Post-COVID Transit Fares for Riders and Recovery

by

Daniel M. O'Neil, Jr.

B.A., University of Oregon (2013)

Submitted to the Department of Urban Studies and Planning
in partial fulfillment of the requirements for the degree of

Master of Science in Transportation

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2023

Authored by: Daniel M. O'Neil, Jr.
Department of Urban Studies and Planning
May 12, 2023

Certified by: John P. Attanucci
Research Associate, Center for Transportation and Logistics
Thesis Supervisor

Certified by: Frederick P. Salvucci
Senior Lecturer, Center for Transportation and Logistics
Thesis Supervisor

Accepted by: Jinhua Zhao
Associate Professor of City and Transportation Planning
Chair, Transportation Education Committee

Post-COVID Transit Fares for Riders and Recovery

by

Daniel M. O'Neil, Jr.

Submitted to the Department of Urban Studies and Planning
on May 12, 2023, in partial fulfillment of the
requirements for the degree of
Master of Science in Transportation

Abstract

In the face of persistent large-scale changes in travel behavior spurred by the COVID-19 pandemic, mass transit agencies face a landscape full of new challenges. Transit ridership, often used as a primary measure of agency success, remains diminished. Nevertheless, the purpose of and benefits provided by a well-designed and well-operated transit network remain unchanged. This thesis investigates one powerful tool at the disposal of transit providers: fare policy. Fare policy can be used to spur transit usage, to fund agency operations, and to respond to societal goals. Rider-centric fare policies that simultaneously increase transit travel volumes while showing only small negative fare revenue impacts can be identified. Implementation of such policies is key moving forward to maintain public investment and individual engagement.

This thesis presents four case studies that analyze fare equity, new fare products, and multi-agency regional fare integration. First, fare equity is considered through a case study of Washington, DC's Metrorail transit fare structure, residential and employment geography, and user demographics. The results highlight policy elements that consistently improve fare equity regardless of structure type, including peak pricing differentiation and removal of penalties for circuitous travel. The second case study designs and evaluates novel fare products using post-pandemic travel patterns on the CTA. The hypothetical products considered differ from traditional offerings by changing the usage restrictions and the validity periods. A flexible pass that confers a set number of CTA journeys at a discounted per-trip price is found to be the most promising, as it would provide the most utility to riders for whom pay-per-use travel is currently the most economical choice. The third case study considers single-day fare capping as an alternative to traditional 1-day passes for transit users in Chicago, identifying benefits to reduced-fare and bus-only riders while providing opportunities to boost agency ridership. Finally, the results of a recently introduced fully-integrated, multi-agency transit pass in the Chicago region are analyzed. Fare structure changes are used to estimate post-COVID commuter rail fare elasticity, and the elasticity for integrated passes. Additional findings include large increases in cross-agency travel, new customers accessing secondary transit agencies, and continued opportunities to integrate.

Thesis Supervisor: John P. Attanucci
Title: Research Associate, Center for Transportation and Logistics

Thesis Supervisor: Frederick P. Salvucci
Title: Senior Lecturer, Center for Transportation and Logistics

Acknowledgments

My time here at MIT has been an incredible experience, filled with joy and excitement. This work attempts to capture what I've done and what I've learned, and it would not be what it is without the guidance, support, and friendship of many people.

John and Fred, I've appreciated our regular meetings over the last two years. John, your certainty in the importance of this research has always been reassuring and your investment in the details has helped strengthen my work. I hope someday you finally manage to golf your age. Fred, you've helped push the boundaries of this research farther and farther. I think I'm looking at the big picture, and you make the picture even bigger. Thank you for helping make sure my investigations encompassed all that they could, and for always taking the time to check in on me as a person. To Jinhua, thank you for your consistent encouragement and infectious passion, and thank you for bringing me into the JTL-Transit Lab. The field of transportation is lucky it has you. To Anson, Awad, and the other faculty and advisors in the lab – thank you for your presence, your feedback, and your engagement.

To my classmates and labmates, I've enjoyed getting to spend time with you all. You are doing important, exciting work and you're great to be around. In particular, I have to thank the other MSTs in my year (Devin, Dingyi, Emma, Yuzhu) for helping me throughout, my various research partners (Ehab, Nineveh, Seamus) for helping develop and execute so much interrelated research, and every lab member for making our Friday meetings such a welcoming place. It's often the highlight of my week.

Thanks also go to WMATA and the CTA for supporting this work. To Tom, Laura, Maulik, Molly, Liam, Michael, Jonathan, Gabe, and the rest of the staff at each, it's been a pleasure working with you.

I owe a debt to everyone who helped me get to MIT in the first place. Unfortunately, that's too many people to list, but it includes my friends (you know who you are), my colleagues and mentors, my undergraduate professors. It also includes my in-laws the Shryocks, who have been nothing short of great throughout this entire journey, and my second family the Cases, who have been there all my life.

The people I do need to thank individually include my mother Yvonne, who is excited and proud and taught me conviction and persistence. My sister Kelly, who pushes me to be better in all ways and is simultaneously one of the most fun people around. Our pets, Winnie and Dottie, who always let me know that the sun is rising and it's time to get up for breakfast. My wife Devon, who agreed to come on this adventure with me and has been there through the joys and sorrows with unwavering love – you'll always be my superhero.

And to my father, who always had a kind word to give and went out of his way to make time for others. He saw me leave home for this experience, but not much more. We miss you.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivation

The COVID-19 pandemic stands out as a monumental turning point in our most recent history, leading to a great number of changes across society. Stay-at-home orders and work-from-home policies may both have been introduced and understood as temporary measures to curb the spread of pathogens and combat a virus spiraling out of control, but as the magnitude of the pandemic continued to be better understood, and the timeline over which it continued to unravel surpassed original estimates, society became more comfortable operating in its newly adapted ways. The spread of technology only served to solidify many changes that occurred over the course of the pandemic, especially with regard to the traditional office environment, work hours and commuting patterns. The spread of teleworking, and the strong preference many employees have for continuing remote work post-pandemic, has raised a plethora of questions with which society will have to tussle. Within this thesis, one particular subset of COVID impacts is considered – public transit systems in the United States of America.

There was a gargantuan decrease in all travel during the pandemic, which led to a corresponding drop in public transit ridership. Nevertheless, many riders were still reliant on public transportation to get to work and take care of other basic needs. Transit-dependent riders became a larger share of transit trips while commuting trips, one of the traditional backbones of public transit ridership, shrank drastically as those with the ability to work remotely took advantage of the option. Additionally, many individuals with alternative options switched to non-transit travel modes to avoid exposure to the COVID-19 virus. Throughout the earlier periods of the pandemic, many transit agencies hoped that they were weathering a brief storm as opposed to staring down a whole new paradigm of travel patterns and transportation choices. To help keep agencies afloat through the early periods of lower ridership and lower fare revenues, which was expected to be temporary, and to hopefully counteract the possibility of cost savings achieved through extensive service reductions, the federal government stepped in to provide two infusions of funds as part of national COVID relief bills.

It is important to note that while the COVID-19 pandemic was having enormous effects on transit ridership, our society was simultaneously grappling with a long history of racial injustice. Catalyzed by the murder of George Floyd by a Minneapolis police officer, protests and rallies occurred across the country to advocate for racial justice. This movement spurred reflection in many individuals and organizations, and increased awareness of widespread inequities. One outcome has been a reappraisal of existing structures through an equity lens. The importance of equity, and specifically as it applies to equitable public transportation in this thesis, is already an area subject to considerable research. However, the importance of equity was magnified at the same time rider trends were shifting.

As pandemic-era restrictions were lifted, and health recommendations were updated in light of effective vaccines and improved understanding of the COVID-19 virus, aspects of everyday life began moving back towards of pre-pandemic norms.

People gradually returned to public places, and travel volumes slowly ticked up. Transit ridership similarly rose, but nowhere near as dramatically as it had fallen off. Across the country, it has failed to recover to near pre-pandemic levels. Data from the American Public Transportation Association ('APTA') and the Transit app noted a new highpoint of post-pandemic transit ridership was reached in September of 2022, with transit trips for the month totaling 70% of pre-COVID levels [1]. Two and a half years after the first stay-at-home orders were issued, and more than a year after vaccines became widely available, it is clear that the ridership levels transit agencies were accustomed to have as of yet not come close to fully returning. Many agencies have realized that a paradigm shift has fully arrived, and policies that made sense before the pandemic must be reconsidered for a post-COVID world.

While the exact contours of post-COVID travel remain uncertain in the long-term, three years on from the first stay-at-home orders transit travel volumes have not rebounded back to pre-pandemic levels. It remains possible that transit riders return to earlier travel frequencies and trip volumes. Nevertheless, for practical purposes this thesis focuses on a continuation of the current scenario and assumes we have reached a new equilibrium. Current evidence supports the idea of a new equilibrium. This assumption does not preclude other possibilities; in actuality there exist a range of potential long-term transit ridership scenarios. While the case studies presented here assume that we are in a steady state regarding transit ridership, the methodological frameworks developed are flexible enough to be applied to other transit scenarios as needed. Further, this uncertainty in and of itself suggests active intervention by transit agencies to support their desired outcomes, which includes a larger transit rebound moving forward.

Given the continued development of a more complete picture of travel pattern changes and an understanding of the COVID-era norms that have become entrenched, transportation policies must be reconsidered. Public transit exists to move people efficiently from origin to destination. There are aspects of geometry to system design; what path should routes follow? How can the most people be moved across the same space? Transit provision can also be thought of as a maximization problem. Agencies have a finite number of resources, and they use those resources to move people around cities. Agencies accumulate additional resources by charging customers to travel on their system, and the decisions about how to price travel can have major implications in terms of trip volumes, travel frequency, and rider equity.

Given the ridership and operational implications of transit fares, the changes to travel behavior due to COVID, and the increased societal recognition of equity's importance, traditional notions of an effective fare policy for transit systems are important to reevaluate. Fare policy is one lever agencies should explore both as an opportunity to encourage transit use for trip making and to ensure that the benefits of transit use are well aligned with its costs, ensuring that fare policy and fare products are not only rider friendly, but also equitable and efficient.

1.2 Research Questions

This thesis explores the components of an effective fare policy as they apply in our post-COVID world. It is an attempt to tie together different strands of analysis to arrive at a holistic understanding of how to determine an equitable, effective fare policy; how such a policy provides benefits to riders; and how such a policy affects agency fare revenues.

