http://www.nber.org/papers/w31629>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2023

We personally received no funding for this paper. We would like to thank Ryan Greene-Roesel at Bay Area Rapid Transit (BART) and the BART team for partnering with us to conduct this research and taking a leading role developing the experiments reported in this paper. We also thank the Federal Highways Agency and the San Francisco County Transportation Authority (especially Joe Castiglione and Camille Guiriba), and BART for funding the program. We greatly appreciate the thoughtful comments of Leo Bursztyrn, Thomas Chaney, Raj Chetty, Nathan Hendren, Jonathan Libgober, Victor Mylonas, Matthew Tarduno, Shosh Vasserman, ReedWalker, Zi Yang Kang, and seminar participants at Harvard University, University of California, Berkeley, and the University of Southern California. We also thank those at Metropia who helped to implement the field experiment, especially Yi-Chang Chiu and Ali Arian. The usual caveat applies. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by Robert W. Hahn, Robert D. Metcalfe, and Eddy Tam. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Welfare Estimates of Shifting Peak Travel
Robert W. Hahn, Robert D. Metcalfe, and Eddy Tam
NBER Working Paper No. 31629
August 2023
JEL No. H20,R41

ABSTRACT

We develop novel estimates of peak and off-peak price elasticities for urban mass transit demand in San Francisco using a large natural experiment with 3.6 million trip sessions and a natural field experiment that both have exogenous price subsidies. We then estimate the welfare impacts for these price subsidies using a sufficient statistics approach. Our analysis suggests that off-peak subsidies can increase welfare, but the positive effects are reduced when consumers take the decisions of others into account compared to when they do not. We also find a large variation in the welfare impacts of shifting travel to different periods, which is explained by differences in demand and congestion characteristics. Finally, we show that the targeting of subsidies can increase welfare, but need not do so if the regulator does not have accurate information on demand.

Robert W. Hahn
University of Oxford
and Carnegie-Mellon University
robert.hahn@smithschool.ox.ac.uk

Eddy Tam
Kings College London
London
United Kingdom
eddy.hftam@gmail.com

Robert D. Metcalfe
Department of Economics
University of Southern California
Los Angeles, CA 90007
and NBER
robert.metcalfe@usc.edu

1 Introduction

There are many problems in which economic welfare may be increased through shifting demand across peak and off-peak periods (Boiteux, 1960). Applications include automobile travel, mass transit, electricity, energy and telecommunications. In this paper, we focus on mass transit. The general problem we address is estimating the economic welfare impacts of a price change on consumption in the peak and off-peak periods. Our welfare estimates rely on well-identified demand elasticity estimates within and across time periods. We also allow for prices to differ from marginal cost and for externalities in consumption, such as congestion and pollution.

Using a large natural experiment and a natural field experiment, we develop novel estimates of peak and off-peak price elasticities for urban mass transit demand in San Francisco. The experiments were conducted by the Bay Area Rapid Transit (BART) in California.¹ Both experiments allow us to obtain causal estimates of the price elasticities of demand. The first, which includes 17,500 BART riders (7% of BART riders) and over 3.6 million trip sessions, is a natural experiment with a price change in the off-peak period for six months. Riders were staggered into the pricing program over the course of six weeks. Prices were reduced an hour before the peak period (630am-730am) and an hour after the peak period (830am-930am).

The second is a natural field experiment with randomized prices across 1,900 riders over four months (i.e., 77,000 trip sessions). The experiment incentivized treated customers to shift their riding demand into the least congested train within a 40 minute time window from their usual departure time by lowering prices in the new time window. With detailed consumer level data on demand and individual characteristics, both of these experiments yield own and cross price elasticities, which are used in the welfare analysis. Across both experiments, we have train-level data on the weight of the train and the time the consumer enters and leaves a station, allowing us to measure congestion. We estimate the welfare impacts of the price subsidies in the experiments using a sufficient statistics approach (Chetty, 2009).

The subsidies in both experiments influenced consumer choice in the expected direction. They reduced travel in the period that did not receive the subsidy, and increased travel in the period that received the subsidy. We estimated an own price elasticity for the period in which consumers received the subsidy and a cross-price elasticity for the period in which they did not. The own and cross price elasticities are around -0.86 and 0.44 respectively in the natural experiment and -0.94 and 0.54 respectively in the field experiment. These estimates suggest

¹BART is a rail and subway system that serves the greater San Francisco area. It is the urban public transit system that Daniel McFadden used to estimate various discrete choice models that contributed to his Nobel prize work (McFadden, 1974; Domencich and McFadden, 1975; McFadden et al., 1977; McFadden, Tye and Train, 1977; McFadden, 2001).

some departure time flexibility for BART users. Other non-experimental elasticity estimates for the BART system as a whole are comparable to our own elasticity for off-peak demand (See, e.g., [McFadden \(1974\)](#)). But our own elasticity estimate is a little higher in absolute value than the -0.6 to -0.75 range in the review of estimates in the US by [Litman \(2004\)](#) and [Holmgren \(2007\)](#), and those estimated in Mexico by [Davis \(2021\)](#).²

Our two experiments offer different insights into the impact of providing subsidies at particular times for economic welfare. We find that the subsidy in the natural experiment generally increases welfare. In the base case, we show that the marginal value of public funds (MVPF) is 1.6 in trying to shift travel to the off-peak period from the peak period (730am-830am).³ This means that for each additional dollar of net cost to the government, the marginal benefit to all parties is \$1.60. We also calculate the net benefits per dollar of subsidy in the natural experiment, and find that the net benefits are \$0.36 per dollar of subsidy expenditure.⁴

In the field experiment, we had price variation across more time periods and show that the welfare benefits varied dramatically across routes and time periods due to differences in demand and congestion characteristics. In particular, we show that net benefits per person in the field experiment are sometimes less than zero. We also show that some MVPFs in the field experiment are less than one, which means that the willingness to pay associated with the policy is less than the net cost to the government.⁵

We do several sensitivity analyses to examine the robustness of our main findings. In particular, we examine three sets of behavioral assumptions about consumer response to see how they would change our welfare results. The first relates to how welfare results would change if we scaled up the subsidy to all BART customers, but allowed for some customers to be less responsive to the subsidy than others, perhaps because of selection effects. The second relates to measuring the impact on consumer welfare under the assumption that consumers take the full impact of others' behavior into account in making their transit choices—a case we refer to as endogenous congestion.⁶ In both the scaling and endogenous congestion simulations, the effect on per capita welfare is to reduce the net benefits from the subsidy by up to 47%. The third assumption relates to the level and shape of the congestion damage function. We find

²We have not found other well-identified *cross-price* elasticity estimates for peak and off-peak in public transit, and thus do not provide a comparison

³See [Hendren \(2020\)](#) and [Finkelstein and Hendren \(2020\)](#) for excellent reviews of this literature. Our MVPF of 1.6 is quite favorable in comparison to the marginal cost of public funds of around 0.8 to 1 through changes in the linear tax rate ([Hendren, 2020](#)).

⁴The net benefits measure we use in our welfare model is based on total net benefits for a discrete change in the subsidy, and is in dollars. In contrast, the MVPF is a marginal measure that is dimensionless. Both can be useful in comparing policies.

⁵We also do a calculation on the optimal subsidy and find that the optimal subsidy may be substantially higher than the actual subsidy based on the natural experiment.

⁶This concept could apply to any externality where consumers consider the impact of others' behavior in their decisions.

that for our parameter estimates the shape of the damage function does not appear to have a large effect on our welfare results.

A final set of findings on welfare relates to how subsidies should be targeted. We examined this issue in two ways. First, using our natural experiment, we examined how net benefits per person might change if the most congested route were targeted. We found that targeting users of the most congested route increases average net benefits by 15%. Second, our entire natural field experiment is focused on the issue of targeting. While the field experiment shows that targeting can be effective in shifting demand, the main takeaway is that these shifts may not always increase welfare. A key reason that targeting does not necessarily increase welfare is because the machine learning algorithm used for targeting in the field experiment only takes into account information on crowding. It does not take into account information on unobserved demand characteristics, which are also central for determining the welfare impact of the subsidy. Our experiments suggest that targeting may require a substantial amount of information on economic fundamentals, such as demand, to increase the likelihood that welfare will actually increase with targeting.

Our paper relates to three important research areas. First, our empirical analysis builds on work in sufficient statistics. Our paper is most closely related to [Jacobsen et al. \(2020\)](#). These authors examine how sufficient statistics can be applied to externality-correcting policies; however, they assume markets are perfectly competitive and that consumers take the externality as given, and do not adjust their behavior in response to the level of the externality. We relax those assumptions in [Jacobsen et al. \(2020\)](#), but retain the assumption that consumers are fully informed.⁷ Our model also relates to work by [Kreindler \(2022\)](#) on estimating endogenous congestion, but our model differs in terms of its focus on urban mass transit and its theoretical underpinnings.

Second, we build on a large literature on congestion and the welfare effects of peak pricing schemes. It has long been argued that many sectors could benefit from peak-load pricing ([Vickrey, 1963](#); [Mohring, 1972](#); [Glaister, 1974](#); [Train, 1977](#); [Glaister and Lewis, 1978](#); [Winston, 1985](#); [Jansson, 1993](#); [Parry and Small, 2009](#)). There have been attempts to estimate the welfare implications of changes in peak or off-peak prices for mass transit (*e.g.*, [Parry and Small \(2009\)](#); [Joskow and Wolfram \(2012\)](#); [Knittel and Sandler \(2013\)](#); [Winston \(2013\)](#);

⁷We do not address issues of estimating welfare with adverse selection, information asymmetries, market structure issues and behavioral agents. For examples of such work, see [Bundorf, Levin and Mahoney \(2012\)](#); [Hendren \(2013\)](#); [Weyl and Fabinger \(2013\)](#); [Mahoney and Weyl \(2017\)](#); [Farhi and Gabaix \(2020\)](#); [Ito, Ida and Tanaka \(2021\)](#). Several studies that focus on externalities and internalities are also related to our work, including [Allcott, Mullainathan and Taubinsky \(2014\)](#); [Piketty, Saez and Stantcheva \(2014\)](#); [Allcott and Taubinsky \(2015\)](#); [Allcott, Lockwood and Taubinsky \(2019\)](#). These studies fit within the general framework developed by [Kleven \(2021\)](#).

Levin, Lewis and Wolak (2017); Kreindler (2022); Tarduno (2022)).⁸ Moreover, there appears to be a dearth of well-identified estimates of demand elasticities and cross elasticities for mass transit that have been used to estimate welfare effects. For instance, Parry and Small (2009) estimate the welfare effects of urban transit subsidies using a structural approach, but note that “[l]ittle information is available about shifts of transit riders across time periods” (p. 17). Our empirical analysis helps to address this gap in the literature.⁹

Third, we add to the literature on optimal targeting. Several papers examine how targeting price changes or interventions could increase welfare (Rodrik, 1987; Chassang et al., 2012; Dupas, 2014; Allcott and Taubinsky, 2015; Cohen, Dupas and Schaner, 2015; Alatas et al., 2016; Byrne, Martin and Nah, 2019; Polyakova and Ryan, 2019; Farhi and Gabaix, 2020; Gerarden and Yang, 2021). Our field experiment targeted subsidies using detailed data on congestion and historical consumer behavior, which is typical of many machine learning algorithms in transport (Tizghadam et al., 2019) and other markets (Mullainathan and Spiess, 2017; Athey, 2018; Knittel and Stolper, 2019; Burlig et al., 2020; Christensen et al., 2021; Aiken et al., 2022). However, not having information on price elasticities of demand for travel and consumers’ willingness to shift their travel times can limit the effectiveness of targeting subsidies using machine learning. For instance, we show in some time periods that targeted pricing using a machine learning algorithm actually lowered welfare.

The paper is structured as follows: Section 2 provides an overview of the theory that allows us to estimate welfare impacts of the experiment. Section 3 develops demand estimates for the natural experiment and section 4 develops demand estimates for the field experiment. Section 5 implements our welfare analysis. Finally, Section 6 concludes.

2 Overview of the Theory

We develop a theoretical model that allows us to estimate the welfare effects of a subsidy using sufficient statistics. We present the details of our welfare model in appendix A.1. The model identifies key demand parameters of interest that are obtained from our experiments. In

⁸There have been many papers that analyze the impact of peak-time pricing on overall demand and welfare (Borenstein, 2005; Wolak, 2007; Joskow and Wolfram, 2012; Jessoe and Rapson, 2014; Ito, Ida and Tanaka, 2018; Burkhardt, Gillingham and Kopalle, 2019; Fowlie et al., 2021; Williams, 2021). Our paper builds on this work by developing an analytically tractable approach for estimating welfare changes of peak and off-peak pricing in the presence of externalities, where the externalities can affect demand.

