
Moving towards a more equal world, one ride at a time: Studying Public Transportation Initiatives using interpretable causal inference

Gaurav Rajesh Parikh*
Duke Kunshan University
gaurav.rajesh.parikh@duke.edu

Jenny Huang*
Duke University
yijian.huang@duke.edu

Albert Sun*
Duke University
albert.sun@duke.edu

Lesia Semenova
Duke University
lesia.semenova@duke.edu

Cynthia Rudin
Duke University
cynthia.rudin@duke.edu

Abstract

The goal of low-income fare subsidy programs is to increase equitable access to public transit, and in doing so, increase access to jobs, housing, education and other essential resources. King County Metro, one of the largest transit providers focused on equitable public transit, has been innovative in launching new programs for low-income riders. However, due to the observational nature of data on ridership behavior in King County, evaluating the effectiveness of such innovative policies is difficult. In this work, we used seven datasets from a variety of sources, and used a recent interpretable machine-learning-based causal inference matching method called FLAME to evaluate one of King County Metro’s largest programs implemented in 2020: the Subsidized Annual Pass (SAP). Using FLAME, we construct high-quality matched groups and identify features that are important for predicting ridership and re-enrollment. Our analysis provides feedback for policy-makers on what type of interventions might be effective in improving access to transit. In particular, we found that SAP is effective in increasing long-term ridership and re-enrollment. Notably, there are pronounced positive treatment effects in populations that have higher access to public transit and jobs. Treatment effects are also more pronounced in the Asian population and in individuals ages 65+. Insights from this work can help broadly inform public transportation policy decisions and generalize broadly to other cities and other forms of transportation.

1 Introduction

Access to high-quality affordable transportation is important for making city resources more equitable and promoting fair access to job opportunities, education, and other essential services [1]. However, forming policies to make transport equitable is difficult. Unclear definitions of measures of equity, difficulty in determining who will benefit from any given transit policy intervention, and challenges in obtaining robust data where causal effects can be disentangled make it hard to decide on effective interventions [2]. Although city planners have attempted to address equitable transportation through various policy interventions like subsidy programs and fare free transit, causal effects of policy interventions are hard to estimate, and it is hard to see if interventions have improved equity.

In 2020, the King County Metro began providing an income-based fare program called the Subsidized Annual Pass (SAP). The Subsidized Annual Pass program aims to ensure that people with low income

*equal contribution




Figure 1: Dataset descriptions and questions that we focus on in this work. Note that Person ID uniquely identifies each eligible person, Stop ID is defined as the unique ID for every bus stop or point from where transit may be available, Trip ID is defined as one trip made by a bus along a specific route, with outbound and inbound trips having unique Trip ID's.

can access the bus. Depending on how effective SAP currently is, the King County Metro is currently deciding whether to expand the qualification category to include more low-income people in SAP as early as 2022 [3]. This work is an attempt to answer one of these important King County Metro challenges – that of predicting re-enrollment and ridership of bus service in response to the Subsidized Annual Pass.

To address this important problem, ideally, we would conduct a randomized controlled trial (RCT), the gold standard to assess policy interventions. However, it is hard to design a proper RCT due to the difficulty of observing outcomes and establishing a control group in the first place [4]. Interestingly, King County *did* attempt to conduct an RCT, where they provided individuals with cards loaded with different amounts of cash incentives or SAP passes. However, the data provided by the King County Metro lacked necessary information – for instance, we did not know which individuals were in the control group; we have records for individuals assigned to the control (given a card with money loaded onto it) who were and were not recruited for the RCT, and we could not distinguish between them. Hence, it is difficult to determine answers to questions such as: did the people who received the SAPs need it most in the first place? What if they are already the most frequent riders? We are thus left back where we started: grappling with complicated data from any source we could get our hands on.

In our search to find useful data, we examined a diverse range of datasets detailing ridership, local demographic information, job access via transit and unequal access to jobs. Despite the complexity of the problem, our entire data pipeline is interpretable. For the treatment effect estimates, we leverage an interpretable machine learning-based causal inference technique called FLAME [5], which uses machine learning to determine the important variables to use for matching, and matches units exactly on as many of these important variables as possible. This process creates matched groups that can be manually checked for trustworthiness, allowing us to form reliable conclusions. Finally, to ensure the validity of our approach, we compared our results to other black box approaches, as a sensitivity analysis.