To investigate these questions, specific aspects of fare policy, fare products, and fare equity are explored using transit data from US public transportation agencies serving two major urban areas: Washington, DC and Chicago, Illinois. The analyses conducted in these cities can be distilled into the following four questions:

1. **How are the costs and benefits of transit experienced across differ-**

ent rider segments? Transit users are not a monolith; the riding public is composed of thousands of individuals with different incomes and backgrounds seeking travel options that meet different requirements. Similarly, the cost for using transit is not uniform across cities or even within agencies. There may be an average cost per trip, but that masks the distributional impacts felt across various customer segments. To move towards an equitable fare policy, an understanding of how different rider groups bear different costs for transit travel is necessary.

2. **What aspects of existing fare structures improve overall equity, and what aspects detract from system equity?** This question is closely tied to the preceding question. Local agencies determine the rules that govern how the cost of a given trip is determined. Fare determination can be as complex or as simple as the agency prefers. However, understanding differing travel behaviors when selecting a fare structure can have the effect of either improving or harming rider equity.

3. **What fare product design could meet new and shifting travel patterns?** It is clear that travel patterns have changed over the COVID-19 pandemic, with fewer work commutes each week and a less singular division between weekday and weekend transit use. The relative utilities of fare products, and in particular fixed period transit passes designed and priced for pre-COVID commuting, have seen a corresponding change. This opens the door for consideration of new fare product designs that complement current travel patterns.

4. **What role can improved cross-agency fare integration play in increasing regionwide transit use, and does it benefit all cooperating agencies?** Many metro areas have multiple providers of public transportation, some offering different and/or complementary modes. In some cases, these agencies

have not integrated their fare offerings making it difficult for a single traveler to utilize multiple providers during a single trip. This presents an opportunity to exploit the benefits of closer collaboration. Riders and operators both stand to benefit from a more integrated environment through lower travel costs and increased boardings, respectively.

While each of these questions is important to understand, final decisions about fare policy are ultimately proposed by high-level decision makers at each transit agency and approved by the agency's governing board. However, by exploring each of the questions listed in the context of a single transit agency, planners and analysts can make informed recommendations to guide fare policy in ways that complements agency priorities.

1.3 Analytical Approach

The research questions considered by this thesis are explored through a series of analyses conducted for the transit in systems Washington, DC and Chicago, IL using data provided through the major agency in each.

1.3.1 Data Sources

The analyses contained in this thesis are built primarily from Automated Fare Collection ('AFC') databases compiled at each agency. The available AFC data at both the CTA and WMATA provide extensive information on trip start location, mode choice, fare product used, cost to ride, whether the transaction is the start of a new journey or a transfer, and unique, anonymized customer identifiers that can be used to link travel at the customer level. The exact parameters of the data used to address each research question will be explained in detail in the relevant chapter.

AFC data from Chicago also contains fare information for Metra (commuter rail) and Pace (suburban bus) systems, as the same vendor, Ventra, is used across all three agencies. Riders are able to buy fare products for all three systems from multiple sales channels, but the smartphone app maintained by Ventra has become more popular over time and, particularly with respect to commuter rail, is responsible for a large number of fare transactions. Transactions on Pace and Metra record fewer AFC variables for retention in the database, although time of use and cost information are recorded for all transactions.

WMATA has additional detailed information on rail riders; given the nature of its distance-based fares each rider must tap into and out of the rail system resulting in a recorded destination station alongside the origin. This AFC data is the basis of WMATA's Trace model, which provides inferred origin, destination, and transfer information for trips taken on the WMATA network.

Beyond AFC records, additional data sources were used to supplement and support aspects of the analyses. Two particularly important sources for this thesis include AFC card registration records, some of which contain important rider demographic information, and local travel surveys conducted around the analysis periods. These additional sources have been integrated with AFC data in order to accurately understand questions of fare policy, as well as to complement the AFC sources where integral information was unavailable from AFC sources.

1.3.2 Methodologies

Different analytical methods are used to evaluate individual aspects of post-COVID fare policy in each case study presented. Detailed descriptions of each methodological approach are provided at the start of each chapter to frame the analysis that follows. A brief overview of each approach is as follows:

- **Fare equity and rider fare burdens** are evaluated in the context of Washington, DC. The findings rest on a new approach to synthesizing rider demographic information and AFC travel records. Existing fare equity is then contextualized against a suite of realistic, revenue neutral fare structures.

- **Post-COVID fare product utility and novel fare product design** are considered in Chicago. Travel frequency changes before and after the onset of COVID-19 are considered at the level of individual CTA users, and the utility of existing fare products are probed to illustrate the magnitude of COVID's impact on fare product choice. Multiple designs for pass products governed by more flexible regulations ('flex-passes') are developed and evaluated against existing offerings to estimate potential rider benefits. Simultaneously, the fare revenue impacts under estimated product switching are considered for all hypothetical products.

- **Single-day fare capping** is analyzed as another novel fare product for the CTA. Single-ride pay-per-use fares have long been the dominant fare choice for CTA users. Potential opportunities to use a single-day fare cap to induce discretionary travel by pay-per-use travelers are explored with a focus on maintaining the primacy of longer-duration pass products. The fare implications of such an implementation and the riders most likely to benefit are determined, both geospatially and by current fare group.

- **Opportunities to boost ridership through regional fare integration** are investigated in the Chicago area, particularly between two transit providers that historically have had limited fare integration: the commuter rail system operated by Metra and the mass transit services provided by the CTA. Given a new multi-operator fare product introduced in the summer of 2022, alongside major pricing changes for monthly commuter rail passes, post-COVID price elasticities are estimated. These elasticities are then leveraged to inform the potential travel gains from expanding the number of multi-operator products offered.

1.4 Thesis Organization

This thesis is organized around four case studies conducted in Washington, DC and Chicago, IL. This chapter details the background and genesis of this research, as well as elaborating on its application.

Chapter 2 is a brief overview of the agencies involved in this research, their respective fare policy and fare structures, and additional background useful to consider this thesis.

Chapter 3 is a synthesis of the existing body of research on fare policy, fare equity, and the data sources used to investigate both.

Chapter 4 analyzes questions of fare equity and approaches to developing an equitable fare structure across rider segments.

Chapter 5 considers the impacts of COVID-19 induced travel pattern changes on fare product choice and design.

Chapter 6 analyzes potential costs and benefits of a single-day fare cap on the CTA system, as well as considers multiple cap thresholds.

Chapter 7 explores the untapped benefits to agencies and riders when fare integration is increased between transit agencies in the same region.

Chapter 8 synthesizes the findings across each case study into a set of principles and findings regarding fare policy in the post-COVID era. It summarizes the analyses completed here and suggests future work.

Chapter 2

Case Study Background

2.1 Chicago Transit Authority

2.1.1 Agency Overview

The Chicago Transit Authority ('CTA') is a public agency that provides bus and rapid rail service in Chicago, Illinois and nearby suburbs. It is one of the highest-ridership transit agencies in the United States, with a pre-pandemic average daily ridership of 1.5 million weekday boardings in 2019 [2]. Average weekday ridership stands at 874,000 in October 2022 [3]. Bus boardings comprised 55% of all CTA trips during that period. The CTA operates hundreds of bus routes and eight rail rapid transit lines. The CTA's rail rapid transit system is known as the 'L' and is designed in a radial pattern with all lines serving the downtown business core known as "The Loop". The rail network can be seen in Figure 2-1 [4].

Funding is allocated by the Regional Transportation Authority ('RTA'), which oversees the CTA and two additional public transportation providers in the Greater Chicago metropolitan area. Those two additional public transportation providers are Metra, which provides commuter rail service into Chicago via eleven rail lines serving the surrounding suburbs, and Pace Suburban Bus, which provides suburban

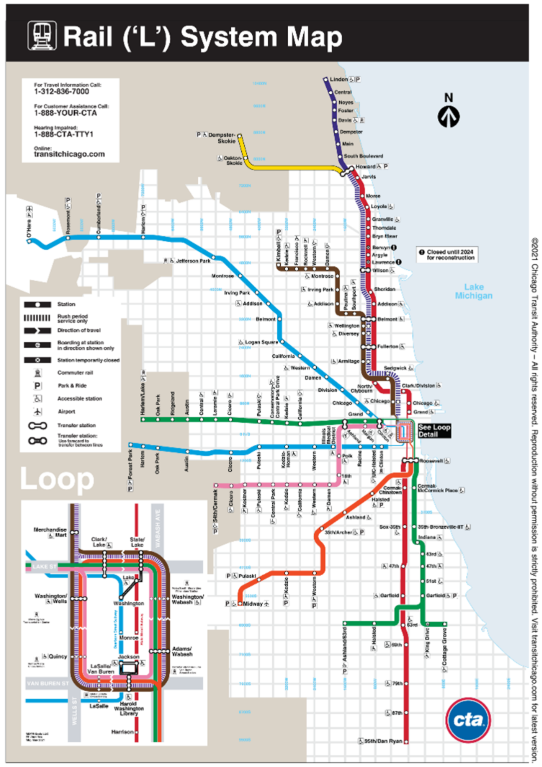

Figure 2-1: CTA Rail System Map

bus service and paratransit service across Greater Chicago.

The farebox recovery ratio for the CTA is governed by Illinois state law that stipulates a minimum of 50% of agency operating expenses must be covered by fare revenues [5]. This requirement applies to Pace and Metra, as well. Nevertheless, in 2019 the farebox recovery ratio for the CTA was 40.7% [4]. This was further eroded by pandemic ridership changes, with a farebox recovery ratio of only 18.3% for the third quarter of 2022 [6]. This value represents an increase from mid-pandemic farebox recovery levels for the same month in 2021.

2.1.2 Fare Policy at the CTA

The CTA offers many payment options, although its AFC system is the predominant choice among transit riders in Chicago. The AFC payment system is branded as Ventra; account-based smart cards and a smartphone app are both available to tap into the system. Open-loop payment technology is in place, giving the option to pay for individual trips by tapping a personal bank card on any fare validator. For unbanked individuals, or those that are unable to utilize the Ventra smartphone app, ticket vending machines ('TVMs') can be found throughout the system and there is a network of retailers acting as a sales channel. Additionally, cash payments can be made onboard buses. Ventra is also used at Metra and Pace for AFC payments. The CTA offers many pass options for different numbers of days alongside traditional pay-per-use travel. Single-ride fares are dependent on the mode utilized, with full-fare bus trips priced at $2.25 and rail trips at $2.50. Pass prices were reduced first as a temporary price reduction to attract riders back to transit in May 2021, but these promotional prices were made permanent in November 2021, with additional pass price reductions instituted at the same time.