⁹Other papers that estimate changes in subway demand with a change in prices include Davis (2021), who estimates the price elasticity of demand for subways in Mexico. Davis’ analysis differs from ours in that he focuses on the impact of general price decreases in subway systems, and not on differences in peak and off-peak travel aimed at reducing an externality. Another paper by Yang and Long Lim (2018), estimates the impact of temporary free rides on early morning subway demand in Singapore, and finds a meaningful shift to earlier trains.

addition we consider whether consumers take the behavior of others into account in making their decisions.

The basic model has two types of agents, identical consumers and a firm (e.g., a utility). A representative consumer maximizes utility over peak and off-peak travel, and a numeraire good. We assume that consumer has a well behaved utility function and quasi-linear preferences. In addition to the consumers, a regulated utility is assumed to minimize the cost of meeting demand. Consumer choices generate an externality that affects the utility, which can be thought of as congestion or delay in our application.

We allow prices to deviate from their marginal private cost (which is consistent with pricing in many regulated markets), and we allow for subsidizing particular consumers. The price subsidy is aimed at encouraging off-peak travel in our first experiment (experiment 1), and shifting particular individuals into less crowded trains in our second experiment (experiment 2). The primary aim of the subsidy is to reduce an externality, which can be thought of as crowding on a subway or delay. In our application, we consider both.

The key results from this model are intuitive. The welfare effect of a subsidy depends on how consumers respond to the subsidy, and the extent to which prices deviate from the marginal social cost. If the price is above the marginal social cost in a particular period, and travel increases as a result of the policy intervention (e.g., a subsidy), then it contributes to an increase in welfare. Summing the welfare changes across each period gives the overall change in welfare.

The equation we use to estimate the per capita welfare effect of an off-peak subsidy is given by:

$$\frac{1}{N}(W(s') - W(0)) = (p - MSC_1)\frac{dx_1}{ds}s' + (p - \frac{1}{2}s' - MSC_2)\frac{dx_2}{ds}s' \quad (1)$$

where N is the number of consumers; $W(s')$ is the economic welfare of a subsidy, s' ; and $W(0)$ is the welfare associated with no subsidy.¹⁰ $\frac{dx_1}{ds}$ and $\frac{dx_2}{ds}$ are the demand responses to the subsidy in the peak and off-peak periods, respectively. These are assumed to be constant. The price of the good is p , assumed to be constant in both periods. The marginal social cost in periods 1 and 2 are MSC_1 and MSC_2 , respectively. Marginal social cost is the sum of the marginal private cost (MPC), which is assumed to be constant across periods, and the marginal external cost, MEC_1 and MEC_2 , which can vary across periods.

The derivation of equation (1) makes use of the envelope theorem for the consumer maximization problem. When the off-peak subsidy increases, consumers may alter their travel

¹⁰We define the change in welfare as the change in producer plus consumer surplus, assuming the subsidy is a lump sum transfer. See appendix A.1 for details.

demand because of an equilibrium change in the subsidy, the price or the congestion level, but the resulting changes in consumption do not have a direct impact on consumer utility. This is consistent with the sufficient statistics approach ([Harberger, 1964](#); [Chetty, 2009](#); [Jacobsen et al., 2020](#)).

An important issue for welfare is how consumers respond to the subsidy. We consider the two polar cases of exogenous congestion and endogenous congestion. Exogenous congestion assumes that consumers ignore the impact that the subsidy has on the level of congestion. Endogenous congestion assumes consumers take into account how others will respond to the externality. The two cases can provide insights into the welfare impacts of how interventions scale. The exogenous congestion case may be more appropriate when estimating the welfare impacts of a small-scale subsidy, such as a pilot. The endogenous congestion case may be more appropriate when measuring the welfare impacts of a relatively large-scale intervention, such as providing a subsidy to all customers.¹¹

These two cases will typically result in different values for demand, as measured by $\frac{dx_1}{ds}$ and $\frac{dx_2}{ds}$. Empirical estimates for these cases are presented in section 5. The theoretical relationship between the endogenous case and the exogenous case is discussed in appendix A.16.

3 Natural Experiment

To estimate the welfare impacts of the subsidy, we will need causal estimates of the impact of a subsidy on peak and off-peak travel. This section explains how we derive these estimates for the natural experiment. We provide a brief overview of the BART network (3.1), a description of the price subsidy and natural experiment (3.2), a discussion of our identification strategy (3.3), and present the key results on peak and off-peak demand (3.4).

3.1 The BART network

The natural experiment was conducted with BART, which is a large public transit system serving the San Francisco Bay area in California.¹² BART is the fifth-busiest heavy rail rapid

¹¹This issue applies in many settings, as [List \(2022\)](#) notes.

¹²The heavy rail elevated and subway system connects San Francisco and Oakland with Alameda, Contra Costa, and San Mateo counties across 112 miles of track connecting 46 stations. BART has an average of 423,000 weekday passengers and 124.2 million annual passengers in fiscal year 2017 (BART, 2017),

transit system in the United States. Figure A8 in Appendix B shows the network map.¹³ The prices on BART are comparable to those of other U.S. commuter rail systems and are higher than those of most subways, especially for long trips. The minimum price is \$1.95 (except for San Mateo County trips) under 6 miles (9.7 km), and the maximum one-way price including all possible surcharges is \$15.70, the journey between San Francisco International Airport and Oakland International Airport (2017 prices).¹⁴ Figure A9 in Appendix B plots the distribution of fare by routes. The mean fare for non-airport routes is \$4.16, with a median fare of \$4.25. The mean and median fare for airport route is \$9.60.

3.2 The Perks Treatment

The treatment changed the relative peak to off-peak prices, and Bart called this the *Perks program*. To be eligible for the new pricing program, customers had to have a Clipper ID number.¹⁵ We used Clipper card data for our study. The card records the amount paid, the time the person entered the departing station, and the time they left the arrival station. Linking each user with their Clipper card was necessary to give them the pricing subsidy. This subsidy was based on the frequency, timing, and length of their trips.

The Perks program was advertised over a four week period through direct outreach at stations, media coverage, and employer partnerships. There was some natural time variation in some stations receiving ads before others, and we use this staggered variation later in this section. An example of the ads for direct outreach is shown in Figure A10 in Appendix B. The left hand side of the figure is an example of the basic ad. The right hand side provides a fuller description of the program with eligibility criteria. The program started on August 23, 2016 and ran until the end of February 2017.

The price subsidy changes based on the time the journey starts. For every mile traveled between 630am and 730am or 830am and 930am, the consumer would get a 11.25% reduction in the off-peak price. The peak time at BART was between 730am and 830am. As part of the program to incentivize people to take part, it also gave a modest 2.25% reduction in the peak

¹³BART has five rapid transit lines; most of each line's length is on track shared with other lines. Trains on each line run approximately every 15 minutes on weekdays and 20 minutes during evenings, weekends and holidays; stations on the section of track between Daly City and West Oakland are served by four lines and therefore have 16 trains an hour on each track. BART service begins around 4:00 am on weekdays, 6:00 am on Saturdays, and 8:00 am on Sundays.

¹⁴The price is based on a formula that takes into account both the length and speed of the trip. A surcharge is added for trips traveling through the Transbay Tube, to Oakland International Airport, to San Francisco International Airport, and/or through San Mateo County, a county that is not a member of the San Francisco Bay Area Rapid Transit District. The farthest possible trip, from Pittsburg/Bay Point to Millbrae, costs less because of the \$4 additional charge added to SFO trips and \$6 additional charge added to Oakland trips.

¹⁵Most regular BART customers have a Clipper card, which is a contactless smart card accepted on all major Bay Area public transit agencies, and may be used in lieu of a paper ticket.

price. A relative price difference between peak and off-peak price of 9% was thus created by the Perks program. That is the exogenous change in prices that we will leverage to estimate demand. To receive the subsidy, participants needed to have an active PayPal account that used the same email address as their BART Perks account.

3.3 Identification

Identification of $\frac{dx_1}{ds}$ and $\frac{dx_2}{ds}$ comes directly from this natural experiment. We obtain identification of peak shifting due to changes in off-peak prices using two strategies. First, we compare the self-selected Perks users with the rest of the network using a difference-in-differences framework (we have data on all individual rides on BART before and during the natural experiment). We can determine the share of rides that are in the peak and off-peak hours and use this to obtain a causal estimate of the Perks program on demand. We test for parallel trends in peak and off-peak shares before the Perks program began, and find that the peak and off-peak trends for the Perks and non-Perks users are statistically indistinguishable.

Second, we use the fact that not all BART customers were brought onto Perks at the same time. As mentioned above, some stations advertised the subsidy program before others, so we have quasi-random variation in when people saw the advertisements to the Perks program. The program officially started on August 23th, 2016, but 18.9% of the sample were enrolled in the program after one month from the start of the program. This is our ‘late’ sample. Comparing ‘early’ versus ‘late’ enrollers allows us to estimate the impact of Perks on the number of BART trips, as opposed to just the shares of peak and off-peak travel. Again, we find that the trends for early and late enrollers are quite similar prior to the treatment starting (i.e., they are statistically indistinguishable from each other). We estimate the cross-price elasticity using a difference-in-differences approaches to estimate our welfare model.

We have individual level trip data for the 17,545 BART users who signed up for Perks for the six months before the Perks program began (February 2016) and six months during the Perks program (up to February 2017). This is a total of 2,119,588 individual journeys on the BART network for twelve months. We also have data on the rest of the BART network at 15 minute intervals for the same time period which corresponds to 72,098,160 individual rides for non-Perks customers.

Because prices were not perfectly randomized across BART users, a potential issue with our causal estimates could arise if those who signed up for the program knew that they were more likely to reduce their peak demand during the program. While we cannot rule out this possibility, the near perfect parallel trends in both identification strategies suggest that this is unlikely. To address any concerns, we provide an external check on the results obtained

here with our field experiment described in section 4, which randomizes prices across BART consumers. Importantly, we will compare the elasticities from the field experiment to those generated from the quasi-experiments considered here. We show that they are similar.

3.4 Results


This results section is based on the two identification strategies presented above. Before we present the results, we describe the data of the Perks users and how it compares to the BART non-Perks users. For the rest of the customers using BART, we know what rides (time, day, stations) have been taken in the same time frame as our opt-in sample, but we can not identify individual accounts.

We first compare the distribution of travel demand for our 17,545 BART users in the Perks program with the demand from non-Perks users. Figure 1 below shows the raw distribution of rides taking place in five minute intervals before and during the Perks program. Panel (a) shows demand before the Perks program and panel (b) shows demand during the Perks program. For both panels (a) and (b), the y-axis is the share of trips. The figure reveals that the overall shape of the dashed line for the non-Perks sample is very similar in panels (a) and (b). In contrast, the line for the Perks sample appears to be steeper and higher for the shoulder hours in the morning (i.e., the subsidized hours). This provides some visual evidence that the increase in demand in the shoulder hours from the Perks commuters during the Perks program does not happen with the rest of the BART network customers during the Perks program. This evidence suggests that once customers received the new price changes, they shifted some of their demand to the shoulder off-peak hours from out the peak hour.

3.4.1 Impact of Prices on Demand: Empirical Approach 1

For our first identification approach in comparing the Perks users to the rest of the BART network users, we must ensure that we have parallel trends in the peak and off-peak shoulder hours for the Perks BART users and the non-Perks BART users before the experiment began. This is the necessary condition for our difference-in-differences approach. Figure 2 presents the daily share of peak hour weekday journeys for the Perks and non-Perks users before and during the experiment. This figure shows that there is a very similar trend in the raw peak hour demand data for the Perks users (black line) and the non-Perks users (the gray line) before the Perks program started on August 23, 2018. The trends for the two groups are not statistically significantly different (see the econometric test of pre-trends in peak and off-peak demand in the Perks and non-Perks samples in Figure A12). This allows us to estimate the

Figure 1: Comparing Perks program participants with the rest of the BART network



Note: The figure plots the average daily share of trips taken by the participants in the first experiment in each time interval from 530am-1130am. It also plots on the same figure the average share of daily trips at each time interval taken by the rest of the BART users. The sample for Panel (a) includes April 1, 2016 to August 22, 2017 for all weekdays. The sample for Panel (b) includes August 23, 2016 to February 28, 2017 for all weekdays.

impact of the price change on peak demand using a difference-in-differences framework.


It is also clear from Figure 2 that there is a very similar trend in the shoulder off-peak hour demand data between the Perks users (black line) and other users (the gray line) before the Perks program started. The trends for the two groups are not statistically significantly different before July 2016, as suggested in Figure 2 (see the test of pre-trend in Figure A12). This allows us to estimate the impact of the price change on shoulder demand using a difference-in-differences framework.