From our study, we find very interesting results: (1) the free pass programs are effective in increasing long-term ridership and enrollment, which means that when given a free pass incentive in the short term, more riders will be inclined to ride on public transportation long term. (2) The free passes influence some subgroups more than others, in particular, they have a larger impact on those who have more access to transportation and jobs, older age groups (65+) and Asians. That is, perhaps handing out free passes in the short term might lead to more long-term public transportation use and improved equity and access across the country.

1.1 Data

In our research, we used several datasets on specific aspects of our target population including survey data, census data as well as automatically collected transit data in the form of GTFS Feed data on

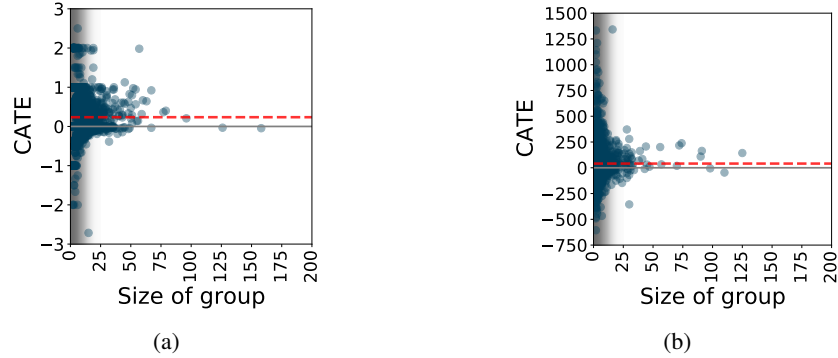


Figure 2: This figure shows Conditional Average Treatment Effects (CATEs, each matched group is represented by a dot) and Average Treatment Effect (ATE, in lines) of SAP for (a) re-enrollment and (b) ridership. SAP appears to increase both re-enrollment and ridership. More specifically, the estimated ATE of SAP on (a) re-enrollment is +0.27 times, which means users enroll 0.27 more times on average when given a SAP than given other incentives. The estimated ATE of SAP on (b) ridership is +40.70 rides during our time period. Matched groups in which there are more units (towards the right) are more trustworthy (gray indicates smaller, less trustworthy matches).

public transit usage. Our datasets are useful for insights on **demographic information, ridership and re-enrollment, and unequal commute.**

American Community Census Data. This open source data provides demographic information on household income, educational attainment, employment, and housing characteristics, etc [6]. **Spatial Data.** Spatial catalog data from King County Metro. This data is open source, and can be obtained from [7]. **APC Data.** Data on public transit schedule and activity. This data can be requested from the King County Metro/CISIL Data Challenge organizers. **ORCA-LIFT Card Registry.** ORCA-LIFT is a program created by the King County Metro to give different subsidy incentives to King County Metro users. This dataset gives information on incentives including monetary benefits and the Subsidized Annual Pass. There are 98,228 unique individuals in this dataset. **ORCA-LIFT Lift Boardings.** This dataset records the boardings made for ORCA-LIFT users [2]. **Urban Institute’s Unequal Commute Dataset.** The Unequal commute data gives data on job access via transit. We look at differential effects of access to jobs and public transportation in order to understand the difficulty for low-income riders to commute to their jobs [1].

In the appendix, we included a discussion of limitations of the data and improvements for future data collection.

1.2 Model Choice

Data on SAP enrollment is purely observational, meaning that individuals may be selected for treatment in a way that may depend on the response. For example, individuals who are already more inclined to ride may be more likely to be chosen for the SAP program.

To account for these confounders, we choose to take the approach of matching; we match individuals in SAP (treatment) and individuals in all other incentive groups (control) based on demographics such as age, race, and job access via transit. We use a method known as Fast Large-scale Almost Matching Exactly (FLAME) [5] that uses a learned (rather than pre-defined) distance metric to create matched groups on as many relevant covariates as possible.