Additional fare product improvements and pricing changes have continued to occur over time. Prior to July 2022, Metra riders that wanted to travel on the CTA in conjunction with commuter rail were able to purchase the "Link-Up" pass, which provided CTA travel only on weekdays during the morning and evening peak hours. This product cost $55, but was superseded by the "Regional Connect Pass" which was introduced in July 2022 and cost $30 for unlimited CTA travel regardless of day or time. The CTA and Pace previously offered a small number of pass products that were usable on both systems, each of which required a slight premium as compared to the CTA only versions. However, in December 2022 it was announced that all CTA passes would become jointly valid on CTA and Pace at the single, lower price point. A selection of CTA fare products and their prices as of January 2023 can be seen in Table 2.1.

Fares	Full	Reduced	Student
Rail	2.50	1.25	0.75
Bus	2.25	1.10	0.75
Unlimited ride passes			
1-Day CTA	5.00	*n/a*	*n/a*
3-Day CTA	5.00	*n/a*	*n/a*
7-Day CTA	5.00	*n/a*	*n/a*
7-Day CTA + Pace	5.00	*n/a*	*n/a*
30-Day CTA + Pace	5.00	*n/a*	*n/a*

Table 2.1: CTA Fare Prices, January 2023

2.2 Washington Metropolitan Area Transit Authority

2.2.1 Agency Overview

The Washington Metropolitan Area Transit Authority ('WMATA') is a public transportation agency headquartered in Washington, DC that provides rail rapid transit (branded 'Metrorail'), fixed-route bus ('Metrobus'), and paratransit service ('MetroAccess') across the Washington, DC area. In August 2022 average weekday ridership on the system was 500,000, with 52% of boardings occurring on Metrobus and 48% on Metrorail [7]. In August of 2019, before the COVID-19 pandemic, average weekday ridership on the system was 946,000 with 36% of boardings occurring on Metrobus and 64% occurring on Metrorail [8]. While WMATA does not have a legally set target for farebox recovery, the recovery ratio in the pre-pandemic period was 33% in 2019 [9].

WMATA is under the joint jurisdiction of Virginia, Maryland, and Washington, DC. It is governed by a board composed of directors appointed by one of the three governments. There are other small local transit providers in the region that primarily provide bus service (e.g., DASH, DC Circulator, and others) as well as commuter rail systems operated by Virginia (Virginia Railway Express, or 'VRE') and Maryland (Maryland Area Rail Commuter, or 'MARC'). WMATA's Metrorail network is composed of six lines providing service to Washington, DC and surrounding counties in Maryland and Virginia. The rail network is considered by some a hybrid subway/commuter rail system given the length of its rail lines. See Figure 2-2 for a visual representation of the network [10].

Figure 2-2: WMATA Rail System Map

2.2.2 Fare Policy at WMATA

At WMATA, the fare structures for bus and rail trips differ. Local Metrobus travel is charged a flat $2.00 fare for full-price riders. On Metrorail, the cost of a single trip varies based on the distance traveled, the time the trip is taken, the day of the week, and until recently, whether a transfer to bus was needed as part of the same journey. Reduced fares are also available to riders based on age and disability status.

Metrobus riders are able to pay their fares using either the WMATA smart card payment system, called SmarTrip, or cash upon boarding. Metrorail requires a tap-in and tap-out on each trip using SmarTrip. As a SmarTrip card is used to enter and exit each Metrorail station, AFC data provides information on the origin and destination

Table 2.2: Metrorail Fare Inputs

[a]	Peak ($)	Off-Peak ($)
Minimum fare (first 3 miles)	2.250	2.000
Per-mile rate 1 (miles 3-6)	0.326	0.244
Per-mile rate 2 (miles 6+)	0.288	0.216
Maximum fare	6.000	3.850

[a]Effective July 2017 on Weekdays; as of 27 June 2022 weeknight travel after 9:30 pm is charged a flat fare

station for each trip on Metrorail, and whether that trip is linked to bus travel on the same journey. As there are no alternative fare media on Metrorail, utilizing AFC data captures all paying rail riders.

At the time of writing, Metrorail's fare structure is distance-based on weekdays before 9:30 pm with flat $2.00 fares after 9:30 pm and on weekends. During weekdays, fares are further differentiated by time of travel, with higher prices in the peak. For both peak and off-peak travel on weekdays, the process of calculating the required fare remains the same, but the base costs, incremental charges, and maximum fares differ. When tapping out of the Metrorail system at the end of a trip, the distance from the origin station is used to calculate the mileage cost, which is rounded to the nearest $0.05. A minimum fare covers the first 3 miles of travel, beyond which a distance-based charge is levied. The per-mile cost is charged at one level for miles three to six; distances beyond six miles incur a lower per-mile rate until the maximum fare is reached. Table 2.2 shows the minimum, maximum, and per-mile rates.

The distance measure used to calculate trip length for fare purposes is the mean of the straight-line ("crow flies") distance and the distance along Metrorail's rail network ("track-miles" distance) between the origin and destination stations. This mean distance is termed the "composite distance"; its use as the measure of distance influences fares for trips of differing directness.

Given the complexity of Metrorail cost per ride, monthly pass products offered by WMATA do not have a uniform price; riders must choose the tier that covers their travel distances. This has led to low adoption of Metrorail pass products at the agency. Conversely, weekly passes conferring unlimited travel on Metrobus are offered at a single rate. Unlimited use Metrorail passes are also offered for 1-day, 3-day, and 7-day periods, but these passes are priced using a baseline assumption of longer distance journeys, resulting in high pass-multiples for short-distance riders.

Chapter 3

Literature Review

3.1 Fare Structures and the Cost to Ride

While riders must pay a fee to use public transit, that cost is not uniform across or even within transit networks. The fare is required of the traveler, but the set of policies governing how that fare is determined are selected and approved by each transit agency individually, which can result in wildly disparate payment schemes. One obvious example in the context of this thesis is the contrasting approach to fares chosen by WMATA and the CTA. For rail trips, WMATA uses a distance-based pricing structure with peak-pricing differentiation whereas a flat fare is charged its bus. Compare this to the CTA, which has flat fares for both bus and rail services, but at different rates.

These examples help illustrate some of the many dimensions along which transit fares can be determined. Many aspects of transit travel that are commonly considered when determining the cost to can be sorted into five general categories: 1) journey distance, 2) time of travel, 3) mode of travel, 4) rider demographics, and 5) travel path characteristics. These categories cover the questions Walker suggests each agency must answer when setting fares [11]. Those questions are: Should people pay more to travel farther? Should transfers require an extra charge? What is the role and extent of discounts for frequent use? Should there be pricing differentiation based on system capacity?

Previous explorations of fare policy and payment technology have identified the functional aspects of a fare system [12], summarized by Stuntz as consisting of four distinct elements: the fare structure, the fare level, product distribution and fare collection, and fare validation and enforcement [13]. These elements demonstrate how fare policy extends beyond determining the fare for a given trip. Within these four functional areas, agencies make a variety of decisions that collectively form the fare policy and the fare structure that customers experience.

When determining fares, it is also to remember that there is a natural and constant tension between two primary goals of fare policy, each of which exists in contrast to the other. The purpose of transit agencies is to transport people, and lower fares encourage ridership. However, there are significant costs associated with operating mass transit systems; the revenues collected from riders in the form of fares help pay for the continued provision of transit services. Given this tension, although not solely because of it, perspectives on the primary goals of fare policy differ. TransitCenter's 2019 report A Fare Framework specifically laid out policies they find counterproductive to equitable transportation [14]. These policies are listed as:

- "Transfer fees that penalize riders because their trip happens to involve a connection
- Distance-based fares that make prices difficult to discern
- Fare structures that impose higher costs per trip on riders with low incomes than on riders with means
- Criminal penalties for fare evasion that are disproportionately applied to people of color."

On the other hand, Cervero found in 1981 that adopted fare policies were inefficient in financial terms, suggesting that fare structures could be redesigned to raise more revenues for transit operators [15]. A particularly important cause for this financial inefficiency, according to Cervero, is the move towards flat fare structures. Elsewhere, Cervero [16] has noted that agency fare priorities can frequently diverge or overlap, and he lists several potential motivations for the fare policy an agency selects, such as a need to:

- "capture the cost of services;
- generate as much revenue as possible;
- reflect the quality or value of service to the user;
- promote equity objectives, such as helping the poor;
- encourage commuters to shift travel from the automobile to transit; and
- redress problems stemming from the underpricing of the automobile."

While there were many changes between these two publications (e.g., land use, residential locations, transit ridership, etc.), they illustrate the dichotomy of fare pricing. Both of these publications are part of the broad literature on fare policy. Cervero touches on other aspects of fare policy in the same paper, exploring not only how fare structures impact agency revenues but also how they correspond to rider benefits. Nevertheless, Cervero and TransitCenter take a different focus and come to different conclusions, including whether pricing differentiation is positive (Cervero) or negative (TransitCenter).

3.2 Methods and Data Sources for Fares Research

In order to analyze fare policy at the level of individual riders, variables that allow for differentiation between these transit users is imperative. The advent of Automated Fare Collection ('AFC') systems, and their widespread adoption, has made such differentiation possible for a majority of transit users. Riders that travel using limited tickets or pay cash fares do not have any unique, anonymized identifier attached to their travel records making identification difficult if not impossible. The number of such riders has decreased steadily over time, however. There have been other advances in fare technology, including the ability to pay for transit fares with contactless bank cards. These bank cards are also differentiable in AFC records.

Consumer preference for fare media has been explored in the past [17], but the ability to pay by tapping a debit or credit card varies across agencies. In Chicago, the CTA provides the ability to tap a bank card directly on a farebox or turnstile to pay a fare, and an anonymized identifier is then assigned to that particular bank card. However, for passes and fare products other than a one-ride fare, a Ventra account is necessary. Such Ventra account-based transactions use a different identifier for each transit user. Given the proportion of riders that have adopted Ventra, there is widespread AFC data in Chicago even without rider identifiers for cash fares and limited-use tickets. In comparison, WMATA has a closed rail system and requires use of their SmarTrip card to enter and exit all rail stations. Metrobus fares can be paid with cash in addition to using a SmarTrip card. However, open-loop payments using bankcards are not accepted. There are many useful benefits to AFC-based data records. One such opportunity of such record keeping is the ability use of AFC data to dissect ridership in novel ways, such as by fare medium or sales channel as demonstrated by Stuntz [13], or longitudinal analysis of pre-defined rider groups over long periods of time as done by Chan et al. [18].