Given these two parallel trends in peak and off-peak shoulder hours before the Perks program started for Perks and non-Perks users, we can now estimate the following panel fixed effects model:

$$y_{gt} = \beta_1 * Perks_t + \beta_2 * Treated_g * Perks_t + \beta_3 Treated_g + \gamma_t + \epsilon_{it} \quad (2)$$

where y_{gt} is the share of peak hour (730am - 830am) trips in the morning from 5:30am-10:30am on date t for group g . $Treated_g$ is an indicator for Perks users with the omitted group being non-Perks users. $Perks_t$ is an indicator for the period in which the Perks program is implemented from August 23, 2016 to February 28, 2017. γ_t is a date fixed effect that controls for unobserved changes in travel patterns such as a date-specific service disruption.

Figure 2: Share of peak and off-peak shoulder hour demand for Perks users versus the rest of the BART network users (non-Perks)



Note: Panel (a) shows the five week moving average of the share of peak-hour demand (730am-830am) among all trips from 5.30am-10.30am for Perks and non-Perks users. Panel (b) shows the moving average of the share of shoulder off-peak demand (630am-730am and 830am-930am) among all trips from 5.30am-10.30am for Perks and non-Perks users. The sample includes weekdays from April 2016 to March 2017. The vertical line represents the beginning of the first experiment on August 23, 2016.

Our sample period covers April 1, 2016, to February 28, 2017. It covers five months before the Perks program started and a six month period when the Perks program was running. Our null hypothesis is that the difference-in-difference estimator, β_2 , is equal to zero for both the peak period and the off-peak shoulder hour.

Table 1 presents the estimates of equation 2. Column (1) reports estimates on the demand share of peak hour. We find that comparing with other users, Perks users reduce their share of peak hour trips by 2.1 percentage points during the program period. Column (2) reports estimates of the share of off-peak shoulder hours. We find that using non-Perks users as the reference group, Perks users have a 2.6 percentage point increase in the share of off-peak shoulder hour trips during the program period compared to before.

We also estimate an alternative specification that controls for differences between Perks and non-Perks users in terms of the routes they take (defined by the entry and exit station). We compute the weekly share of peak hour (and off-peak shoulder hour) trips for Perks users and for non-Perks users for each individual route in our data. This specification allows for some route has sparse number of rides at daily level. We then estimate equation 2 using the weekly share of peak hour (and off-peak shoulder hour) trips for each route and group as outcome, controlling for route-group fixed effects (i.e., the route-specific travel pattern for Perks and

Table 1: Difference-in-difference estimates - Participants and rest of network

	Outcome: Share of all trips in morning	
	Peak Hour (1)	All dates Off-peak Shoulder Hour (2)
Treated users $\times$ Perks Period (β_2)	-0.021*** (0.0029)	0.026*** (0.0023)
Treated users (β_3)	-0.0030*** (0.0008)	0.0744*** (0.0016)
Observations	466	466

Note: Includes daily observations for Perks users and non-Perks users. Sample includes weekdays from April 1, 2016 to February 28, 2017, excluding public holidays. Standard errors in parentheses are clustered by week.

non-Perks users). Table A13 reports the estimates. We find very similar effects as those shown in Table 1.

We next analyze whether the effect of the subsidy applies to the two busiest stations in the BART network, Montgomery Street and Embarcadero. In Table A15 we split the routes into two directions - those that travel from the east of the two stations, thereby crossing the Transbay tube, and those that travel from the west of the two stations. We find that the treatment effect in shifting riders from the peak to the off-peak hour is much larger for the congested routes travelling from the east of the Transbay exiting at Montgomery Street and Embarcadero, compared to those travelling from the west that are less congested.

We also examine the change in demand for peak and off-peak demand when the Perks incentives are removed, which allows us to understand if the incentives created a persistent change in behavior. After six months of the price subsidy program, the program ended. We have BART demand data for those that signed up versus those that did not for six months after the program concluded. We find that while the peak to off-peak shift is lower after the program, there is a persistent and significant treatment effect. We find that the long-term change in peak hour demand without the subsidy was -0.9% (see appendix Table A12). One possible explanation is that people change their demand during the experiment, and once the price changes have been removed, it is costly for users to change their habits or schedules to revert back to their pre-treatment levels. Another explanation is that people may find better ways to commute that increase their welfare, which is supported by evidence from London and Singapore (Larcom, Rauch and Willems, 2017; Yang and Long Lim, 2018).

Despite parallel trends prior to Perks, the customers who selected into Perks may have known prior to Perks that they might want to change their early morning commute time in the future.

This is impossible to test with our data but given how strong the parallel trends are, we find this outcome is unlikely. Another caveat here is that we only observe shares of demand and not absolute demand (i.e., number of trips). This is where the next section and identification strategy helps us.

3.4.2 Impact of Prices on Demand: Empirical Approach 2

For the second empirical approach, we estimated how the Perks price program changed demand in the first month in the program, when 18.9% of users had not already signed up (through the staggered roll-out of the advertisement). We define 'early' enrolled Perks users as those who enrolled in Perks from August 23 to September 2, 2016, where our 'late' control group consist of users who eventually enrolled between Oct 1 to Nov 5, 2016. Therefore in our treatment time period (August 23 to September 31, 2016), all the late enrolled Perks users would have not enrolled in the program yet and thus constitute an effective control group for the travel behavior in the absence of the Perks program.¹⁶

A benefit of using early versus late Perks users is that we can analyze absolute changes in demand at the individual level as a result of the price change. Moreover, both early and late customers all enrolled into the price change, so unobserved characteristics are even less likely to drive the change in peak demand. In addition, the approach of comparing an early treatment group to a late control group is suggested by [Callaway and Sant'Anna \(2021\)](#) as a way to overcome some of the issues in two-way fixed effects models. Figure A6 shows the cumulative distribution for BART enrollment as a function of time.

We can analyze whether the pre-Perks trend before the Perks program is the same for the early and late enrollers. Figure 3 suggests that the trend on peak and shoulder hours before the Perks program is the same. In Figure A13 we test for the difference of the pre-trend by each of the week, we find that for peak hour trips, the pre-trend for the two groups are not statistically significant. For off-peak shoulder hour trips, we find no systematic difference in the trend for the two groups as well.


We can now estimate the following equation:

$$y_{it} = \beta_1 * E_i * Perks_t + \sigma_i + \mu_t + \epsilon_{it} \quad (3)$$

where y_{it} measures the number of trips during Peak hour on date t for user i . E_i is an indicator for individual i being enrolled early from August 23 to September 2, 2016. $Perks_t$ is an

¹⁶Our analysis is robust to the time period selected for early versus late. We present in Appendix B Table A14 results that define late enrolled users as those enrolled between Oct 15 to Nov 5.

Figure 3: Early and late enrolled



Note: The sample includes trips from July 12, 2016 through September 30, 2016. The figure plots the two week moving average for off-peak shoulder hour trips (630am-730am and 830-930am) and peak hour trips (730am-830am) of early enrolled and late enrolled users. The vertical line represents the date August 23, 2016, when the Perks program started.

indicator for the perks program period. σ_i is an individual fixed effect and μ_t is a date fixed effect. Our null hypothesis is that the difference-in-difference estimator, β_1 , is equal to zero in both the peak hour and the off-peak shoulder hour.

Table 2 presents the estimates for equation 3. Columns (1)-(2) report the estimates on the number of trips per day in the time window of peak hour and off-peak hours respectively. We find that compared with the late enrolled users, during the first month of the program the daily number of trips during the peak hour by early enrollers decreases by 0.007; for the off-peak shoulder hours the daily number of trips by early enrollers increases by 0.015. Column (3) reports the estimates on the total number of trips in the morning.¹⁷ We find that the Perks program does not increase the total number of trips taken by early enrollers.

To convert these numbers to percentages, we need to know the baseline number of trips. The baseline number of trips for the late enrollers during September is 0.178 and 0.273 during the peak and shoulder hours respectively. The DiD estimate of -0.00705 for the peak hour represents a 4% decrease in peak demand. The DiD estimate of 0.0151 for the off-peak shoulder hours represents a 5.5% increase in off-peak shoulder hour demand.

¹⁷The estimates correspond closely to the estimates in share in Table 1. The share of peak (off-peak) trips implied among all trips from 5.30am-10.30am is 0.303 (0.535), with reduction of trips of -0.007 (increase of 0.0151), and the share of peak (off-peak) trip during the program would be 0.284 (0.566), which is an implied change of shares of -0.019 (0.0304).

Table 2: Average treatment effects by time period: Early vs. late comparison

	OLS		
	Peak	Off-peak shoulder (Either)	Total (5.30-10.30am)
	(1)	(2)	(3)
Early enrolled users $\times$ Perks Period (β_1)	-0.0071** (0.0031)	0.0151*** (0.0038)	0.0040 (0.0044)
Observations	600148	600148	600148
Date FE	x	x	x
User FE	x	x	x
Sample mean (before Perks)	0.18	0.31	0.57

Note: The sample includes daily trips from June 1 to October 1, 2016, excluding August 23 to September 2. The Perks period is defined as dates after September 2. Early enrollers are defined as those who enrolled between August 23 and September 2. The control group includes those who enrolled between October 1 and November 5, 2016. Standard errors are clustered at the user-week level. The sample means are for late enrolled users in periods before the Perks period.

To check whether the treatment effect is robust to different specifications, we relax the specific definition of early versus late enrollment. Instead, we make use of the fact that the activation time into treatment for each user differs. We estimate the change in travel behavior of each user before and after their Perks program activation using the following equation:

$$y_{it} = \eta_1 \text{PerksActive}_{it} + \mu_t + \sigma_i + \epsilon_{it} \quad (4)$$

where PerksActive_{it} is an indicator that turns on when user i activated their Perks status on date t . σ_i are user fixed effects and μ_t are time fixed effects. As above, our null hypothesis is that the difference-in-difference estimator, η_1 , is equal to zero in both the peak hour and the off-peak shoulder hours.

Table 3 presents the estimates for equation 4 and shows that the effects of the price subsidies on peak and off-peak demand are very similar to our staggered specification. We find that the after activating their Perks user status, peak hour journeys decreased by 0.007 per day for each user and bonus hour journeys increased by 0.024. The latter result is significantly larger than that from the early versus late enrolment specification. We also estimate the treatment effect following [De Chaisemartin and d'Haultfoeuille \(2020\)](#) and [Callaway and Sant'Anna \(2021\)](#) and report the estimates in Appendix Table A18. Using these approaches does not alter our estimates, providing more confidence in our identification strategy.

Table 3: Average treatment effects by time period: staggered activation

	OLS		
	Outcome: Number of daily trips		
	Peak (1)	Off-peak shoulder (Either) (2)	Total (5.30-10.30am) (3)
Active status level	-0.00713*** (0.00221)	0.0244*** (0.00389)	0.0130*** (0.00454)
Observations	1426285	1426285	1426285
Date FE	x	x	x
User FE	x	x	x
Sample mean (before Perks)	0.17	0.30	0.56

Note: Sample includes dates from July 1, 2016 through November 6, 2016. Standard errors are clustered by users and dates.

3.5 Own and cross-price elasticity of demand for peak and off-peak travel

Given our estimates above, we can now estimate the own and cross-price elasticity of demand for peak and off-peak travel. We estimate $\hat{\eta}$ from equation 4 from estimating the demand effects in our two difference-in-difference specifications. We estimate the causal impact of changing relative prices on the demand for rides taken per day in peak or off-peak periods (we denote η_j as the estimates for the equation where the outcome is travel in period j).

We define the relative price of off-peak, p_2 , as the ratio of price per mile of off-peak to that of peak. The estimated marginal price of a mile demanded on BART is \$0.112 in the control condition, where there is no difference between peak or off-peak prices. With no price subsidy, this relative price of off-peak is $p_2 = \frac{0.112}{0.112} = 1$, and it applies to the control group. For the treatment group, the price of one mile on BART is reduced by 9.2%. We calculate the average off-peak shoulder hour trip taken per day by rider before the program as $\bar{x}_2$. Our estimate on own-price elasticity of off-peak hour equals to $\epsilon_{22} = \frac{p_2}{\Delta p_2} \frac{\eta_2}{\bar{x}_2}$. Our estimate on cross-price elasticity of peak hour demand to off-peak hour price equals to $\epsilon_{12} = \frac{p_2}{\Delta p_2} \frac{\eta_1}{\bar{x}_1}$. Our estimates for own-price elasticity of off-peak demand is -0.86, and the cross-price elasticity of peak demand with respect to off-peak price is 0.44.¹⁸ We find that this own-price elasticity does not meaningfully vary by whether we examine consumers who shift to the early off-peak hour as opposed to the later off-peak hour (see appendix Table A16).