The FLAME method uses a machine learning algorithm on a held-out training set to drop irrelevant confounders. This method thus ensures that units are matched on confounders relevant to predicting re-enrollment and ridership [5]. We use FLAME as it does not introduce a distance metric on the covariate space, a priori, like other matching methods such as current coarsened exact matching, mixed integer programming matching, and network flow methods, instead, it *learns* the distance metric. Further, it improves over regression, propensity and their variable selection methods by not forcing a model form on either the treatment or the outcome, instead it matches on covariates

Figure 3: Conditional average treatment effects (CATEs) for race and job access via transit on re-enrollment (a) Asian individuals are more likely to ride King County Metro as a result of enrolling in SAP than non-Asian Individuals. (b) Treatment effect of SAP on re-enrollment by job access via transit, where job access via transit is given in quartiles. Individuals with higher job access via transit are more likely to ride King County Metro as a result of enrolling in SAP than individuals with lower job access via transit.

Figure 4: Conditional average treatment effects (CATEs) for race and job access via transit on ridership. (a) Treatment effect of SAP on long term ridership by race. Asian individuals are more likely to ride King County Metro as a result of enrolling in SAP than individuals of other races. (b) Treatment effect of SAP on long term ridership by job access via transit. Individuals with higher job access via transit are more likely to ride King County Metro as a result of enrolling in SAP than individuals with lower job access via transit.

directly (non-parametrically). Unlike other non-parametric methods like black box machine learning approaches like Causal BART, FLAME is interpretable. [5]

We are able to see the covariates each individual was matched on in Figure 8. After the matched groups are constructed from FLAME, conditional average treatment effects (CATEs) can be calculated from each matched group by comparing the treatment and control outcomes of units in the matched group. The CATEs for the individual matched groups can be combined to determine average treatment effect estimates for any desired subgroup. Interpretability is essential for real-world problems because black box causal inference are difficult to troubleshoot. With FLAME, because you have interpretable matched groups, you can manually check the trustworthiness of those matched groups. Our analysis has the advantage of being interpretable for policy-makers to identify which demographics benefited most from the SAP program to understand how future resources should be allocated.

More formally, assuming common conditions for causal inference of (1) unconfoundedness, (2) SUTVA, and (3) conditional ignorability, our matched groups can provide valid estimates of the causal effects of the SAP program. Within each matched group, we used the difference between the average outcome of the treated units and control units to estimate the conditional average treatment effect (CATE) as $E[Y(1) - Y(0) | X^{REL}]$ given covariates X^{REL} associated with that group [5].

In order to calculate the number of times individuals re-enrolled, we aggregated the number of times a particular user appears in the data set after October 1, 2020. In order to measure long-term ridership,

(a) Treatment effect of SAP on Asian population. Matched groups with individuals of Asian descent (dotted in yellow) showed a more positive conditional estimated average treatment effect of SAP on ridership compared to other races (ATE = +112.12 more rides). Gray indicates smaller, less trustworthy matches than the larger ones on the right.

(b) Treatment Effect of SAP in Hispanic populations. Matched groups with Hispanic individuals (dotted in pink) show a negative conditional estimated average treatment effect of SAP on ridership (ATE = -56.60). Gray indicates smaller, less trustworthy matches than the larger ones on the right.

Figure 5: Matched groups for Asians and Hispanic populations

we aggregated the total number of times individuals boarded (focusing exclusively on King County Metro), starting six months after they registered in ORCA-LIFT.

2 Findings

2.1 Re-enrollment

Overall, there was a positive effect of SAP on ORCA-LIFT re-enrollment and long-term ridership, shown in Figure 2. Based on the estimated treatment effect (ATE), which is +0.27 after matching, the subsidized annual pass (SAP) incentive causes individuals to re-enroll 0.27 more times, on average, compared to individuals given other ORCA-LIFT incentives. For instance, if a non-SAP incentive user enrolls 2 times, a SAP user would enroll 2.27 times, on average. We show the sizes of the matched groups in this figure; matched groups formed from a larger number of units are more reliable.