As data collection on public transportation has become commonplace and the systems used to collect and record that data have become standardized, additional opportunities to leverage these sources have been developed. AFC systems record fare transactions associated with individual travelers as identified through the payment card used as previously described. Automatic Vehicle Location ('AVL') systems are installed on transit vehicles and record vehicle location at discrete intervals. Beyond solely tracking the ways in which individuals pay their fare, the records generated by these systems, supplemented by additional automated data collection systems, have been used to generate detailed datasets used by transit agencies to better understand how the public accesses and utilizes transit. Wilson et al. demonstrated the application of these data sources to estimate origin and destination pairs for transit users in Chicago [19]. Gordon et al. extended on this work in the context of London, demonstrating that AVL and AFC data could be used in conjunction to not only estimate origin and destination matrices for transit users, but could further be used to estimate the transfer points within a transit network along the path of each rider's journey [20]. Gordon's Origin-Destination-Transfer ('ODX') model was found to be both accurate when compared to the ground-truth provided by survey responses as well as efficient enough to populate transit records on a daily basis. Later work by Gordon, Koutsopoulos, and Wilson succeeded in scaling the results of the ODX ridership matrices by total system ridership to accurately estimate the travel paths of riders whose travel was not inferable either because the rider did not pay using a fare card or because some records required for inference using the ODX model were missing [21]. Later work on the ODX model by Sanchez-Martinez resulted in enhancements to improve accuracy and provide efficiency gains when performing the ODX inference procedure [22]. Taken together, these strides have dramatically increased in the ability of public transportation providers to understand how riders actually move through the transit networks.

There is a substantial benefit to understanding the passenger flows available using the ODX model as opposed to traditional metrics such as, say, only knowing the average vehicle load at specific nodes in the transit network. Potential applications to agency planning is clear – understanding where people start, end, and transfer helps direct the appropriate resources throughout the network. Vanderwaart et al. used the availability of ODX data for MBTA travelers to demonstrate a new approach to evaluating design of bus networks and transit service planning [23]. The applications of this model are not restricted to service planning, however. Understanding the travel paths of transit users augments traditional AFC data and allows insight on the exact impacts of fare policy decisions given rider travel patterns. Accordingly, the case studies in this thesis will utilize travel information from the ODX model to explore fare equity at the level of individual riders, who can be aggregated into various rider segments to understand how fares impact different rider groups as desired.

After identifying rider travel and capturing fare payments, the extensive data available through AFC systems can be used to model customer sensitivity to fare changes. Transit agencies have a history of using the economic concept of elasticity to anticipate travel volume changes in response to fare policy. Such elasticity measures can be used to estimate rider responses to any pricing changes undertaken by the agency. The economic concept of price elasticity is well-suited for these kinds of analysis and has been adapted and fully incorporated into the field of transportation. The Demand for Public Transport: A Practical Guide, provides a detailed review of the historical transit fare elasticity landscape including elasticity differences by travel mode, by region, and by urban population [24]. This guide also details different methodologies used to calculate transit fare elasticities, and highlights where elasticity studies have differentiated based on rider and travel characteristics. For instance, estimated elasticities have been found to differ based on the original fare level, with a proportional price change having larger impacts if the original price was already set at a high level. Elasticity studies have also found that the magnitude of the change can impact results, with larger changes returning higher estimated elasticities. And, finally, the guidebook highlights that studies are not in agreement on whether transit fare elasticities are symmetrical regardless of the direction of price change; some studies find that the elasticity is greater for a price increase than a price decrease, while others found the elasticities were the same regardless.

A second overview of transit elasticity studies by Litman argues that frequently used elasticity estimates originate from studies that are now many decades old and were first conducted when incomes and transportation options were very different, resulting in underestimates when applied in today's context [25]. This paper further argues that the studies used to generate widely accepted transit elasticity estimates predominately estimated short-run and medium-run impacts, whereas long-term elasticities tend to be greater and are potentially more applicable when considering transit's potential to support long-term societal restructuring. Litman reconsiders analyses of transit fares and ultimately suggests an overall transit elasticity range of -0.2 to -0.5 in the short term and -0.6 to -0.9 in the long term.

Regarding differences between short-run and long-run elasticity estimates, Schmiek finds that the temporal difference between these two periods is small, with 90% of adjustment to long-run equilibrium taking place within 1.5 to 3 years [26]. Chen et al. explores how service, gasoline prices, and fare levels impacts transit ridership and estimate cross-elasticities [27]. The authors find the short-run fare elasticity to be -0.4 and in the long-term -0.8. The extensive research that has been done on transit fare elasticities generally support fare elasticity being less than unit elastic, with long-term elasticities increasing from short-term counterparts.

3.3 Fare Products and Sales Channels

Within a transit system, the cost to travel can be affected by the selection of products offered. When people hear "transit fare" they often think of the cost to take a single trip on a system. However, fare payment for each individual trip taken is not the only way to access many transit systems. Additional fare products are commonly offered, although the variety of products, alongside their costs, are determined by individual transit agencies. Some common fare products include the pay-per-use single-ride ticket, calendar-based pass products available for a flat rate to all customers (e.g., 1-day passes, monthly passes, etc.), and discounted fare offerings for certain rider classes (e.g., students, senior citizens, etc.).

Common pass products offer potential benefits and drawbacks to both the transit rider and the local agency. The rider generally must first predict how many times they will ride transit over the period of the pass offering, second determine if they believe will travel enough to make passes a beneficial purchase, and then finally pay the full price of the product upfront. After purchasing a pass, the marginal cost of each transit trip reaches $0, which economics tells us should encourage increased use. The drawbacks to pass purchases are the inherent uncertainty of future travel expectations and the often high upfront cost for a pass. The uncertainty of transit travel has only increased for most individuals following the COVID-19 pandemic, and some agencies have responded by lowering pass prices while other agencies have not.

Focusing on the power of passes to spur ridership, it is key to realize that lowering the marginal cost of each transit trip to zero reduces the price "salience". Salience is the prominence of a cost to a consumer and plays a significant role in purchasing decisions. Behavioral and economic research has consistently shown that the more aware individuals are of a cost, the more that cost impacts demand. Rosenfeld demonstrated this finding clearly in the arena of car commuting, where Boston-area employers switched from monthly to daily parking passes, requiring staff to choose to pay each day they chose to drive resulted in a clear decrease in parking demand [28]. This finding has been replicated elsewhere, and there is no economic reason the same findings cannot be extended to transit fare products. In fact, given the low marginal costs provided by pass products, and the importance of price salience to consumers, Stuntz demonstrated the opportunities that exist for transit agencies to increase ridership levels without large corresponding decreases in fare revenues by transitioning customers away from pay-per-use travel and onto pass products [13]. Current levels of post-pandemic transit use offer an opportunity to exploit these findings. The amount of transit service that agencies continue to provide has resulted in excess capacity that can be used to absorb additional ridership without an immediate need to increase service levels. This would allow higher ridership at the same operating cost, which can help support arguments for continued transit funding from local, regional, and national sources.

The impacts of fare products and fare offerings on ridership, and not solely fare levels, is an important consideration. Changes in fare levels can be analyzed using elasticities as previously described and can be used to reliably measure changes in travel volumes at all levels, from the individual to the entire riding public. Some studies have explored a variety of options for fare pricing differentiation that smart card technology made feasible [29]. While such differentiation is possible, not every transit agency has chosen to implement extensive pricing differentiation. Further, the impacts of fare policy on non-riders must also be considered. It is possible these riders have different sensitivity to transit prices than individuals already on transit. One study out of Australia analyzed a fare restructuring in Southeast Queensland that simplified the zonal system, lowered fares, and added usage incentives [30]. This study found this restructuring spurred much more travel among existing users, but brought new users onto the transit network at lesser levels.

Similarly, the fare product offerings themselves, and the way those offerings interact with each other, can be equally powerful in affecting transit use. For example, Allentown, PA changed their fare structure in 1987 by raising the price of single-ride pay-per-use fares from $0.50 to $0.75 while offering pre-paid tickets valid for ten transit trips at a cost of $5.00 [31]. This equates to $0.50 per trip, equal to the previous single-ride fare. Six months after this pricing change, the local transit agency saw increases in farebox revenues and ridership, with 67% of trips utilizing the pre-paid 10-trip ticket.

Another avenue for lowering the cost of transit use is available to employees through the transit tax benefit created by the US federal government. Through this program, companies can offer employees tax-advantaged transportation benefits although participation is not required. These tax-advantaged benefits can be fully paid for by an employer, they can be fully paid for by the employee through a pre-tax payroll deduction, or they can be partially paid for by the employer with employee pre-tax payroll deductions covering the remainder [32]. This program can substantially reduce the effective cost of transit products, reducing prices by the level of an individual's income tax burden. The monthly limit for pre-tax deductions has increased over time to reflect changes in transit fares. It is not only transit users that benefit from such programs, however; transit agencies are receiving the full fare due from users enrolled in such programs. Riders experience a lower effective fare, which incentivizes increased transit travel which in turn leads to higher agency revenues, which agencies collect at the full posted price, which is higher than that experienced by the riders involved in these programs. Transit benefits programs are a win-win for riders and agencies. They are an even bigger win when employers choose to cover transit expenses as a benefit for employees. Enrollment in this program are only available through an employer, however, and employers are not required by the US federal government to offer transit benefits. In fact, a relatively recent study found that only 7% of US employees have access to transit benefits with only 2% using them [33], which seems even smaller when one considers the population that is not in the workforce. Usage of these programs is not uniform across agencies, with some regions seeing higher employer engagement than others. All of this conveys the need to increase the number of organizations and individuals enrolled in these programs. The potential for significant increases in transit ridership and transit revenues has so far been underexploited, and deserves significant attention.

Research has supported these assumptions. As one would anticipate, the availability of commute benefits has been found to increase transit usage in a variety of studies. A 2005 Transportation Research Board report on commuter benefits found across survey data that most employers introducing transit benefits see an uptick of at least 10% in transit use [34]. The impacts differed, with low-transit-usage areas seeing relatively larger gains in new ridership while a few regions saw most of the benefits flow to existing transit users [34]. In 2014, TransitCenter published a report that studied the attitudes and behaviors of national transit users; the report found that individuals offered pre-tax commute benefits have a much higher probability of regular transit use, being five times as likely to be regular riders [35]. A 2017 study focused on 22,000 commuters in the New York-New Jersey metro area included individuals with differing benefits and concluded that employer transit-pass and bike benefit programs increased travel share for those modes. This program can have not only a large impact on transit cost, but also transit use. It is important to remember that it is only available through participating employers, restricting its reach to the wider spectrum of ridership.