¹⁸The program offered 5 points per mile of off-peak trip travelled. Each point has monetary value of \$2.52/1000, therefore the subsidy provides a monetary value of $5 \times 2.52/1000 = 0.0126$ per mile in off-peak. The subsidy also gives 1 mile per peak trip travelled. This implies a new relative price of off-peak as $\frac{0.112 - 0.0126}{0.112 - 0.00252} = 0.9079$. Therefore, the relative price change is $\frac{\Delta p_2}{p_2} = -0.0921$. From Table 3, the estimate for η_2 is 0.0244. The average off-peak trip per day is $\bar{x}_2$ is 0.307, therefore, $\epsilon_{22} = -0.86$ (s.e. = 0.137). As $\bar{x}_1 = 0.176$, and η_1 from Table 3 is -0.00713, $\epsilon_{12} = 0.44$ (s.e. = 0.136).

We find that the cross-price elasticities of demand change in ways that one would think based on economic theory. For example, the closer the consumer typically travels to the off-peak hour before the subsidy, the greater their own-price response (see appendix Figure A14).

We also find that the more congested the network, the greater the cross-price elasticity (see Table A15 in the appendix). For example, we examine the shift for any route exiting at Montgomery Street and Embarcadero (the two most congested stations of the BART network). Demand from the east (i.e., eastbound Transbay tube trips that are in the busy corridor) towards these two stations shift more (change in absolute share) than that from the west (i.e., westbound transbay tube trips). While all of these heterogeneous effects go in the direction that we expect, they are still correlations and not the causal estimates of changing these parameters on the cross-price elasticities of demand. Nevertheless, we can estimate welfare changes with these different samples in the BART network.

4 Field Experiment

In this section, we estimate the change in peak and off-peak demand from the field experiment. We use these demand estimates to develop a causal estimate of the impact of the subsidy on welfare in the next section. Here, we provide a brief overview of the field experiment and a discussion of our identification strategy, and then present the key results on demand. We then compare the estimates with those obtained in experiment 1.

4.1 Identification

The identification of price elasticities in this field experiment comes from different variation than in the natural experiment. First, we control the assignment mechanism of price changes in this field experiment, whereas we did not control assignment in the natural experiment. The clear advantage of this experiment was the ability to randomize some BART customers into the control group where prices did not change and some BART customers into the treatment group where prices were changed.

Second, in contrast to the natural experiment, which focused on generally shifting riders from peak (730am-830am) to off-peak (630am-730am and 830am-930am) travel, this field experiment focuses on shifting riders to less crowded trains in time periods adjacent to their normal commute. For example, a rider may receive a subsidy to shift their travel from their usual time period to an adjacent twenty minute interval, when a less crowded train is available at their departing station.

Third, in comparison with the natural experiment, this field experiment differed in terms of how price changes were determined. Prices were customized based on the individual's previous BART demand and location. Prices were also updated at least monthly based on changes in congestion levels and other factors. In contrast, the natural experiment offered the same incentives to all participants for the duration of the pilot.¹⁹

Overall, the general size of the incentives, the way they are paid to consumers, and the type and selection of consumers are very similar across both experiments. The price changes started in mid-December 2018 and ended in March 2019. We could not examine any persistence in travel behavior in this experiment as those in the control group received offers in April 2019. Customers registered for the price subsidies via BART's website and the BART Official Mobile Application.

In terms of defining our field experiment, we class it as a hybrid between a natural and a framed field experiment, as defined by [Harrison and List \(2004\)](#). While we have selection of people into the experimental sample who wanted some incentives, people did not know that they were part of an experiment and did not know that they would get price subsidies for changing their commute time.²⁰ Therefore, we believe that there is little selection on peak and off-peak price elasticities due to the opaque advertising in the field experiment about what incentives would be given.


4.2 Recruitment

Participants were recruited from two sources. First, twelve people employed by BART distributed flyers advertising that people could sign up for the program at the Embarcadero, Montgomery and Civic Center stations in downtown San Francisco from 8:00 – 9:30 AM on December 13th, 2018. These three stations were chosen because they are the most congested stations in downtown San Francisco. Second, a subset of those who participated in Perks natural experiment opted in to receiving notifications from BART about future incentive programs. We analyzed the travel histories of these individuals, and identified customers who made at least 50 percent of their BART trips from downtown Embarcadero, Montgomery, or Civic Center stations. An email invitation to join the Perks Phase II program was sent to a random subset of qualifying individuals. Emails were sent in batches between November 20 and December 12, 2018. The overall recruitment led to 1900 customers, of which 63% were recruited through flyers and 36% through email. See appendix Table [A19](#) for the summary statistics for our sample.

¹⁹Price subsidies were paid out through gift cards as opposed to PayPal (as in the natural experiment).

²⁰The program was advertised as giving monetary payments for completing surveys and traveling on the BART on evening and weekends. See the recruitment flyer in Figure [A15](#).

Figure 4: View of the incentives offer on the app



4.3 Price changes

Experimental subjects in the treatment group were offered a price reduction of around 25% per ride in an adjacent time window to their current commute time. The price subsidy had two components: a random term and a deterministic term. The latter was based on how far the time interval was from their usual departure time. The price changes were targeted to each customer's frequently-visited station at a specific 20-minute time window during the morning and/or evening commute period. The time window where the subsidy was offered was up to 40 minutes before or after their typical entry time, and was optimized to reduce overall crowding on the BART network. Users could receive up to four commute-related point offers (shift early AM, shift late AM, shift early PM, shift late PM).

The price subsidies were determined by a machine learning algorithm that drew upon a participant's travel history and a predictive model for crowding provided by the transit company [Metropia](#). Figure 6 provides an example of an offer to shift commuting time. The price subsidies were offered in terms of points that were to be redeemed by the consumer.

To identify the correct time windows to reduce overall BART congestion, the Metropia algorithm identified a user's typical departure time and then calculated the predicted crowding reduction benefit of shifting this user to one of the adjacent time periods. If no benefit would occur from shifting the user, then no offer would be shown. Similarly, no offer was shown if achieving the crowding reduction would require the user to make shifts beyond 40 min-

utes from their average departure time.²¹ The consumer was emailed the price change at the beginning of the month and the price change was constant for one month, so we are estimating one-month elasticities in this field experiment (in comparison with six month elasticities from the natural experiment). During the course of the field experiment, about \$23,000 in peak shifting price subsidies were given to BART consumers.

4.4 Results

Our analysis of the field experiment consists of two parts. First, we provide the overall treatment effects of the incentives on shifting ride demand later or earlier. Second, we provide a more precise estimate of the welfare effects of shifting travel into or out of specific 20 minute time periods.

To estimate the average treatment effect of the price subsidies for shifting demand earlier or later than the usual travel time, we employ a linear regression of whether a rider travels at the subsidized time period based on the subsidy they receive as the dependent variable.

For the shift early subsidies, we estimate the following equation:

$$y_{it}^{SE} = \alpha + \beta T_{it}^{SE} + \epsilon_{it} \quad (5)$$

Here, y_{it}^{SE} is an indicator variable for rider i who receives the shift early subsidy and travels at the subsidized time on date t . T_{it}^{SE} is an indicator variable for rider i in the treated group who receives the subsidy (while it is zero for control group riders who receive a hypothetical subsidy). The parameter β captures the overall effect of the shift early subsidies on travel at the subsidized time.

Similarly, we define another indicator variable u_{it}^{SE} to capture whether a rider received a shift early offer but still travels at their usual travel time on date t . We estimate the following equation:

$$u_{it}^{SE} = \gamma + \delta T_{it}^{SE} + \mu_{it} \quad (6)$$

Here, u_{it}^{SE} represents the indicator variable for rider i who received a shift early offer but still travels at their usual time on date t . T_{it}^{SE} is the treatment indicator, indicating whether rider i received a shift early offer on date t . The parameter δ measures the average effect of the subsidy in reducing travel at the usual time for riders.

²¹This interval was based on feedback from user focus groups, which found that individuals did not wish to make large shifts in their commute

We follow a similar approach to estimate the effects of the shift late subsidies, using indicator variables y_{it}^{SL} and u_{it}^{SL} to capture riders traveling at the subsidized time and their usual time, respectively.

In Figure 5, we present the likelihood of riders traveling at their usual time or the subsidized time, comparing the treatment and control groups. The results show a clear effect of the subsidies in shifting riders from their usual travel time to the subsidized time. Specifically, for the shift early subsidies, the estimated equation (5) suggests that the subsidies increased the probability of riders traveling at the subsidized time (compared to the control group) by 0.034 (s.e. 0.009), or 3.4%. Equation (6) suggests that the subsidies reduced the probability of riders traveling at their usual time by 0.035 (s.e. 0.016), or 3.5%.

For the subsidies that shift riders to a later time, the analysis indicates an increase in the probability of riders traveling at the subsidized time by 0.022 (s.e. 0.008), or 2.2%, and a reduction in the probability of riders traveling at their usual time by 0.046 (s.e. 0.017), or 4.6%. These estimates demonstrate that the subsidies successfully shifted transit demand from congested time periods to non-congested time periods, resulting in meaningful and significant changes in travel behavior.


Our second part of analysis of the field experiment leverages the detailed part of the experiment, where we have varying off-peak and peak time windows for each consumer, so a flexible specification is required to account for this. We will also turn these estimates into welfare calculations. We estimate the following equation:

$$y_{it}^k = \beta_{k1}S_k + \beta_{k2}S_{-k} + \delta_{kjt} + \delta_{-kjt} + \mu_{ia}A_t + \mu_{id} + \mu_{iq} + \gamma_t + \epsilon_{it} \quad (7)$$

where y_{it}^k is the number of trips that start in time period k by individual i on date t . Our sample includes both the treatment and control groups, and both time periods (before and during the experiment). S_k is the monetary value of the points subsidy offered to individual i for time period k , regardless of the "usual travel time" of individual i (e.g., if k is 8am, S_k is any subsidy for 8am, whether it is for someone who usually travels at 7.20am, 7.40am, 8.20am or 8.40am). $S_k = 0$ for the treated group in the pre-treatment period, for the control group, and for those who are not offered a subsidy for time period k . β_{k1} is therefore the own-price demand effect for time period k .

S_{-k} is the value of the subsidy received by individual i for any time period outside of time period k , targeted for someone who usually travels at time k . It is the value of the subsidy offered to individual i who usually travels at time period k that may shift them away from time period k (e.g., if k is 8am, S_{-k} are subsidies for 8.20am/8.40am/7.40am/7.20am that shift people away from 8am). For individuals who do not normally travel in time k , $S_{-k} = 0$ for any

Figure 5: Field Experiment: Change in demand for subsidized time and usual travel time



Note: The figure presents the share of days riders travelling at their usual time (Panel (a)), and the subsidized time (Panel (b)), by the treatment and control riders, and by the period they receive shift early or shift late subsidies. Each bar represents the share of days the outcome (travelling at the usual time or the subsidized time) takes the value of one for the respective group. The red bar(s) represent the average outcome for treated riders and the white bar(s) represent that outcome for the control riders for the respective groups. The p-values above the pairs of bars indicates the statistical significance for the differences between the treated and control riders in the outcome variable, estimated by the regression equations (the treatment coefficients are 0.0336 (0.00928), 0.0217 (0.0082), -0.0353 (0.0163) and -0.0458 (0.0171) respectively for the shift early usual time, shift late usual time, shift early subsidized time and shift late subsidized time). Standard errors are clustered at the rider level and are in parentheses. The sample includes dates when riders receive a shift early (or late) offer or a hypothetical shift early (or late) offer for the control group riders. The experimental period begins December 2018 and goes through March 2019.

subsidy they were offered.²² β_{k2} is therefore the generic cross-price effect on the demand for time period k when a price subsidy was offered in time periods other than k .

δ_{kjt} is a time-varying fixed effect for the individuals who normally travel in time j and have received a subsidy for time period k (i.e., those who were actually chosen by the Metropia machine learning algorithm to travel on a less congested train within a 40 minute window of their usual travel time). Each type of rider, jk , includes riders from the treated group and the control group. For the treated group it is defined by the price subsidy offered (those who normally travel in time j who received a subsidy for time period k), and for the control group, it is defined using the hypothetical offer that would have been given to the control group if they were in the treatment group (those who normally travel in time j who received a hypothetical subsidy for time period k).²³ The fixed effects therefore control for any time-varying factors that affect the travel pattern of similar types of riders in the treatment and control group. For the control group, $S_k = 0$. Hence, the identification of S_k comes from a comparison of riders of similar "types" as defined by Metropia offers in both the treatment and control groups.²⁴

δ_{-kjt} is the time-varying fixed effect for the individuals who normally travel in time k and were offered a subsidy for time period j . Each type of rider, jk , includes both riders from the treated and control groups. For the control group, $S_{-k} = 0$. Hence, our identification for S_{-k} comes from a comparison of riders of similar type in the treatment and control groups that were identified as having similar offers.