Our results show that individuals who are Asian and individuals belonging to census tracts with highest job access via transit and overall job access showed a greater effect of SAP on re-enrollment. That is, there is a higher positive conditional average treatment effect (CATE) for individuals located in regions with better transit access and better overall job access. Additionally, we found that matched groups with individuals of Asian descent showed higher CATE estimates, compared to matched groups of White individuals. We also note that the size of matched groups with Asian individuals tended to be large, giving us reliability in the CATE estimates due to the larger size of the groups.

The most important characteristics in predicting enrollment include overall job access, job access via transit, and spatial mismatch (Fig. 8a). Spatial mismatch is the number of job opportunities within a 30-minute commute from a job seeker over the sum of the competition for those jobs. Figure 6(a) summarizes the subgroups for which the treatment effect estimates were the largest for SAP on re-enrollment.

2.2 Ridership

SAP also increases long-term ridership (total boardings counted after 6 months of enrollment). The SAP incentive causes individuals to board a total of 40.70 more trips per year on King County Metro, on average, compared to individuals given other incentives (average treatment effect of +40.70).

For conditional average treatment effects, individuals belonging to census tracts with highest job access via transit and overall job access showed a greater effect of SAP on ridership. Individuals of Asian descent show the largest treatment effect. In terms of age, older individuals see more positive treatment effects of SAP on ridership, with the effects of those in the oldest age group (65+) showing a treatment effect nearly double that of those in the 55-65 age group. These results are consistent with our results on re-enrollment.

Figure 6: These tables outline the characteristics of the groups who saw the greatest increase in (a) re-enrollment and (b) ridership from receiving the SAP incentive.

The most important characteristics in predicting ridership include race, stage of card (the number of cards an individual has used), and spatial mismatch (Fig. 8). People of Asian descent, who are already long term users of the metro and who live in areas with high spatial mismatch tend to be the individuals who benefited most from the free passes.

(a) Re-enrollment (b) Ridership

Figure 7: To ensure the robustness of our analysis to the matching method chosen, we compared our estimated treatment effects obtained using FLAME to a state-of-the-art black box model, Bayesian Additive Regression Trees [Causal BART], Average treatment effect of (a) re-enrollment and (b) ridership by model. Treatment effect estimates are very similar between FLAME and Causal BART [8], a state-of-the-art black box model.

3 Conclusion

We find that low-income incentive programs such as subsidized-annual passes are effective for incentivizing ridership in low income populations. Additionally, we find that individuals from census tracts with highest access to jobs benefited the most from the SAP incentive. Follow-up observational studies (particularly long term) would be useful to monitor the effectiveness of the free fare programs.

Since our analysis is reproducible and each matched group can easily be checked for trustworthiness, it acts as strong point of reference for policy planning. We suggest, that cost permitting, reduced and free fare programs may be implemented and expanded in neighborhoods and areas with high job access via transit and job hubs, such as inner-cities. In fact, King County Metro is considering expanding the SAP program eligibility criteria in 2023. [Our work communicates to King County Metro stakeholders that the SAP program has been effective and provides specific insights into directions for future work to make transit more equitable. [2].