However, two issues arise from the focusing on pass products to increase transit. The first issue with utilizing traditional pass products to boost ridership in a post-COVID environment is that rider travel patterns have changed. Historically passes were designed and priced to attract commuters. For a given pass price and single-trip cost there is a corresponding point at which the cost of the pass breaks even with the cost of paying for each trip. This number is called the pass “multiple”, and changes in either the cost of the pass or the cost of a single trip can impact the usefulness of a pass product to a given rider. In the post-COVID environment, many individuals are commuting less due to hybrid work schedules or fully-remote work. Correspondingly, a pass multiple that previously would have been seen as reasonable may be substantially too high for many former pass users.

The second issue is a longstanding concern with the equity of collecting payment upfront to confer lower per-trip costs. Low-income riders who are dependent on public transportation have less ability to pay large upfront sums in order to access the most beneficial fare pass products, as demonstrated by Agrawal et al. [36]. To accommodate these riders, researchers have made many proposals, one of the most common being some way of dividing large lump-sum payments into smaller payments made more frequently to even the burden across income levels. One implementation of this approach is fare capping, wherein a rider pays for each trip until they have paid fares equivalent to the price of a pass. The equity impacts are obvious, but it contradicts a goal of spurring ridership by reducing the marginal cost of each trip to $0. In practice, the marginal cost remains the full cost of a trip until the pass multiple is achieved, at which point additional trips finally incur no further costs.

3.4 Fare Equity

Equity in transportation is a key consideration for every transit system, and its importance has become clearer in recent years. Drops in transit ridership during the pandemic were not evenly distributed across socioeconomic groups. Hu and Chen found that low-income riders, riders of color, and riders without a college degree continued riding transit at higher frequencies than their counterparts [37]. A survey by the Transit app [38] and a study by Liu et al. [39] came to very similar conclusions, both finding that early in the COVID-19 pandemic essential workers, females, riders of color, and low-income individuals were continuing to ride transit at higher rates. The importance of essential workers is a key ingredient in all of these findings, as another study corroborated the findings that disadvantaged communities remained more reliant on mass transit, but also found that adjusting for the number of essential workers in each neighborhood controlled for these findings [40]. Brough et al. analyzed transit use data in King County, Washington, where Seattle is located, and found the largest travel decreases were among those with higher educational attainment and higher-incomes [41]. The travel volume discrepancies that developed during COVID remained stable as all groups began traveling more as things began once again reopening, suggesting it may be an ongoing feature.

The legally defined measure of equity for public transportation networks rests on the application of Title VI of the Civil Rights Act of 1964. Through this act, each agency is required to formally analyze how any major change to any of transit infrastructure, transit service, or transit fares will affect the population, and to ensure that there will be no disparate impacts on demographic groups defined by race and income when compared to the rest of the population [42]. It is important to note that Title VI analyses only measures changes in equity; it has no requirements regarding the existing equity. Compelling suggestions on how to remedy this shortcoming and expand the analysis to current system equity are discussed throughout the existing research literature [43].

While there is only a single legal requirement regarding equity, there are numerous definitions of equity across the literature. An important aspect of the wide-ranging definitions is that application or different definitions can lead to conflicting findings and contradictory suggestions for improvement. Equity has been defined as procedural, considering who is represented in the policy and planning process; distributional, comparing how costs and benefits are borne by different populations; structural, where accountability for decision-making is institutionalized and long-term historical, cultural, and institutional dynamics are recognized as they related to privileged and disadvantaged groups; and transgenerational, concerned with long-term implications of policy decisions as they affect future generations [44]. Taylor further utilizes market equity, opportunity equity, and outcome equity as primary considerations [45]. Market equity occurs when the services provided are proportional to the revenues collected; opportunity equity considers whether a uniform level of investment is made in all areas; and outcome equity is achieved when access, mobility, or transit service are equal for all users. Observing all of the ways equity is designated, it is clear that application of disparate definitions will often result in contradictory findings.

These contradictory definitions in the literature have real-world analogs that can be observed in discussions around agency funding decisions. Consider the boards governing the CTA and WMATA. These boards approve budgetary decisions, and they are composed of members representing different levels of government in each agency's service area. WMATA's board is composed of eight voting directors, with an even number of directors appointed by Maryland, Virginia, the District of Columbia, and the US federal government. The CTA board is composed of seven members, with four appointed by the mayor of Chicago and three appointed by the governor of Illinois. In both instances, the governments that provide funding also appoint members to set agency policy. However, transit service levels are not provided and funding is not generated in perfect proportion to board representation. This can, and has, resulted in differences of opinion regarding whether those providing funding are reaping the benefits of transit to the proper extent. Such tension is symbolic of the many differing conceptions of equity defined by TransitCenter [44] and Taylor [45].

Moving beyond defining what is meant by equity, the existing literature also considers the unit of measurement used for analysis. How should individuals and groups be aggregated, and to what purpose, are key questions for interpreting findings regarding equity. Three units of scale are established by Taylor: individual equity is evaluated at the person-level; group equity compares segments individuals clustered by similar characteristics; and geographic equity is focused on measurement across locales [45].

Although many definitions of equity exist, to operationalize one for analysis necessitate that a unit of scale and evaluation criteria must be selected. Regarding scale of analysis, fare equity studies are most commonly at the group-level to allow comparison across defined rider segments. Selection of a measurement criteria has been less uniform. Cervero has identified diverse criteria that could be used to measure the equity of a fare structure [15]. Nuworsoo et al. adapts and operationalizes these criteria as the "benefit criterion", which holds that fares should be proportional to the benefit the rider receives; the "cost criterion", which holds that a trip's fare should be proportional to the cost of providing that trip; and the "ability to pay criterion", which holds that fares should be proportional to the rider's wealth [46].

The choice of how to define and quantify the benefits of transit leads to divergent results when applying the benefit criterion. For example, one can define the benefit as a "trip", in which case flat fares are judged equitable, or as the travel distance, which requires uniform per-mile rates to be considered equitable. Methodologies used to assess fare equity according to the benefits criterion have used metrics including fare paid per mile [47], the average transit fare [48], and changes in both the average fare and the number of trips taken [46].

There have been a number of studies that assess the equity of either existing fare structures or hypothetical fare structures to understand the interaction of agency ridership, agency fare policies, and the corresponding fare equity. Brown developed six potential fare structures for LA Metro and estimated corresponding fares per mile using survey responses from the 2012 California Household Travel Survey [47]. Nuworsoo et al. used on-board rider surveys to estimate the equity and ridership impacts of five fare structures that were being considered by AC Transit in Oakland, CA [46]. Using the 2012 Utah Household Travel Survey, Farber et al. created a model to understand the joint decisions to travel by transit and the daily distance traveled by transit [48]. The purpose of the developed model was to determine whether a switch to distance-based fares would improve rider equity for the Salt Lake region. For an international example, Rubensson et al. analyzed flat, zonal, and distance-based fare options within Stockholm, Sweden using census data [49].

There is no consistent finding across these studies regarding whether flat or distance-based fares are better for equity. Instead, they all illustrate the importance of the local geographies to understand fare equity for a given agency. It is also important to note the types of comparisons made are not consistent across existing studies. For instance, Brown has developed alternate fare structures for comparison but did not discuss the overall fare revenue differences across hypotheticals. Contrast this to Nuworsoo who is comparing possibilities proposed by AC Transit and specifically evaluating revenue impacts, among other aspects.

A separate approach to a more comprehensive understanding of fare equity was undertaken by El-Geneidy et al. in which travel costs inclusive of travel time and transit fare were compiled to understand social equity [50]. The authors considered total travel costs to be the combination of required travel time and monetary costs. These total travel costs were evaluated against the accessibility of jobs for different rider segments. This study is interesting but is focused on accessibility for rider groups as opposed to actual travel patterns exhibited by demographic groups. Further, its focus on jobs does not consider transit use for non-work trip making. Although the importance of transit fares is integral to this work, it is ultimately an extension of the accessibility literature more so than a review of fare equity specifically.

Within this thesis, I focus on group-level analysis of fare equity as it applies to transit agency networks. Within the services areas of both WMATA and the CTA, group-level analysis sometimes overlaps with geographic analysis given the distribution of demographic groups. Further, this thesis conceives of fare equity as being rider-centric; it therefore focuses on the distributional impacts of transit, considering how its benefits and costs are borne by rider segments. The benefits of transit in terms of fare equity are defined as the fare per linked trip, which is closest in line to how riders experience transit fares. Transit is used to access amenities, not as an opportunity to travel a given distance.

Chapter 4

Analyzing Fare Equity

4.1 Introduction

Washington Metropolitan Area Transit Authority, as described in Chapter 2, uses a distance-based fare with peak-pricing differentiation for its Metrorail service. The Metrorail system is uncommon among US subway systems in that it utilizes both trip characteristics to determine the cost of a trip. However, it is unclear if this pricing regime results in unequal cost burdens for subpopulations of the Metrorail ridership.

When evaluating the distributional effects of a given fare policy, and the fare structure features driving those effects, transit agencies regularly face a dearth of rider demographic information and a challenge in connecting trip characteristics to rider characteristics. Traditional survey-based analysis can provide detailed demographic information for transit users, but the sample sizes are often relatively small. In this chapter a case study is presented that amalgamates multiple data sources containing rider-reported demographic information and detailed AFC-based travel records[1]. This method is applied to Metrorail to understand the equity outcomes of its existing fare

[1] While this chapter develops and applies a methodology to understand fare equity using ridership data from the years preceding COVID-19, the findings are reaffirmed with updated post-COVID ridership data as described in Appendix A The appendix explores two approaches to updating the pre-COVID travel data used in this chapter, with consistent findings between both of those approaches and these pre-COVID findings as explained more fully in Appendix A.

structure.