μ_{id} is a fixed effect for individual-day of the week. It controls for individual travel pattern on each day of the week. μ_{iq} is an individual*year quarter fixed effect, which allows for an individual's travel pattern to change in each quarter. μ_{ia} is an individual fixed effect that takes a value of one for the weeks beginning November 12 and 19, 2018, in which the air quality in San Francisco was a historically poor level due to wildfires. It controls for any idiosyncratic response of each individual during those two weeks. γ_t is the date fixed effect. Our null hypotheses are that the price subsidy in the targeted time period has a zero impact on demand and that the cross-price effect on demand is zero.

²²Specifically, let S_m^k be the subsidy that a targeted traveler who usually travels in time period k and subsidize travel in time period m . As in our setup the subsidy shift traveller earlier and later by 20 or 40 minutes interval, if we index the time interval according to the order of time of the day, $S_m^k = 0$ if $|k - m| > 2$. S_{-k} is defined as $S_{-k} = S_{k-2}^k + S_{k-1}^k + S_{k+1}^k + S_{k+2}^k$. As each traveller has a specific usual travel time k , and receive subsidy for only one time period other than k at each period t , it is equivalent to $S_{-k} = \max\{S_{k-2}, S_{k-1}, S_{k+1}, S_{k+2}\}$.

²³The algorithm created hypothetical offers for all consumers, but we only gave the offers to the treated group and the control group was unaware of such offers.

²⁴Equivalently, these are the fixed effect coefficients for a time invariant indicator D_{kj} interacted with time dummies, where the indicator D_{kj} is 1 if the rider i in the treatment group has ever received subsidy of type S_{kj} in the experiment period, or if the rider i from the control group has ever received a hypothetical subsidy of type S_{kj} .

Figure 6: Field Experiment: Demand Parameter Estimates



Note: The figure plots the estimates of equation 7. The sample include participants in the field experiment. Panel (a) plots the estimate for the effect of a small change in the subsidy at time t on the demand at time t . Panel (b) plots the estimate for the effect of a small change in subsidy for time t , where the usual travel time is time j , on the demand at time j . Standard errors are robust to autocorrelation of 14 days. The average journey price in the control group is \$3.99.

Figure 6 plots the demand estimates, β_{k1} , for each of the time periods in equation 7. Panel (a) plots the demand change in the time period that subsidies were offered. We find, in general, that the estimates for β s are positive, meaningful, and statistically significant, suggesting riders are responsive to price at relatively narrow time intervals (as compared to the wider window in experiment 1). The time periods with the greatest demand shifts toward that period are 740am-8am, 8am-820am, 840am-9am, and 1020am-1040am. The average own-price elasticity of demand from peak to off-peak is -.939. However, there are some time periods where people are less elastic in their demand, such as before 740am and after 9am (excluding the 1020am-1040am time period).

Panel (b) plots β_{k2} of equation 7, which is the change in the usual travel time for the consumer. In general, we find negative and statistically significant coefficients, suggesting that riders were reducing travel in their usual time and shifted into the subsidized time period (panel (a)). The average cross-price elasticity of demand from peak to off-peak is .543.²⁵ It seems

²⁵The cross-price elasticity for each type of treatment subsidy (a subsidy for those who normally travel in period j' receiving a subsidy for time period k') is calculated as $\epsilon_{j'k'} = \frac{dx_{j'}}{dS_{-j}} * \frac{p}{x_{j'}}$ where $x_{j'}$ is the average travel (in number of trips) for those who normally travel in time period j' who received subsidy in time period k' , $\frac{dx_{j'}}{dS_{-j}}$ is the estimate of the coefficient for S_{-k} in equation (7) for the time period $k = j'$, and $p = 3.99$. The average elasticity is arithmetic average of the elasticities $\epsilon_{j'k'}$ all the treatment subsidies (include both shifting early and late).

that consumers are more elastic between 7am and 10am, but less elastic before 7am and after 10am.

From both panels in Figure 6, it is clear that elasticities vary across time periods. This variation in elasticities results from a combination of individual preferences and BART system effects. For example, the own-price and cross-price elasticity at 8-820am are -1.45 and 0.30 respectively.

For those individuals who did not switch at any time period, we sent them a survey asking them why they did not switch earlier or later.²⁶ We found that for those who would not switch earlier, the most chosen reasons were personal preference (55%) and employer would not allow it (29%), and that for those who would not switch later, the most chosen reasons were also personal preference (41%) and employer would not allow it (40%). This suggests possible unobserved reasons for not switching demand to an earlier or later time period.

The previous analysis allows us to compare the elasticities from the field experiment to those from the natural experiment. The own-price elasticity of demand and cross-price elasticity of demand from the natural experiment were -0.86 (se = .14) and 0.44 (se = .14) respectively. Given the respective estimates from the field experiment are -0.94 and 0.54, we cannot reject that the estimates from both experiments are equal for both the own-price elasticity and the cross price elasticity.

5 Welfare analysis

In this section we use the theory presented in section 2 and estimates of demand from the natural and field experiments to derive the change in welfare. We also present an alternative welfare calculation using the marginal value of public funds (MVPF).

We present an overview of the key data used in the calculation in section 5.1. We then present welfare estimates for the natural experiment in section 5.2. We extend this analysis to allow for differences in demand across groups (in this case, customers that received or did not receive a subsidy) in section 5.2.2. Finally we estimate the MVPFs for the natural experiment in section 5.3 and the field experiment in section 5.4.

The welfare analysis yields several novel findings. First, providing the off-peak subsidy for BART users improves welfare both when we use our welfare model and when we compute the MVPF (Finkelstein and Hendren, 2020). We find that the average net benefits per dollar of direct subsidy is \$0.36 in the base case. We also show that the MVPF in this case is 1.6. This means that for each additional dollar of net cost to the government, the marginal benefit

²⁶Survey implementation and results can be found at BART Perks Phase II Evaluation Report (2019).

to all parties is \$1.60.²⁷ The increase in welfare from the subsidy results from an increase in overall ridership and a shift in ridership from the peak to the off-peak period, which reduces congestion.

Second, allowing for riders to take account of how other riders respond to congestion (*i.e.*, the case of endogenous congestion) also leads to an increase in welfare, but less of an increase than for the case of exogenous congestion. The reason for a higher level of welfare in the case of exogenous congestion is that fewer people ride in the peak period (which yields a smaller welfare loss), and more people ride in the off-peak period (which yields a larger welfare gain).

Third, based on the natural experiment, targeting particularly congested routes can increase average per capita welfare. We estimate that per capita welfare gains increase by 15% when targeting subsidies to a highly congested route—the Transbay route.

Fourth, the optimal subsidy may be substantially higher than the subsidy offered in the natural experiment. We estimate that the optimal subsidy could be about eight times higher than the subsidy that was offered, and that welfare could be about four times higher. As we note below, the particular numerical results are likely to be sensitive to our assumption of linearity of demand.

Finally, using the natural field experiment, we show that the targeting of subsidies to time periods that are relatively narrow can increase welfare, but does not always do so. Our experiments suggest that such targeting may require a substantial amount of information on demand to increase the likelihood that welfare will actually improve with a change in the level of the subsidy.

5.1 Natural Experiment: Data

We analyze the welfare effect of experiment 1 using equation (1) in section 2. We consider both the case of exogenous and endogenous congestion. The key parameters used in equation (1) in the “base case” are presented in Table 4.

The expressions for $\frac{dx_1}{ds}$ and $\frac{dx_2}{ds}$ (the impact of a change in the subsidy on peak and off-peak demand) are obtained from Section 3 for the case of exogenous congestion.²⁸ These quantities are measured in terms of the change in the average number of trips per weekday with a change in the subsidy. Price is obtained by averaging the fare for all BART trips in our sample, measured in dollars per trip. This includes all trips in the morning from 5am to noon, excluding those that begin or end in the airport because they have a different fare schedule.

²⁷We compute all MVPFs based on the assumption that the initial subsidy is zero.

²⁸For endogenous congestion estimates, see the discussion in appendix A.6 and A.16.

Table 4: Key parameters used in our base case welfare analysis

Parameter	Exogenous congestion	Endogenous congestion	Percent Change
$\frac{dx_1}{ds}$	-0.0194	-0.0166	-14.68
$\frac{dx_2}{ds}$	0.0664	0.0625	-5.91
p	\$3.99	\$3.99	-
s'	\$0.25	\$0.25	-
MPC	\$1.89	\$1.89	-
MEC_2	0	0	-
MSC_1	\$3.22	\$3.22	-
MSC_2	\$1.89	\$1.89	-

Notes: Demand decreases in the peak period and increases in the off-peak period with the introduction of the subsidy. The magnitude of the decrease is smaller with endogenous congestion in the peak period, and the magnitude of the increase is smaller with endogenous congestion in the off-peak period. See discussion in the text for sources of the parameter estimates. Demand is measured as the average number of trips taken per weekday for the average rider.

The subsidy, s' , is taken from the actual experiment and equals \$0.25 per trip on average. MPC refers to marginal private cost of a BART trip. MSC_1 and MSC_2 refer to the marginal social cost per trip in the peak and off-peak period, respectively.

The marginal private cost in each period is assumed to be the same in both periods. We estimate the MPC from BART's financial report from 2006-2016. We follow [Winston and Maheshri \(2007\)](#), and regress yearly total operating cost (adjusted for inflation) on yearly annual ridership to obtain the marginal private cost. In addition, we assume that there is no externality in the off-peak period, so the marginal external cost in the period, MEC_2 , equals zero. This implies that the marginal social cost in the off-peak period equals the marginal private cost.

To obtain the marginal social cost in the peak period, we need an estimate of marginal external cost (MEC_1) to add to the MPC . The estimate of MEC_1 is based on crowding, delay and the value of time. Our analysis assumes that external costs are proportional to the value of time for both crowding and delay (see [Haywood and Koning \(2015\)](#)). We also assume that MEC_1 is constant for given values of the value of time, congestion and delay. We assume the value of time is half the hourly wage ([Small, 2012](#)), also do a sensitivity analysis with 75% of the wage rate ([Goldszmidt et al., 2020](#)). Congestion and delay are derived from BART data (see appendix [A.4](#) for details of this calculation). Using this approach, we find that MEC_1 is \$1.33 per peak hour trip.

Table [4](#) shows that the assumption about the response to congestion matters for the demand

response to the subsidy in both the peak and off-peak periods. Even though peak travel declines with the introduction of the subsidy in both cases, the decline is about 15% less with endogenous congestion than with exogenous congestion. More people travel during the peak period when they take into account how others will respond to congestion. Similarly, fewer people travel during the off-peak period with endogenous congestion (see the discussion in section 2 and Appendix A.6 for details of the empirical estimation). In both the exogenous case and the endogenous case, the subsidy increases overall ridership. The change in overall ridership, measured by $\frac{dx_1}{ds} + \frac{dx_2}{ds}$, is 0.047 trips per day in the exogenous case and 0.046 trips per day in the endogenous case. Thus, the endogenous case has a two percent lower increase in overall ridership than the exogenous case.

Table 4 also suggests that MSC_1 and MSC_2 are both less than the price in the base case. This means that there is under-consumption of rides relative to the first best in both periods (where price would be set equal to the marginal social cost). This relationship helps explain some of our welfare findings. We also consider a case for a crowded corridor where MSC_2 is less than the price.²⁹

5.2 Natural Experiment: Welfare estimates

5.2.1 Estimates for the Experimental Sample

We consider the welfare impacts of two subsidies in detail: one is a general subsidy for all riders in the experiment 1, and the second is a targeted subsidy for riders on what we call the Transbay route. We define the Transbay route to include trains that lead to the Transbay Corridor³⁰, and also trains that travel on the Transbay route westbound. This route is in the top 10 percent of congested routes for BART in the morning.³¹ We conclude by briefly considering the welfare effects of an optimal subsidy.

Table 5 presents the main welfare results for the base case subsidy and a targeted subsidy for the Transbay route. First, we consider the base case. We measure welfare in dollars per person per weekday. Because the price is above the marginal social cost in both periods in the base case, an increase (decrease) in ridership in a particular period will increase (decrease) welfare in that period (see equation (1)). Ridership decreases in the peak period as a result of

²⁹The average fare exceeding the marginal social cost may be unusual in rapid transit systems. For example, Parry and Small (2009) find that the marginal private cost, and thus the marginal social cost, exceeds price for rail in Washington D.C., Los Angeles and London. Winston and Maheshri (2007) find similar results.

³⁰See “Transbay Corridor Core Capacity Program Train Control Modernization Project”, BART.

³¹This includes routes that depart or arrive at stations: West Oakland, Embarcadero, Bay Fair, San Leandro, Coliseum, Fruitvale, Lake Merritt, Orinda, Rockridge, MacArthur, 19th St/Oakland, 12th St/Oakland City Centre, and Ashby, towards the Transbay tube direction.

the subsidy, and there is a decline in welfare (see row 1). Ridership increases in the off-peak period, and there is an increase in welfare (see row 2).