Code is available at the following GitHub link: Codebase

References

- [1] Wesley Jenkins, Christina Stacy, Alena Stern, Kristin Blagg, Yipeng Su, Eleanor Noble, Macy Rainer, and Richard Ezike. The unequal commute: Examining inequities in four metro areas' transportation systems, Oct 2020.
- [2] Causal Inference for Social Impact Lab, Center for Advanced Study in the Behavioral Sciences. Causal inference for social impact lab's data challenge. Available at <https://casbs.stanford.edu/causal-inference-social-impact-lab-s-data-challenge>, 2022.
- [3] Dow Constantine. King county metro and sound transit offer new subsidized annual pass for riders in greatest need, 2021.
- [4] Louis Wright, Tim Denne, and Aaron Schiff. Ex-post evaluation of transport interventions using causal inference methods. *NZ Transport Agency research report 630*, pages 1–154, 2017.
- [5] Tianyu Wang, Marco Morucci, M. Usaid Awan, Yameng Liu, Sudeepa Roy, Cynthia Rudin, and Alexander Volfovsky. Flame: A fast large-scale almost matching exactly approach to causal inference. *Journal of Machine Learning Research*, 22(31):1–41, 2021.
- [6] US Census Bureau. American community survey 5-year data (2009-2020), Apr 2022.
- [7] KCGIS. King County ACS Select Indicators, King County Spatial Data Catalog. Available at <https://www5.kingcounty.gov/sdc/TOC.aspx?subject=census>, 2019.
- [8] Jennifer L Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, 2011.
- [9] Parth Bhavsar, Ilya Safro, Nidhal Bouaynaya, Robi Polikar, and Dimah Dera. Chapter 12 - machine learning in transportation data analytics. In Mashrur Chowdhury, Amy Apon, and Kakan Dey, editors, *Data Analytics for Intelligent Transportation Systems*, pages 283–307. Elsevier, 2017.
- [10] Shenhao Wang, Qingyi Wang, Nate Bailey, and Jinhua Zhao. Deep neural networks for choice analysis: A statistical learning theory perspective. *Transportation Research Part B: Methodological*, 148:60–81, 2021.
- [11] Yunhan Zheng, Shenhao Wang, and Jinhua Zhao. Equality of opportunity in travel behavior prediction with deep neural networks and discrete choice models, 2021.
- [12] US Census Bureau. In what ways are d.c.-area neighborhoods with rail transit different from other neighborhoods?, Oct 2021.
- [13] Pieter A. Gautier and Yves Zenou. Car ownership and the labor market of ethnic minorities. *Journal of Urban Economics*, 67(3):392–403, May 2010.

4 Appendix

4.1 Related works

4.1.1 Machine Learning in Transportation

Many transportation researchers are beginning to leverage traditional machine learning techniques such as support vector machines, decision trees, and deep neural networks to conduct data analysis of transportation patterns [9, 10]. Recent work such as [11] even use deep neural networks to analyze transit accessibility, an issue that is central to our paper.

Unfortunately, these techniques are black box machine learning models, making it difficult for users to understand why those models make the decisions they do. Further, it is important to note that

Figure 8: These tables show the covariates that matched groups were formed on against the number of matches made on those covariates for (a) re-enrollment and for (b) ridership. “Y” (yes) and “N” (no) indicate whether or not matches were formed using the row’s covariate at the column’s iteration number. The columns represent iterations of FLAME, starting from the left. Histograms above the tables show that a large number of units were matched on a large number of covariates, with the least important covariates dropped first (such as which agency location gave them the card) and important covariates like spatial mismatch dropped later.

prediction methods do not naturally extend to causal analysis [hence the reason Causal BART is different than standard BART, see 8].

On the other hand, FLAME is inherently interpretable, and is designed for observational causal inference. Instead of matching on propensity score, FLAME makes almost exact matches using as many important covariates as possible. This allows us to create interpretable, high-quality matches between units in our public transit research project.

4.1.2 Equitable Access in Public Transit

Social scientists, policy researchers, and sociologists have long analyzed how racial disparities exist for American workers to get to work [12, 13]. Our research builds on this literature to understand how a subsidy plan can have differential effects on aforementioned historically-marginalized groups.

4.2 Limitations in the Data and Future Steps to Improve Data Collection

There are a few limitations to the data that we were given from the King County Metro. First, the SAP program launched in late 2020, giving us only two years worth of data to analyse long term ridership and re-enrollment, which limited the scope of our analysis. We encourage future analyses to look into the motivation behind re-enrollment when longer-term data are available. Second, we did not have the necessary data to ensure that our treatment and control groups had the same income.

Prior to the launch of the SAP incentive program, King County policymakers were interested in how ORCA-LIFT incentives impacted enrollment in SAP. Although our main analysis did not focus on this question, through preliminary exploratory analysis, we observe that those who received ‘monthly passes’ as their incentive showed the highest proportion of SAP enrollments. However, we note that we do not have data on individual income, hence are not able make conclusive statements about whether the individuals analyzed were at or below 80% of the poverty line (which is the SAP enrollment eligibility criteria).