Fare equity is considered here through comparisons of the fare burdens felt by ridership segments as defined by race and income characteristics. To fully contextualize the impact of the chosen fare structure on equity outcomes, a suite of alternate fare structures are developed. The fares under each alternate structure are calculated to be revenue neutral assuming constant ridership, although this constraint is later relaxed to capture realistic rider responses to fare changes. The two key elements of current Metrorail fares are the distance-based structure and peak pricing differentiation; the alternate fare structures are correspondingly designed to explore these two aspects, and how different implementations would impact the fare burden felt by rider groups. Comparison between race- and income-based rider segments help illustrate how residential locations, travel patterns, and fare structure mediate the equity of the fare system as experienced by historically disadvantaged riders.

4.2 Research Goals

The case study provided in this chapter provides clarity on the equity of WMATA's existing Metrorail fare structure. It seeks to determine the fare burden felt by low-income riders and riders of color, and then to contrast those results against higher-income and white counterparts. Through this process, the case study determines a set of travel characteristics that can be employed by the fare structure to improve equitable outcomes low-income riders and non-white riders.

The question of distance-based fares in the Washington, D.C. region is especially important given the complexity of WMATA rail fares. Though this question has received substantial attention in the popular press and transit advocacy community [51, 52, 53], it has not enjoyed equivalent scholarly attention. In considering alternative fare scenarios, agency staff have defended the economic logic and fairness of the rail fare structure, and raised potential Title VI equity concerns should distance-based fares be replaced with a flat fare in a revenue-neutral fashion [54]. Such an assertion can be scrutinized more completely in light of the analysis undertake in this case study.

The argument that flat fares raise a Title VI equity concern rests on a combination of two travel behavior assumptions: that low-income Washingtonians disproportionately take shorter trips and that they disproportionately travel off-peak. However, it is important to empirically test these assumptions, especially in light of changing demographic patterns. Demographic trends, including in the Washington, D.C. region, show that suburban areas are becoming less wealthy, less white, and less-educated relative to the metropolitan area as a whole [55], substantiating what Ehrenhalt terms the "great inversion" [56] or what Renn terms "the new donut" [57]: the idea that the revitalization and gentrification of sections of the urban core and the suburbanization of poverty create a "donut" of poverty with a wealthy outer ring and a wealthy "donut hole" in the city center. The suburbanization of poverty results in the increased concentration of lower-income populations in more auto-centric areas, thereby obstructing their travel and reducing their access [58]. These demographic trends and their adverse effects point to the need for progressive policies, including a reevaluation of fare structures.

4.3 Methodology

4.3.1 Equity Metrics

To assess fare equity, the measurement criteria must be established. A set of equity metrics are necessary to indicate how fare burden is currently felt across Metrorail users. Further, selecting the proper equity metrics allows comparisons across existing and alternate fare structures. Such comparisons can be used to identify aspects of a fare structure that provide benefits to disadvantaged rider groups, or fare structures to avoid. The case study in this chapter is driven by a focus on understanding the way costs are felt by riders. Three distinct equity metrics based on the rider experience have been selected to guide the understanding of fare equity:

1 **Fare per linked trip** captures the direct monetary costs of traveling on the system.

2 **Fare differential** across rider groups compares the average fares paid by different rider segments.

3 **Trip volume changes** account for the expected changes in travel by riders given a new fare level for their trips. Increases or decreases in travel volume are indicative of rider benefits, as trips not taken on transit due to the fare burden are a form of inequity that are not captured without a comparison to some baseline level of travel, which in this chapter is set as the current Metrorail fare structure.

4.3.2 Metrorail Travel with Known Rider Demographics

As discussed in Chapter 3, many previous analyses exploring the equity of fare structures have generally either used surveys or spatial distribution of demographics as the dataset for analysis. Survey-based analyses utilize actual transit trips reported by riders, often through travel diaries. Spatially-based analysis often rely on general geographic distributions of demographic groups; one common source of this information is census data.

In the analysis presented here, however, the focus is placed on actual trips taken by representative riders with known demographic information. This necessitates first identifying riders for whom demographic information is known with a high level of confidence, and second determining the trips taken by those riders on the Metrorail system. To compose a large, representative sample of transit users and their Metrorail travel, disparate data sources were collated to leverage the complementary aspects provided by each. These data sources include traditional household travel surveys and WMATA fare card data. Traditional travel surveys provide a good source of demographic information on respondents, but sample sizes are often small, especially when there is a need to focus on subpopulations of interest or distinct rider segments. On the other hand, fare card data provides detailed travel information that can be used to create a very large sample, but that sample is often lacking in important demographic details. This case study uses a method to jointly synthesize and use the information available from both travel survey and fare card data, ensuring comparable data.

Our first source of demographic information is the Metropolitan Washington Council of Governments (MWCOG) Regional Travel Survey (RTS) [59]. RTS reports detailed travel and demographic information, including race and income, from a large sample of over 16,000 Washington-area households. Each household was assigned a single weekday between October 2017 and December 2018. On that day, information on all journeys and travel modes was gathered for each respondent. When reporting a Metrorail trip, respondents provided the origin and destination rail stations. From this detailed trip-level travel information, a framework of logical rules based on stated origin station, destination station, travel mode sequence, and time of journey was used to match a representative trip from WMATA's implementation of ODX (referred to internally as 'Trace') to each RTS response record.

Our second source of demographic information is WMATA's internal SmarTrip card registration records. Riders are able to use the optional registration process to ensure transit value can be recovered from lost or stolen cards or to access tax-advantaged transit benefits. Of the SmarTrip cards that have been registered with WMATA, there is a subset for which the registrant has provided demographic information including race and income. This subset of just under 4,000 registrants, anonymized so that only race, income, a unique customer identifier, and hashed identifiers for each SmarTrip card are available, provides another avenue to link Metrorail trips in ODX with stated income and race. Many of the Metrorail customers in this registered, labeled subset of cards had registered multiple SmarTrip cards. For this reason, it was necessary to set an upper limit on the number of registered cards a user could have and still be included. The maximum number of cards registered by a single user was 886, as an extreme example. The upper limit set was 7 SmarTrip cards, which captured 3,271 unique SmarCard users.

To evaluate fare equity for riders based on race and income characteristics, the parameters for each group were established. Riders in this analysis are classified by race and ethnicity as either white or as a person of color. The latter group, also referred to as "riders of color", includes people of non-white race or of Hispanic ethnicity. Passengers are additionally classified by household income as earning above or below the livable wage. The livable wage is defined using the Living Wage Calculator for Washington, D.C., at the beginning of the RTS data period in October 2017 [60]. This calculator adjusts for regional variation in costs and spending. The income thresholds used for classification in these two data sources are not uniform. The RTS breaks income groups at $25,000 and $35,000; WMATA's registered SmarTrip card database has an income break at $30,000. Given these differences, riders whose household income is below the dataset threshold nearest the living wage of $32,686 are designated as "low-income". That corresponds to RTS respondents making less than $35,000 and registered SmarTrip users with incomes under $30,000.

Given the subset of registered SmarTrip cards with labeled demographics, the RTS travel diary process was mimicked to create a synthetic RTS dataset, ensuring a direct comparison with the collected travel diary responses. The synthetic RTS data was generated by randomly assigning each registered, labeled SmarTrip customer a single weekday during the RTS period. Once assigned, WMATA's ODX travel records database was queried for all journeys by each customer on their specified travel day. These travel records became "Synthetic RTS" entries available to be used in conjunction with RTS travel responses matched to ODX travel records.

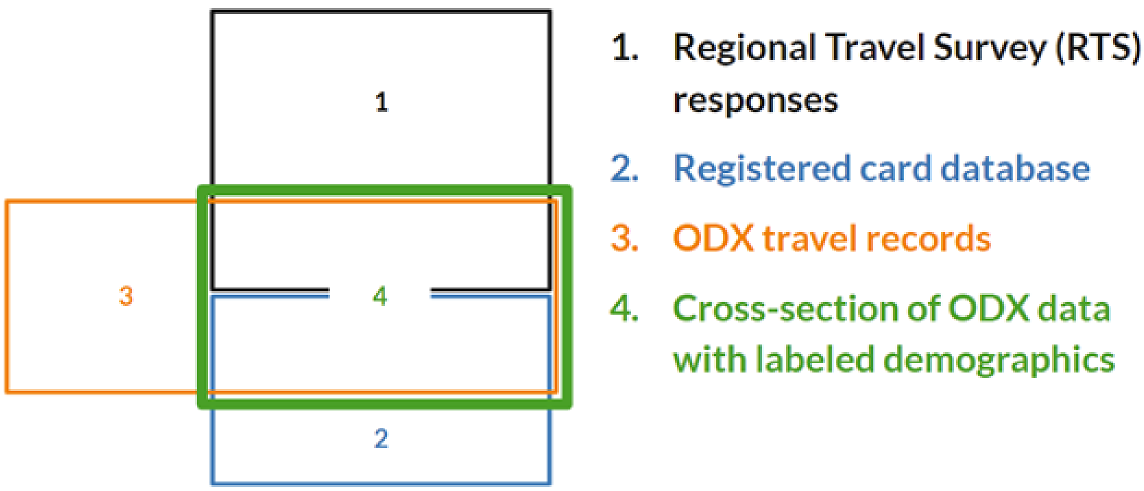

Figure 4-1: Schematic of Compiled Metrorail Trips Data

After matching travel records representative of the RTS travel records, and creating synthetic RTS data, the information was combined into a final dataset of Metrorail trips that contain the necessary travel and demographic identifiers. The resulting combined dataset is representative of systemwide Metrorail travel. It comprises 6,174 unique riders taking 12,129 Metrorail trips. 2,916 of these riders were identified using RTS responses and 3,258 were identified using registered SmarTrip cards. Similarly, 5,486 of the Metrorail trips are from RTS data; the remaining 6,643 come from registered SmarTrip cards. All 91 active Metrorail stations were represented as a trip origin and a trip destination. Within the dataset, there are 1,973 unique station pairs, representing 48.2% of potential rail station origin-destination combinations. However, origin-destination pairs do not experience equal travel volumes; correspondingly, the 1,973 station pairs captured in the sample data are responsible for 86.5% of trips taken on the Metrorail system during the RTS period.

To ensure that the sample of travel data is representative of all Metrorail riders, iterative proportional fitting ('IPF') was employed to calculate appropriate weights for each individual Metrorail trip within this dataset. The initially compiled dataset underrepresented riders below the livable wage and overrepresented the number of trips taken during the peak periods. To make sure travel was accurately represented along all dimensions relevant to this analysis weights were calculated to control for income, race, and peak travel. Each of these characteristics of travel is a necessary dimension given the focus on fare equity and a rail system with peak-pricing.