Table 5: Welfare impacts for the base case and a targeted subsidy

Parameter	Base case			Targeted subsidy		
	Exogenous conges- tion	Endogenous conges- tion	Percent Change	Exogenous conges- tion	Endogenous conges- tion	Percent Change
(1) Welfare change associated with peak travel change	-0.0037	-0.0032	-14.7	0.0006	0.0005	-14.7
(2) Welfare change associated with off-peak travel change	0.0328	0.0309	-5.9	0.0328	0.0309	-5.9
(3) Total Welfare change	0.0291	0.0277	-4.8	0.0334	0.0314	-6.1

Notes: Welfare is measured in dollars per person per weekday. Overall welfare increases with the subsidy. The decline in ridership in the peak period reduces welfare, and the increase in ridership in the off-peak period increases welfare. See text for details. The total welfare change is calculated from equation (1).

The overall welfare effect for the base case is the combination of the welfare loss associated with reduced peak travel and the welfare gain associated with more off-peak travel. For both the exogenous and the endogenous congestion cases, this net gain is positive. A subsidy of \$0.25 per trip in the off-peak period results in a net welfare gain of about \$0.03 per rider per day for the exogenous congestion case (row 3 of Table 5).

Welfare in the endogenous case declines by 4.8% compared with the exogenous case because of differences in demand response (see row 3 of the table). This percentage change masks more significant differences in actual travel patterns in the peak and off-peak periods (see the percentage change in row 1 and row 2 for the base case in Table 5).

The welfare analysis for the Transbay route is similar to the base case qualitatively. The main point of this analysis is to illustrate that there could be potential gains from targeting. The average net benefits are about 15% higher when targeting Transbay riders under the assumption of exogenous congestion.³²

Average net benefits per dollar of direct subsidy, which sometimes is used as a measure of welfare by decision makers, is also higher.³³ This measure ignores the fiscal externality, which is discussed below in estimating the MVPF. We estimate average net benefits by taking the total welfare change and dividing by the average subsidy expenditure per person per weekday.³⁴

³²Take the average welfare gain of .0334 for Transbay and divide by .0291 for the base case. Note that this percentage is higher in the endogenous case.

³³The measure is similar to that used in the tax literature that estimates the welfare change per dollar of tax revenue. See, e.g., Saez, Slemrod and Gieritz (2012).

³⁴See appendix A.10 for details.

For the base case, we find that an essential difference between the analysis of the Transbay route and our analysis of the base case is the level of congestion, which is higher. This, in turn, increases the MSC associated with the Transbay route during the peak period. The average peak period congestion on the Transbay route is 108 passengers per car, while it is 57 passengers per car for all other routes.³⁵ We also find that delay is more severe on the Transbay route than other routes. Using equation (1), we obtain a welfare estimate of providing a subsidy that exclusively focuses on the Transbay route during the morning (see appendix A.8 for details of the calculation).

For the Transbay route, the marginal social cost per peak period trip is \$4.12. This is higher than the marginal social cost when we use the average level of congestion across all routes (\$3.22). The welfare effect of a \$0.25 subsidy for an off-peak trip is \$0.033 per person per day in the case of exogenous congestion, and \$0.031 in the case of endogenous congestion. These welfare estimates are greater than the welfare effect effects estimated in the base case. The difference is due to the higher marginal social cost in the peak period for the Transbay route.

With additional assumptions about the functional form of demand, we can compute the optimal subsidy using equation (A5).³⁶ We provide details of this computation in appendix A7. The analysis relies on the assumption that demand is linear over a wide range, which may not be the case. With this caveat in mind, we find that the optimal subsidy is about eight times higher than the actual subsidy (\$1.87 in the exogenous case and \$1.90 in the endogenous case versus an actual subsidy of \$0.25 per trip). Total welfare is about the same for the exogenous and endogenous cases, but is about four times higher than welfare associated with the actual subsidy (\$.12 vs. \$.03). The reason that total welfare is similar in the endogenous and exogenous cases is because the optimal subsidy is slightly higher in the endogenous case, but the overall demand increase ($\frac{dx_1}{ds} + \frac{dx_2}{ds}$) is slightly lower. These changes tend to cancel each other out (see equation 1) (see appendix A7 for details).

We also considered what would happen with changes in the shape of the *MEC* function. In particular, we allow for the *MEC* function for peak period travel to be an increasing linear and non-linear function of density, e , instead of a constant function of e . Our main conclusion is that in this particular example, the results do not change much in moving from a constant *MEC* function to a linear *MEC* function or in moving from a constant *MEC* function to a quadratic *MEC* function for the peak period. For example, the total welfare change for the exogenous case is about .029 per person per day with a constant *MEC* function, .029 for a

³⁵The units of the congestion measure used in our calculation is passengers per square meter in an average car. See appendix A.7 for details.

³⁶We consider the optimal subsidy for the off-peak market. This is second-best because prices do not equal the marginal social cost in both markets. We could in principle compute optimal subsidies in both markets. However, to compute the welfare change, we would also need an estimate of how demand changes with price in the peak period, which we do not have.

linear *MEC* function and .028 for a quadratic *MEC* function. The results are also similar for endogenous congestion (.028, .028 and .027). Overall welfare decreases between exogenous and endogenous congestion are also similar, on the order of 5%. We discuss our methodology and present the empirical results in appendix [A.12](#).

The preceding analysis of welfare changes in Table 5 does not explicitly consider how changes in other modes of travel could affect our calculation. In Appendix [A.10](#), we present a calculation of how intermodal substitution could affect the welfare from the subsidy. In general, we argue that welfare would improve because of decreases in emissions and reduced congestion that are not priced in the market, supporting previous research by [Anderson \(2014\)](#). As an illustrative example, we assume that the overall miles traveled remains constant, but that the increase in the use of BART as a result of the subsidy is accounted for by a reduction in car travel. We compute the benefits associated with these improvements for the case in which the subsidy includes all riders, and all riders are assumed to respond similarly to those in the experiment. The estimated benefits include local pollution reduction benefits of \$0.07 million per year and global pollution benefits associated with CO_2 reductions of \$0.25 million per year, for a total of about \$0.32 million per year. These pollution reduction benefits may be compared with the benefits of the subsidy from the experiment, which are about \$2 million per year. We derive this estimate by multiplying the average benefit per person per weekday by the number of weekdays per year by the number of BART riders. See Appendix [A.10](#) for details of the calculation.³⁷

5.2.2 Estimates for the BART Network

We extend our theoretical framework to allow us to estimate the welfare benefits when riders may differ in their responsiveness to a subsidy. We do this to develop a more realistic welfare estimates when data are available.

Formally, we model two types of riders with different demand characteristics. We allow for differences in demand characteristics because people participating in the experiment may, for example, be more sensitive to price changes on average than those who do not. In appendix [A.17](#), we show that this can lead to a straightforward modification of equation 1, where the average per capita welfare change for all riders depends on the demand characteristics of both groups and their respective sizes. For example, when demand characteristics are the same for the two groups, average per capita welfare benefits from the subsidy are the same.

To empirically implement this framework, we first explored how the sample in the natural

³⁷We do not analyze how the introduction of the BART network affects pollution. Previous research suggests this could be important ([Gendron-Carrier et al., 2022](#)).

experiment differed from the general BART population. As discussed below, the sample does not appear to be that different. Nonetheless, because there may be unobserved differences between populations, we present a sensitivity analysis to assess how welfare could change if the sample under study was not representative of the general population.

A key question is how participants in the natural experiment may differ from the general population that uses BART. We relied on a survey administered by BART to address this question. Our main finding is that, based on the survey, the observable characteristics of our sample and the BART customer sample are no different. For instance, the percent of individuals who have household income less than \$50,000 in our natural experiment sample was 13%, versus 14% for all those using BART who commute downtown. The percent of individuals who are BIPOC in the natural experiment sample was 57%, versus 64% for all those using BART who commute downtown. Thus, in terms of external validity on observables, the sample is not that different from all those using BART.

However, there may be unobserved demand differences between those who signed up to the natural experiment and the rest of the BART network. To address this issue, we compared our base case results with a case in which the rest of the BART network is assumed to be one-half as responsive to the price subsidy as participant in the experiment. In the base case, recall that the average net benefit per dollar of direct subsidy was \$0.36. If the demand by those not in the experiment is one-half that of people in the experiment, this measure drops by about 46% to \$0.2. The reason for the large drop is that other people not in the experiment represent about 93.5 percent of the entire network, and thus have are the major factor driving the per capita welfare calculation. This analysis highlights the need to understand the demand characteristics of different groups of riders to develop a more precise estimate of the welfare impacts of a large-scale subsidy.

5.3 Natural Experiment: MVPF estimates

The preceding section used a measure of welfare based on our welfare model. It assumes that the method of financing the subsidy is a lump sum tax. Net benefits are computed in terms of dollars per person per week day. In contrast, the MVPF calculation for the subsidy is silent on the source of the revenues that fund the intervention.³⁸

The MVPF benefits measure is given in terms of the change in net benefits per net additional dollar of government cost. It is, thus, not a welfare measure in the sense of providing an estimate of net benefits (*i.e.*, benefit minus costs). However, if we assume that the net costs to

³⁸This source can be considered in a separate MVPF – for example, the MVPF of raising a dollar of funds through an income tax on a particular group.

the government are funded by a lump sum tax, it is possible to show that the MVPF measure and the measure for welfare we obtain in our basic model are related.³⁹ The intuition is that a lump sum tax converts the net cost to the government into a cost to society.

Equation 8 describes our MVPF calculation.

$$MVPF = \frac{\text{Inframarginal benefit from the subsidy} + \text{Change in congestion benefits}}{\text{Change in direct subsidy cost} + \text{Change in the transfer to BART}} \quad (8)$$

In our application, the benefit to BART users is the numerator and the net cost to the government of the subsidy is the denominator. The numerator consists of two terms: inframarginal benefits of the subsidy for those off-peak users who did not switch from the peak, and the benefits to all peak travelers from reduced congestion.⁴⁰ The denominator also consists of two terms: the direct subsidy costs that are paid to BART consumers, and the change in the transfer to BART resulting from the change in ridership.

The two terms in the numerator are positive. The subsidy provides benefits to those who do not change their usage patterns and there are benefits (in terms of reduced congestion) for those who continue to travel during the peak period. The two terms in the denominator have different signs. The direct subsidy cost to the government increases as a result of introducing the subsidy; however, in our particular example, discussed in more detail below, BART's revenues profits actually increase as a result of the subsidy, which means that the transfer from the government to BART, which is needed to cover BART's losses actually decreases. Thus, the second expression in the numerator is negative. This situation arises because ridership increases as a result of the intervention, and price exceeds marginal cost for each ride. BART's profits thus increase from the increased ridership, and the transfer needed by the government to keep BART's profits at a particular level decreases.

Table 6 shows the MVPF calculation for the cases of exogenous and endogenous congestion. In both cases the MVPF is close to 1.6. This means that for each additional dollar of net cost to the government, the marginal benefit to all parties is \$1.60. The 95% confidence interval for the MVPF, obtained by bootstrapping both the peak and off-peak demand response estimations, ranges from about 1.4 to 1.7 for both cases.⁴¹

There are two general points to be made about the MVPF approach based on our example.

³⁹They are related in the sense that the welfare measure $\frac{dW(s)}{ds}$ is the same using both approaches under this assumption. See appendix A.15 for details. A key feature common to both approaches is their reliance on the envelope theorem.

⁴⁰We assume the net benefits to those who shift from peak to off-peak travel as a result of the subsidy are negligible or zero because of the envelope theorem, and do not include this term in net benefits. If we included some part of these benefits, it would have the effect of increasing the MVPF.

⁴¹If we assume riders not in the experiment were half as sensitive to the subsidy as those in the experiment then the MVPF drops from 1.6 to 1.26.

Table 6: MVPF's for the subsidy

Numerator	Exogenous Congestion	Endogenous Congestion
1. Inframarginal benefits	0.307	0.307
2. Congestion benefits	0.0258	0.0220
Numerator total	0.3328	0.329
Denominator		
1. Change in direct subsidy cost	0.307	0.307
2. "Change in transfer to BART"	-0.0987	-0.0965
Denominator Total	0.2083	0.211
MVPF = Numerator/Denominator and error bounds	1.598 [1.479, 1.716]	1.563 [1.444, 1.681]

Notes: The MVPF is computed for a subsidy of \$1 per trip in the off-peak period using equation (8). The VOT is set equal to 50% of the wage rate. The 95% confidence interval is calculated by bootstrapping the estimation for the demand responses with 100 replications. Details of the calculation are provided in appendix A.7.

First, the subsidy in this case involves more than simply the direct subsidy to passengers. Because we are dealing with a regulated entity, one must also consider the impact of the subsidy on the revenues needed to support the regulated entity – in this case, BART. This is a feature that is endemic to many regulatory activities, where prices depart from marginal cost (e.g., [Hahn and Metcalfe \(2021\)](#)).