To utilize iterative proportional fitting, a three-dimensional matrix (2x2x2) was seeded with each dimension corresponding to one of the defined controls. The values in each cell were taken directly from the compiled Metrorail travel sample. The marginals for each of the controls were taken from contemporaneous Metrorail data covering all travelers. The proportion of riders by race and by income were set using the 2016 WMATA Metrorail Rider Survey. The ratio of peak to off-peak rail travel was set using total SmarTrip records on weekdays during the RTS period. IPF was then used until the values converged to the ground-truth values for ridership by race, income, and peak period. Upon convergence, the resulting weights were used to transform the travel sample such that all demographics precisely match contemporaneous Metrorail ridership in an unbiased way. Applying the calculated weights to each trip record resulted in a weighted travel sample which both matches Metrorail ridership demographics and covers the majority of all Metrorail travel in spatial terms as detailed above. The weights calculated for trips based on these three dimensions can be seen in Table 4.1.

4.3.3 Developing Alternate Fare Structures

Given a set of rider travel patterns and a formula for fares, one can easily calculate average fares across rider segments. However, that alone does not provide a complete

	Riders of Color		White Riders	
	Below Livable Wage	Above Livable Wage	Below Livable Wage	Above Livable Wage
Peak travel	3.796	1.276	2.065	0.694
Off-peak travel	4.391	1.476	2.388	0.803

Table 4.1: Calculated Weights by Metrorail Trip Dimensions

understanding of whether a fare structure is performing equitably; there is a need for other potential fare regimes against which the performance of the current fare structure can be compared. To fulfill this need it was necessary to create realistic fare structures and estimate equity metrics under each.

Alternate fare scenarios were developed through consideration of previous fare policies, review of formal and informal fare proposals, attention to prior WMATA board discussions, and conversations with WMATA staff. To ensure direct, unbiased comparisons between fare scenarios, fare levels within each fare scenario were first calculated to be revenue neutral, assuming unchanged ridership. Once fare rates were determined for each scenario the assumption of constant ridership was relaxed in order to estimate rider response.

Alternate fare scenarios were developed to consider of a diversity of realistic structures. The constructed scenarios vary the fare impacts ascribed to a Metrorail trip's distance, time of travel, and specific journey characteristics in order to isolate the importance of each. Altogether, these variations provide clarity on how the interplay between trips and fare structures combine to affect fare equity outcomes. Alternate fares were grouped into five scenarios, each comprising fare implementations that share a high-level organizing principle. A summary of each structure can be seen in Figure 4.2.

Fare structure	Type	Distance basis	Peak surcharge	Core surcharge
1a	Distance-based	Composite miles	*X*	
1b	Distance-based	Composite miles		
1c	Distance-based	Composite miles	*X*	*X*
1d	Distance-based	Composite miles	*X*	
2a	Distance-based	Crow-flies miles	*X*	
2b	Distance-based	Track miles	*X*	
3a	Zonal	Zonal difference	*X*	
3b	Zonal	Zonal difference		
3c	Zonal	Zonal boundary crossings	*X*	
3d	Zonal	Zonal boundary crossings		
4a	Flat-fare	–	*X*	
4b	Flat-fare	–		

Table 4.2: Summary of Alternate Fare Structures Analyzed

Revenue Neutral Alternate Fares

To determine the per-mile costs under all scenarios given differing measures of distance, trip-level data from WMATA's Trace model for October 2019 was used. The data comprises weekday Metrorail fares and origin-destination stations for individual trips taken by riders paying full fare rates. For scenarios that maintain peak-pricing differentiation, per-mile costs were estimated independently for peak and off-peak periods.

The process of calculating revenue-neutral fare rates for all hypothetical scenarios began from the principle that the same portion of riders should be paying the minimum and maximum fare in both peak and off-peak hours. The percentage of Metrorail trips paying a minimum fare under the current structure was calculated for both peak and off-peak periods. Under network and straight-line distance measures, the distance at which the same portion of riders would fall into the minimum fare bucket was determined (which will be lower for straight-distance fares and higher for track-based distance fares); this distance was then rounded to the nearest half-mile and deemed the point at which per-mile charges would commence. While distance-based fares allow for arbitrary cutoffs, WMATA's fare principles call for legibility and simplicity. The distance at which the maximum fare is reached was similarly calculated such that the same proportion of riders would be charged the maximum fare, although no adjustments were made to this distance. The revenues that would be generated by trips paying the minimum or maximum cost, for which riders essentially pay a flat fare, were summed. The resulting sum was subtracted from total fare revenues actually generated by Metrorail, and the remaining difference composes the actual fare revenues not yet recovered by the given alternate structure.

The next step was to partition all trips between the minimum-fare and maximum-fare distances into the set of "per-mile" trips. The miles traveled by these per-mile trips were aggregated. Each per-mile trip pays a fare made up of two components: 1) the minimum fare and 2) the additional per-mile charge. The minimum-fare component of each per-mile trip was subtracted from the remaining revenues to be covered by mid-distance trips, leaving only the fare revenues to be apportioned on a per-mile basis.

To reflect Metrorail's existing fare structure, one per-mile fare cost was calculated for the first few miles traveled beyond the minimum-fare distance and a second per-mile cost for all miles beyond. Under Metrorail's existing system, the higher per-mile charge applies to rail distances from 3 miles to 6 miles; the same principle has been applied here. The higher per-mile charge applies up to the point at which a traveler has ridden twice as far as the distance covered by the minimum fare, and the lower charge applies for all miles thereafter until the maximum fare is reached. In Metrorail's system the second, lower per-mile cost is 88% of the higher per-mile cost; this ratio is very similar in the peak, the off-peak, and at previous WMATA fare levels. Given this consistency, the same ratio between per-mile costs is applied here.

Figure 4-2 is a visual representation of the calculations described above. The x-axis displays each individual trip taken on Metrorail, not the distance of trips. It is essentially an index where trips have been ordered from shortest to longest. The y-axis is the distance-based fare for each Metrorail trip. Region 1 represents fare revenue from riders who only pay the minimum fare, regardless of miles traveled. Region 2 represents the fare revenue generated by passengers who reach the maximum fare cost. Region 3 is the minimum-fare component of trips that are charged based on distance. Region 4 is the remaining fare revenue to be recovered on a per-mile basis from the total distance traveled by trips charged between the minimum and maximum fares.

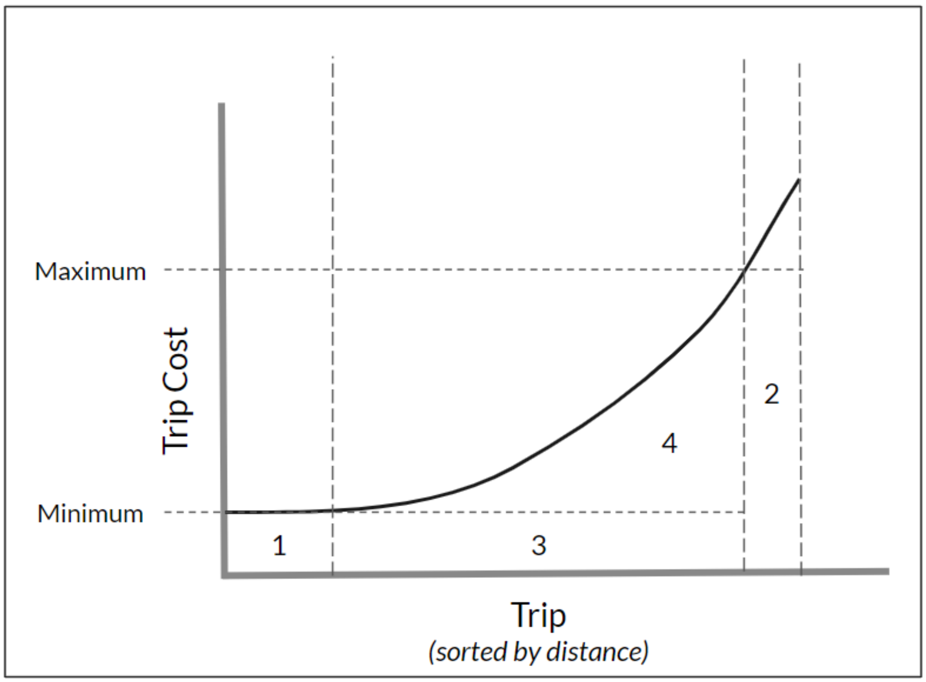

Figure 4-2: Visual Representation of Revenue Neutral Calculations

Scenario 0 - Previous WMATA Fare Structure

This scenario is a previous WMATA structure, which was replaced by the current fare structure in September 2021. Passengers transferring from bus to rail or rail to bus paid the full fare for both systems, minus a $0.50 discount. That transfer penalty was permanently eliminated in 2021, meaning that passengers now pay the rail fare and enjoy a free transfer to or from Metrobus, equivalent to a $2 discount. The net effect was to reduce the cost of a multi-modal trip by $1.50. This structure is contemporaneous to the sample of travel data, which comprises data collected prior to 2021. Under this fare scenario, multimodal trips have the transfer penalty of $1.50 intact but otherwise match the current fare structure.

Scenarios 1 - Variations on the Composite Distance-Based Structure

The current Metrorail fare structure, as previously described in the background section, is the baseline scenario and referred to as scenario 1a.

Fare structure 1b diverges from the existing structure in that there is no peak-pricing differentiation. The off-peak minimum fare ($2.00) becomes the minimum fare for all Metrorail trips. The maximum fare is set at $5.50; this is nearer the current peak maximum fare ($6.00) than the off-peak maximum fare ($3.85), although it is between the two. The first per-mile charge, for miles 3 to 6, and the second per-mile charge, for distances further than 6 miles, are calculated to be $0.312 and $0.275, respectively.

Fare structure 1c mimics the current fare structure, but incorporates a small surcharge of $0.30 for all trips entering into, terminating at, or traversing the system core during the peak periods. The core is defined as the section of the Metrorail network comprising stations that experience maximum passenger loads along each rail line. To offset this core surcharge, the minimum peak fare is reduced to $2.00. This matches the off-peak minimum, resulting in a single minimum fare. The maximum fare in the peak is adjusted down by $0.25 to $5.75. The distance-based costs are unchanged from the existing structure.

Fare structure 1d utilizes existing peak and off-peak differentiation, but with a single, consolidated fare per-mile for all Metrorail journeys beyond 3 miles. The existing peak per-mile charge for distances between 3 and 6 miles is applied to all distance-based charges. The per-mile charge for all off-peak trips is recalculated to be $0.110 given the additional revenue generated during the peak. There are no changes to the minimum or maximum fares.