Second, more work needs to be done in this area to rank various subsidies and pricing interventions using the MVPF framework in the regulatory domain, and connecting them to tax policies that might be used to assess their effectiveness. That will allow us to make more informed assessments of the actual welfare impacts of policies. A review of [Hendren \(2020\)](#) reveals the MVPFs we estimate are comparable to, and in some cases higher than, many educational and labor policy interventions.⁴²

5.4 Field Experiment: MVPF Estimates

This section analyzes the welfare effect of the targeted subsidies used in experiment 2. In the interest of brevity, we focus on results using the MVPF here, and present results on the net

⁴²A higher MVPF does not necessarily imply the policy is "better." One would need to make further assumptions about equity weights to compare policies.

welfare benefits in the appendix. The results are qualitatively similar.⁴³

There are two main results. First, targeting results in many MVPFs that exceed one (meaning the increase in willingness to pay exceeds the net cost to the government), but in some that do not. This is because targeting parameters are set without reference to changes in demand, and these changes are critical for welfare. Second, we explain the subtle relationship between a change in the value of time and a change in welfare. The key point is that the value of time interacts with the demand changes in both periods to affect the change in welfare. As a result, the impact of a change in the VOT on economic welfare cannot be signed *apriori*.

In experiment 2, thirty groups receive a subsidy in 12 time periods that are 20 minutes long. The subsidies are to encourage travel that is 20 minutes or 40 minutes away from the rider's usual travel period of 20 minutes. A group is defined by their subsidized time period and their usual period of travel. Because of data limitations, we assume that the targeted subsidy for a specific group only affects travel in the targeted period and the usual travel time period.

We identify each targeted group, j , by the usual travel interval and the subsidized interval, which can be 20 or 40 minutes before or after the usual travel time. For example, one group is riders who usually travel from 8:20 am to 8:40 am and receive a subsidy to travel between 8.40 am and 9 am.

We focus on evaluating the welfare impact of the subsidies for experiment 2 using the MVPF. This is because each of the targeting subsidies involves a small fraction of riders among all BART users. We consider the case in which the riders take congestion as exogenous in their decision making. In this case, the MVPF provides a simple measure that allows us to rank the welfare impact of each of the subsidies.⁴⁴

In calculating the MVPF for subsidies in experiment 2, we allow for the possibility that the marginal external cost is positive in the targeted interval. We calculate the MVPF for each subsidy using the following formula:

$$MVPF_{jt} = \frac{x_{jt} - MEC_{j1} * \frac{dx_{j1}}{ds_{jt}} - MEC_{jt} * \frac{dx_{jt}}{ds_{jt}}}{x_{jt} - (p - MPC) * (\frac{dx_{j1}}{ds_{jt}} + \frac{dx_{jt}}{ds_{jt}})} \quad (9)$$

$MVPF_{jt}$ is the MVPF of the subsidy given to group j for the targeted time interval t ; p is the price for a trip; MEC_{j1} is the marginal external cost for the route taken by group j on the usual time before the subsidy, and MEC_{jt} is the marginal external cost for the same route in the time targeted by the subsidy. x_{jt} is the average number of trips taken by group j in the

⁴³See appendix B Figure A16.

⁴⁴Our qualitative findings are the same when we use our welfare model in section A.1 to estimate welfare changes.

targeted travel time. The numerator consists of two parts: the inframarginal benefits for those who receive the subsidy in time t (the first term) plus the change in external congestion costs (the last two terms). The denominator consists of two parts as well: the direct costs to the government of the subsidy for customers (the first term) plus the change in profits to BART (the second term).


We measure the congestion level using the passenger density on the route travelled by group j , at the usual travel time and also at the subsidized time. We also consider the travel delay in estimating the marginal external cost in each time period. This allows us to calculate MEC_{j1} and MEC_{jt} .

To calculate the marginal external cost, we need a measure of the value of time (see equation (A9)). Using survey data from the sample, the median wage for a single person household is \$50 dollar per hour. We focus on these households because the survey did not ask for the number of working individuals in a household, but did ask for income and the number of people in the household. This wage estimate is higher than the median wage used to evaluate experiment 1, which was based on all San Francisco residents. We consider a value of time equal to 50% of the wage, as we did before, and also one equal to 75% of the wage. This allows us to calculate the marginal external and social cost using equation (A9), as we did for experiment 1.

We use the estimate of marginal private cost per trip discussed in section 2. We assume the price of an average trip is about \$4 for all groups. The behavioral parameters, $\frac{dx_{j1}}{ds_{jt}}$ and $\frac{dx_{jt}}{ds_{jt}}$, are based on the analysis in section 3. Table A11 in section A.13 presents the parameter values we used. Figure 7 plots the MVPFs for each of the subsidies in experiment 2, by group. The figure shows that the point estimates of the MVPFs of these targeted subsidies can be above or less than 1 for both the case where VOT is 50% of the wage rate (panel a) (Small, 2012) and 75% of the wage rate (panel b) (Goldszmidt et al., 2020). For example, the MVPF for the subsidy that shifts riders who usually travel in 8:00-8:20am to travel during 8.40-9.00am exceeds 1; the MVPF for the subsidy that shifts riders who usually travel from 9:20-9:40am to travel during 8:40-9:00am is less than 1. The average MVPF for a value of time at 50% of the wage rate is 1, but for a value of time at 75% of the wage rate the average MVPF is 1.17.

The impact of a change in the value of time cannot be signed apriori. A change in the value of time enters into the analysis through its impact on the MEC in both periods. If the change in the value of time entered into the calculation in only one period, the comparative statics would be straightforward. However, when we allow for a change in the MEC in both periods, the comparative statics depends on the MEC *and* demand parameters in both periods, and this effect cannot be signed a priori. We find that in 75.7 percent of cases, an increase in value of time from 50 to 75 percent increases welfare, and in 24.3 percent of cases an increase in

Figure 7: MVPF of subsidy in experiment 2



Note: This figure shows the MVPF of a targeted subsidy. See text for details. The x-axis represents the beginning time of the targeted time period. The shape of each point shows the usual travel time of the group of riders, which is 20 or 40 minutes earlier or later than the targeted time. The value of time is assumed to be 50% of the wage in panel (a) and 75% of the wage in panel (b).

value of time decreases the welfare effect. ⁴⁵

A comparison of the MVPFs from the two experiments suggests that the average MVPF from the field experiment is lower than the MVPF obtained in the natural experiment. This may at first seem unusual because one might expect net benefits to be higher in the case of targeting (experiment 2) than with a general off-peak subsidy (experiment 1). However, there are several possible reasons for this finding. First, the MVPF provides a measure of marginal and not total net benefits. A social planner could in principle use the information from a finely targeted experiment like the field experiment to achieve higher total net benefits than an experiment that simply provided the same subsidy to all users (e.g, experiment 1). This increase in net benefits could in principle be achieved with differentiated subsidies across different time intervals. Second, the two experiments are measuring different things. The natural experiment measures the marginal benefit of moving people from peak to off-peak, and the field experiment measures the marginal benefit of moving from one relatively narrow time interval to an adjacent time interval, which is not necessarily off-peak. Finally, the estimate of congestion used in the field experiment varies across all of the relatively small time intervals, whereas we assume for simplicity in our natural experiment that there is one average level of congestion that applies to the peak period, and there is no congestion in the off-peak period. The bottom line is that there is no reason to assume a particular relationship between the MVPFs estimated in the two experiments, but the more finely tuned approaches could lead to higher net benefits from targeting.

6 Conclusion

Using a large natural experiment and a natural field experiment, we develop novel estimates of peak and off-peak price elasticities for urban mass transit demand in San Francisco. Both experiments allow us to obtain causal estimates of demand.

Using the demand estimates, we assess the welfare impacts of the price subsidies using a sufficient statistics approach. Our analysis suggests that off-peak subsidies can increase welfare, but the positive effects are reduced when consumers take the decisions of others into account compared to when they do not. We also find a large variation in the welfare impacts of shifting travel to different periods, which is explained by differences in demand and congestion

⁴⁵The precise formula for measuring the impact of VOT on MVPF is obtained by differentiating equation (9) with respect to value of time. It is given by $-(\frac{dMEC_{j1}}{dVOT} \frac{dx_{j1}}{ds_{jt}} + \frac{dMEC_{jt}}{dVOT} \frac{dx_{jt}}{ds_{jt}})MVPF_{jt}$, where VOT represent value of time, and $\frac{dMEC_{j1}}{dVOT}$ and $\frac{dMEC_{jt}}{dVOT}$ are the change in each of the MEC with respect to VOT, that depends on the crowding level during the time period, and how the delay relates to the number of passengers. See appendix A.4 for the derivation and further discussion of the empirical results.

characteristics. Finally, we show that the targeting of subsidies can increase welfare, but need not do so if the regulator does not have accurate information on demand.

The experiments offer different insights into the impact of providing subsidies at particular times for economic welfare. We find that the subsidy in the natural experiment generally increases welfare. In the base case, we show that the marginal value of public funds is 1.6. We also calculate the average net benefits per dollar of subsidy in the natural experiment, and find that the average net benefits are \$0.36 per dollar of subsidy expenditure. In contrast, for the field experiment, MVPFs showed considerably more variation depending on the particular time period being targeted.

There are three areas for future research that we believe would be useful. First, with respect to the issue of managing transport networks, we think it would be instructive to explore the extent to which various forms of congestion pricing result in attractive MVPFs. In particular, we would advocate doing more experiments that allow for causal identification of key elasticities related to peak and off-peak pricing. This would allow scholars to draw more robust conclusions about the generalizability of studies in this area ([List, 2020, 2022](#)). In addition, we think it would be useful to have experiments that focused on both the intensive margin (*e.g.*, making greater use of a transport mode that a person already uses) and the extensive margin (*e.g.*, changing transport modes or changing the city structure ([Desmet and Rossi-Hansberg, 2013](#); [Severen, 2019](#); [Tsivanidis, 2019](#); [Heblich, Redding and Sturm, 2020](#); [Barwick et al., 2021](#))).

Second, as pointed out by [Kleven \(2021\)](#), sufficient statistics is only one of many approaches for estimating the impacts of economics policies. A strength of this approach is its use of the envelope theorem in simplifying the analysis. But this may also be a weakness when dealing with non-marginal policy changes. While non-marginal changes can, in principle, be estimated using our welfare model and the MVPF approach ([Hendren, 2020](#)), it is an empirical question as to whether such estimates are reasonably accurate. It would be instructive to explore how different models yield different results for a particular class of problems, such as the one studied here. It would also be useful to explore the relationship between sufficient statistics approaches designed to measure welfare, such as those developed by [Chetty \(2009\)](#), and the MVPF approach advanced by [Hendren \(2020\)](#).

Third, we noted in our analysis that the targeting of subsidies does not always meaningfully increase welfare when the decision maker does not have perfect information. We think that empirical and theoretical work could help identify when targeting actually helps.

References

- Aiken, Emily, Suzanne Bellue, Dean Karlan, Chris Udry, and Joshua E Blumenstock.** 2022. "Machine learning and phone data can improve targeting of humanitarian aid." *Nature*, 603(7903): 864–870.
- Alatas, Vivi, Ririn Purnamasari, Matthew Wai-Poi, Abhijit Banerjee, Benjamin A Olken, and Rema Hanna.** 2016. "Self-targeting: Evidence from a field experiment in Indonesia." *Journal of Political Economy*, 124(2): 371–427.
- Allcott, Hunt, and Dmitry Taubinsky.** 2015. "Evaluating behaviorally motivated policy: Experimental evidence from the lightbulb market." *American Economic Review*, 105(8): 2501–38.
- Allcott, Hunt, Benjamin B Lockwood, and Dmitry Taubinsky.** 2019. "Regressive sin taxes, with an application to the optimal soda tax." *Quarterly Journal of Economics*, 134(3): 1557–1626.
- Allcott, Hunt, Sendhil Mullainathan, and Dmitry Taubinsky.** 2014. "Energy policy with externalities and internalities." *Journal of Public Economics*, 112: 72–88.
- Anderson, Michael L.** 2014. "Subways, strikes, and slowdowns: The impacts of public transit on traffic congestion." *American Economic Review*, 104(9): 2763–96.
- Athey, Susan.** 2018. "The impact of machine learning on economics." In *The economics of artificial intelligence: An agenda*. 507–547. University of Chicago Press.
- Barwick, Panle Jia, Shanjun Li, Andrew R Waxman, Jing Wu, and Tianli Xia.** 2021. "Efficiency and equity impacts of urban transportation policies with equilibrium sorting." National Bureau of Economic Research.
- Boiteux, Marcel.** 1960. "Peak-load pricing." *The Journal of Business*, 33(2): 157–179.
- Borenstein, Severin.** 2005. "The long-run efficiency of real-time electricity pricing." *Energy Journal*, 26(3).
- Bundorf, M Kate, Jonathan Levin, and Neale Mahoney.** 2012. "Pricing and welfare in health plan choice." *American Economic Review*, 102(7): 3214–48.
- Burkhardt, Jesse, Kenneth Gillingham, and Praveen K Kopalle.** 2019. "Experimental evidence on the effect of information and pricing on residential electricity consumption." National Bureau of Economic Research.

- Burlig, Fiona, Christopher Knittel, David Rapson, Mar Reguant, and Catherine Wolfram.** 2020. "Machine learning from schools about energy efficiency." *Journal of the Association of Environmental and Resource Economists*, 7(6): 1181–1217.
- Byrne, David P, Leslie A Martin, and Jia Sheen Nah.** 2019. "Price Discrimination, Search, and Negotiation in an Oligopoly: A Field Experiment in Retail Electricity."
- Callaway, Brantly, and Pedro HC Sant'Anna.** 2021. "Difference-in-differences with multiple time periods." *Journal of Econometrics*, 225(2): 200–230.
- Chassang, Sylvain, Padró I Miquel, Erik Snowberg, et al.** 2012. "Selective trials: A principal-agent approach to randomized controlled experiments." *American Economic Review*, 102(4): 1279–1309.
- Chetty, Raj.** 2009. "Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods." *Annual Review of Economics*, 1(1): 451–488.
- Christensen, Peter, Paul Francisco, Erica Myers, Hansen Shao, and Mateus Souza.** 2021. "Energy Efficiency Can Deliver for Climate Policy: Evidence from Machine Learning-Based Targeting."
- Cohen, Jessica, Pascaline Dupas, and Simone Schaner.** 2015. "Price subsidies, diagnostic tests, and targeting of malaria treatment: evidence from a randomized controlled trial." *American Economic Review*, 105(2): 609–45.
- Davis, Lucas W.** 2021. "Estimating the price elasticity of demand for subways: Evidence from Mexico." *Regional Science and Urban Economics*, 87: 103651.
- De Chaisemartin, Clément, and Xavier d'Haultfoeuille.** 2020. "Two-way fixed effects estimators with heterogeneous treatment effects." *American Economic Review*, 110(9): 2964–96.
- Desmet, Klaus, and Esteban Rossi-Hansberg.** 2013. "Urban accounting and welfare." *American Economic Review*, 103(6): 2296–2327.
- Domencich, Thomas A, and Daniel McFadden.** 1975. "Urban travel demand-a behavioral analysis."
- Dupas, Pascaline.** 2014. "Short-run subsidies and long-run adoption of new health products: Evidence from a field experiment." *Econometrica*, 82(1): 197–228.
- Farhi, Emmanuel, and Xavier Gabaix.** 2020. "Optimal taxation with behavioral agents." *American Economic Review*, 110(1): 298–336.

- Finkelstein, Amy, and Nathaniel Hendren.** 2020. "Welfare analysis meets causal inference." *Journal of Economic Perspectives*, 34(4): 146–67.
- Fowlie, Meredith, Catherine Wolfram, Patrick Baylis, C Anna Spurlock, Annika Todd-Blick, and Peter Cappers.** 2021. "Default effects and follow-on behaviour: Evidence from an electricity pricing program." *The Review of Economic Studies*, 88(6): 2886–2934.
- Gendron-Carrier, Nicolas, Marco Gonzalez-Navarro, Stefano Polloni, and Matthew A Turner.** 2022. "Subways and urban air pollution." *American Economic Journal: Applied Economics*, 14(1): 164–96.
- Gerarden, Todd D, and Muxi Yang.** 2021. "Using Targeting to Optimize Program Design: Evidence from an Energy Conservation Experiment." Tech. rep., Working Paper.
- Glaister, Stephen.** 1974. "Generalised consumer surplus and public transport pricing." *Economic Journal*, 84(336): 849–867.
- Glaister, Stephen, and Davis Lewis.** 1978. "An integrated fares policy for transport in London." *Journal of Public Economics*, 9(3): 341–355.
- Goldszmidt, Ariel, John A List, Robert D Metcalfe, Ian Muir, V Kerry Smith, and Jenny Wang.** 2020. "The Value of Time in the United States: Estimates from Nationwide Natural Field Experiments." National Bureau of Economic Research.
- Hahn, Robert W, and Robert D Metcalfe.** 2021. "Efficiency and Equity Impacts of Energy Subsidies." *American Economic Review*, 111(5): 1658–88.
- Harberger, Arnold C.** 1964. "The measurement of waste." *American Economic Review*, 54(3): 58–76.
- Harrison, Glenn W, and John A List.** 2004. "Field experiments." *Journal of Economic literature*, 42(4): 1009–1055.
- Haywood, Luke, and Martin Koning.** 2015. "The distribution of crowding costs in public transport: New evidence from Paris." *Transportation Research Part A: Policy and Practice*, 77: 182–201.
- Heblich, Stephan, Stephen J Redding, and Daniel M Sturm.** 2020. "The making of the modern metropolis: evidence from London." *Quarterly Journal of Economics*, 135(4): 2059–2133.
- Hendren, Nathaniel.** 2013. "Private information and insurance rejections." *Econometrica*, 81(5): 1713–1762.

- Hendren, Nathaniel.** 2020. "Measuring economic efficiency using inverse-optimum weights." *Journal of Public Economics*, 187: 104198.
- Holmgren, Johan.** 2007. "Meta-analysis of public transport demand." *Transportation Research Part A: Policy and Practice*, 41(10): 1021–1035.
- Ito, Koichiro, Takanori Ida, and Makoto Tanaka.** 2018. "Moral suasion and economic incentives: Field experimental evidence from energy demand." *American Economic Journal: Economic Policy*, 10(1): 240–67.
- Ito, Koichiro, Takanori Ida, and Makoto Tanaka.** 2021. "Selection on welfare gains: Experimental evidence from electricity plan choice." National Bureau of Economic Research.
- Jacobsen, Mark R, Christopher R Knittel, James M Sallee, and Arthur A Van Benthem.** 2020. "The use of regression statistics to analyze imperfect pricing policies." *Journal of Political Economy*, 128(5): 1826–1876.
- Jacobsen, Mark R, James M Sallee, Joseph S Shapiro, and Arthur A van Benthem.** 2022. "Regulating Untaxable Externalities: Are Vehicle Air Pollution Standards Effective and Efficient?"
- Jansson, Kjell.** 1993. "Optimal public transport price and service frequency." *Journal of Transport Economics and Policy*, 33–50.
- Jessoe, Katrina, and David Rapson.** 2014. "Knowledge is (less) power: Experimental evidence from residential energy use." *American Economic Review*, 104(4): 1417–38.
- Joskow, Paul L, and Catherine D Wolfram.** 2012. "Dynamic pricing of electricity." *American Economic Review*, 102(3): 381–85.
- Kleven, Henrik J.** 2021. "Sufficient Statistics Revisited." *Annual Review of Economics*, 13: 515–538.
- Knittel, Christopher R, and Ryan Sandler.** 2013. "The welfare impact of indirect pigouvian taxation: Evidence from transportation." National Bureau of Economic Research.
- Knittel, Christopher R, and Samuel Stolper.** 2019. "Using machine learning to target treatment: The case of household energy use." National Bureau of Economic Research.
- Kreindler, Gabriel.** 2022. "Peak-hour road congestion pricing: Experimental evidence and equilibrium implications." *Unpublished paper*.

- Larcom, Shaun, Ferdinand Rauch, and Tim Willems.** 2017. "The benefits of forced experimentation: striking evidence from the London underground network." *Quarterly Journal of Economics*, 132(4): 2019–2055.
- Levin, Laurence, Matthew S Lewis, and Frank A Wolak.** 2017. "High frequency evidence on the demand for gasoline." *American Economic Journal: Economic Policy*, 9(3): 314–47.
- List, John A.** 2020. "Non est disputandum de generalizability? A glimpse into the external validity trial." National Bureau of Economic Research.
- List, John A.** 2022. *The voltage effect: How to make good ideas great and great ideas scale*. Currency.
- Litman, Todd.** 2004. "Transit price elasticities and cross-elasticities." *Journal of Public Transportation*, 7(2): 3.
- Mahoney, Neale, and E Glen Weyl.** 2017. "Imperfect competition in selection markets." *Review of Economics and Statistics*, 99(4): 637–651.
- McFadden, Daniel.** 1974. "The measurement of urban travel demand." *Journal of Public Economics*, 3(4): 303–328.
- McFadden, Daniel.** 2001. "Economic choices." *American Economic Review*, 91(3): 351–378.
- McFadden, Daniel, Antti Talvitie, Stephen Cosslett, Ibrahim Hasan, Michael Johnson, Fred Reid, and Kenneth Train.** 1977. "Demand model estimation and validation." *Urban Travel Demand Forecasting Project, Phase, 1*.
- McFadden, Daniel, William B Tye, and Kenneth Train.** 1977. *An application of diagnostic tests for the independence from irrelevant alternatives property of the multinomial logit model*. Institute of Transportation Studies, University of California Berkeley, CA.
- Mohring, Herbert.** 1972. "Optimization and scale economies in urban bus transportation." *The American Economic Review*, 62(4): 591–604.
- Mullainathan, Sendhil, and Jann Spiess.** 2017. "Machine learning: an applied econometric approach." *Journal of Economic Perspectives*, 31(2): 87–106.
- Oum, Tae Hoon, and Michael W Tretheway.** 1988. "Ramsey pricing in the presence of externality costs." *Journal of Transport Economics and Policy*, 307–317.
- Parry, Ian WH, and Kenneth A Small.** 2009. "Should urban transit subsidies be reduced?" *American Economic Review*, 99(3): 700–724.

- Piketty, Thomas, Emmanuel Saez, and Stefanie Stantcheva.** 2014. "Optimal taxation of top labor incomes: A tale of three elasticities." *American Economic Journal: Economic Policy*, 6(1): 230–71.
- Polyakova, Maria, and Stephen P Ryan.** 2019. "Subsidy targeting with market power." National Bureau of Economic Research.
- Reguant, Mar.** 2019. "The efficiency and sectoral distributional impacts of large-scale renewable energy policies." *Journal of the Association of Environmental and Resource Economists*, 6(S1): S129–S168.
- Rodrik, Dani.** 1987. "Policy targeting with endogenous distortions: Theory of optimum subsidy revisited." *Quarterly Journal of Economics*, 102(4): 903–911.
- Saez, Emmanuel, Joel Slemrod, and Seth H Giertz.** 2012. "The elasticity of taxable income with respect to marginal tax rates: A critical review." *Journal of Economic Literature*, 50(1): 3–50.
- Severen, Christopher.** 2019. "Commuting, labor, and housing market effects of mass transportation: Welfare and identification." *Review of Economics and Statistics*, 1–99.
- Small, Kenneth A.** 2012. "Valuation of travel time." *Economics of transportation*, 1(1-2): 2–14.
- Spulber, Daniel F.** 1989. *Regulation and markets*. MIT press.
- Tarduno, Matthew.** 2022. "For Whom the Bridge Tolls: Congestion, Air Pollution, and Second-Best Road Pricing." *Unpublished manuscript*.
- Tizghadam, Ali, Hamzeh Khazaei, Mohammad HY Moghaddam, and Yasser Hassan.** 2019. "Machine learning in transportation."
- Train, Kenneth.** 1977. "Optimal transit prices under increasing returns to scale and a loss constraint." *Journal of Transport Economics and Policy*, 185–194.
- Tsivanidis, Nick.** 2019. "Evaluating the impact of urban transit infrastructure: Evidence from bogota's transmilenio." *Unpublished manuscript*.
- Vickrey, William S.** 1963. "Pricing in urban and suburban transport." *American Economic Review*, 53(2): 452–465.
- Weyl, E Glen, and Michal Fabinger.** 2013. "Pass-through as an economic tool: Principles of incidence under imperfect competition." *Journal of Political Economy*, 121(3): 528–583.

- Williams, Kevin R.** 2021. "The welfare effects of dynamic pricing: Evidence from airline markets." National Bureau of Economic Research.
- Winston, Clifford.** 1985. "Conceptual developments in the economics of transportation: an interpretive survey." *Journal of Economic Literature*, 23(1): 57–94.
- Winston, Clifford.** 2013. "On the performance of the US transportation system: Caution ahead." *Journal of Economic Literature*, 51(3): 773–824.
- Winston, Clifford, and Vikram Maheshri.** 2007. "On the social desirability of urban rail transit systems." *Journal of Urban Economics*, 62(2): 362–382.
- Wolak, Frank A.** 2007. "Residential customer response to real-time pricing: The Anaheim critical peak pricing experiment."
- Yang, Nan, and Yong Long Lim.** 2018. "Temporary incentives change daily routines: Evidence from a field experiment on Singapore's subways." *Management Science*, 64(7): 3365–3379.