Scenarios 2 - Per-mile Structures with Alternate Measures of Distance

The current WMATA Metrorail fare structure uses a composite measure of distance to determine fare, calculated as the average of two components: the crow-flies distance between the Metrorail stations at which a trip begins and ends and the miles traveled along the WMATA rail network between those stations. See Figure 4-3 for a visual explanation of these different measures. The two structures considered in this scenario each use a single component of composite distance as the distance basis. They retain the current minimum and maximum fares in both the peak and off-peak periods.

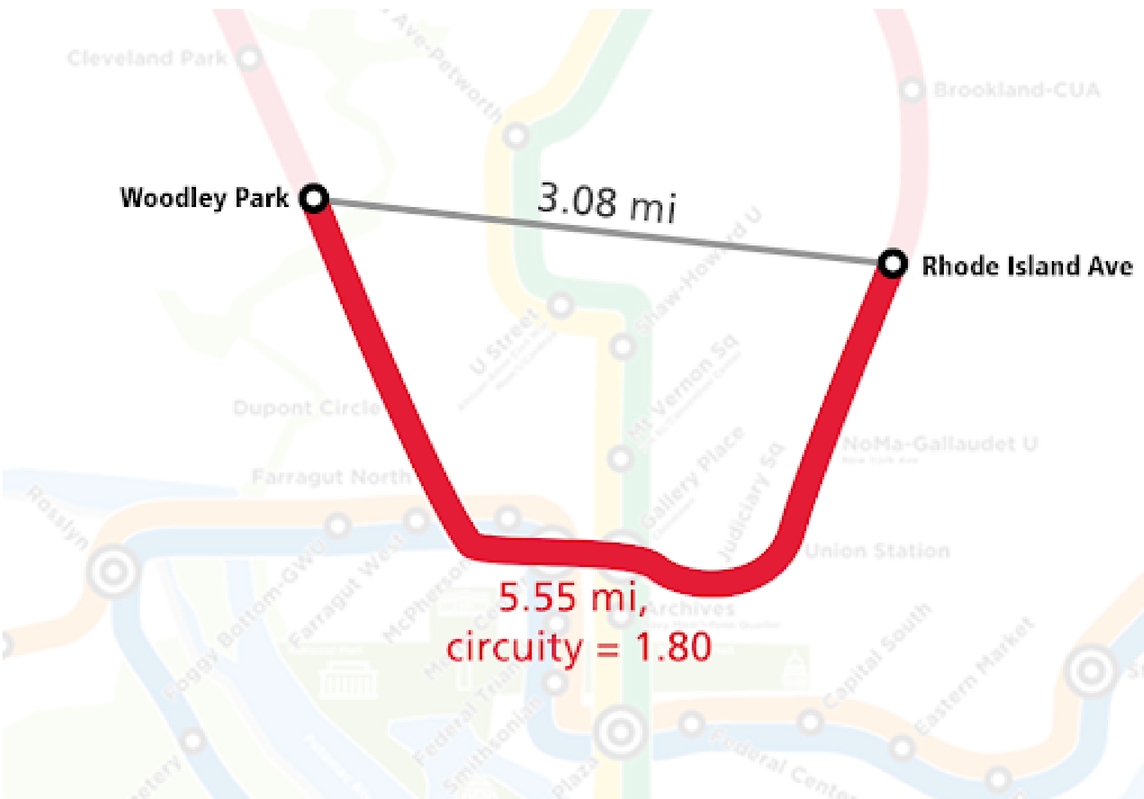

Figure 4-3: Circuity Schematic

Fare structure 2a uses the crow-flies distance between origin and destination rail stations to calculate the distance-based fees. The minimum fare covers the first 2.5 miles traveled, a higher per-mile rate ($0.348 peak / $0.254 off-peak) covers travel between 2.5 and 5 miles, and a lower per-mile rate ($0.307 peak/ $0.224 off-peak) is charged beyond 5 miles until the maximum fare is reached. Such a structure would remove any additional fare burden currently incurred for traveling between stations not directly linked.

Fare structure 2b considers track miles traveled on the Metrorail network between origin and destination as the distance basis. The track-miles distance is assumed to be the shortest path along the rail network, regardless of how close the stations are geographically. The minimum fare covers the first 3.5 miles of travel, a higher per-mile rate ($0.285 peak / $0.205 off-peak) is charged for travel between 3.5 and 7 miles, and a lower per-mile rate ($0.251 peak / $0.180) is charged beyond 7 miles until the maximum fare is reached.

Scenarios 3 - Zonal Fare Structures

Zonal fares would retain some benefits of a distance-based fare system while simplifying the calculation of station-to-station fares compared to the current system. However, such systems can have unexpected consequences, as expensive zonal boundaries make short inter-zone trips costlier than longer intra-zone trips.

Zonal systems can charge passengers based on two fare bases: the absolute difference in the zone numbers of the origin and destination station ("zonal difference fare basis") or the number of zonal boundaries crossed in the course of travel ("boundary crossings fare basis"). The former would discount longer suburb-to-suburb trips relative to trips between the system core and the suburbs, while the latter charges passengers proportionally to the total length of their trip.

In collaboration with WMATA staff, a zonal map was developed for the Metrorail system that includes five zones. We determined zonal boundaries algorithmically based on the track distance of each station to the center of the Metrorail system. The center of the system was designated as the 0.3 mile segment between Metro Center and Gallery Place-Chinatown stations. These two stations are centrally situated in Washington, DC and together serve all Metrorail lines. The distance from a given station to the center of the system is then defined as that station's track distance to the closer of the two stations. The track distance cutoffs defining the zonal boundaries were first determined using the Jenks natural breaks classification method, then heuristically adjusted to achieve specific desirable characteristics, such as restricting Zone 1 to a smaller region of the District of Columbia and ensuring that each non-terminus station is adjacent to at least one station within the same zone.

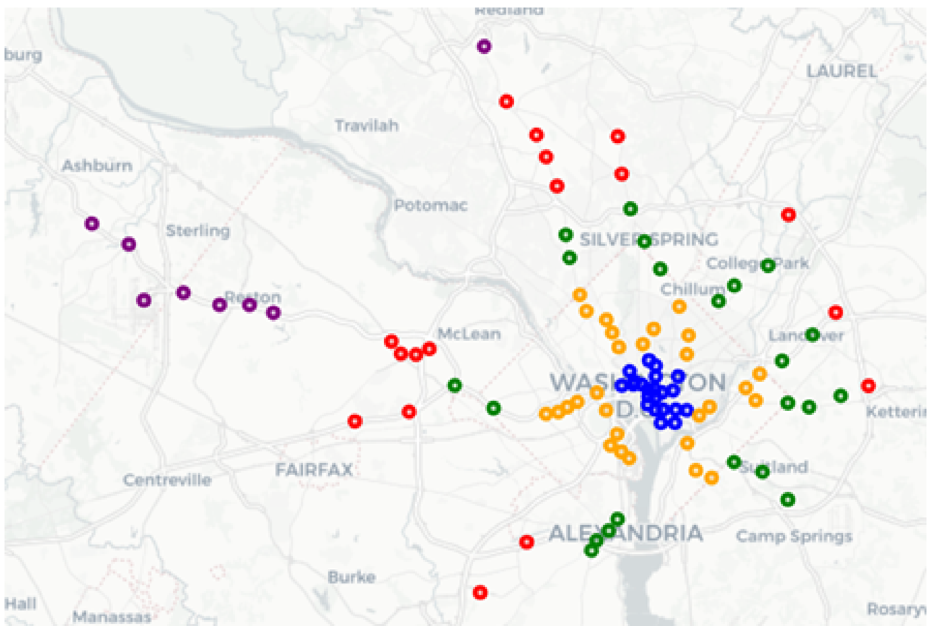

Figure 4-4: WMATA Stations by Scenario 3 Zone

Given the zonal map, as shown in Figure 4-4, and the two fare bases described, four zonal structures were developed: a zonal fare structure with zonal difference fare basis and peak pricing (3a); a zonal fare structure with zonal difference fare basis and no peak pricing (3b); a zonal fare structure with a boundary crossings fare basis and peak pricing (3c); and a zonal fare structure with a boundary crossings fare basis and no peak pricing (3d).

Under scenario 3a and 3b, passengers pay on the zonal difference fare basis, meaning a passenger who crosses 8 zonal boundaries and ends up in the same zone pays the minimum fare. Under 3a, with peak pricing, fares are designed to closely mirror the current fare system, ranging from \$2.00 to \$4.00 in the off-peak periods and \$2.25 to \$6.00 in peak periods. Under 3b, with no peak fare differentiation, the fares range from \$2.25 to \$5.05 in even \$0.70 increments.

Scenarios 3c and 3d utilize a boundary crossings basis. Fares are composed of the minimum fare, one charge that increments for each boundary crossed up to four crossings, and then a discounted charge that increments for all remaining boundary crossings. The result is that long suburb-to-suburb trips traversing the network core are costlier than the longest suburb-to-core trip. Scenario 3c has the same minimum fares in the peak and off-peak periods as the current system, but has a slightly higher maximum fare in both periods due to small additional charges for travel crossing more than four zonal boundaries (\$6.05 peak / \$4.40 off-peak). Scenario 3d follows the same logic as scenario 3c, but with no peak pricing differentiation, resulting in fares that range from \$2.00 to \$5.80.

Scenarios 4 - Flat Fare Structures

Lastly, two flat fare structures are evaluated that remove the distance-based component of Metrorail's fare structure. One scenario includes continued price differentiation between the peak and off-peak periods (4a) and the other charges a single fare for all trips regardless of distance and time of journey (4b).

To calculate the revenue neutral flat-fare equivalent with peak pricing (4a), all fare revenues collected during peak transit hours were spread evenly across all of the journeys taken during peak hours. Likewise, all fare revenues generated in the off-peak were divided among all off-peak trips. This resulted in flat fares of \$3.45 in the peak and \$2.70 in the off-peak.

Finally, the revenue neutral flat fare is determined such that all fare revenues generated by the existing distance-based structure were spread evenly across all Metrorail trips taken (4b). The calculated flat fare in this scenario is $3.20.

4.3.4 Estimating Equity Metrics Under Alternate Fare Structures

With an extensive dataset of Metrorail trips, including detailed information on transfers, time of travel, travel path along the WMATA network, and multiple measures of trip distance, the fare due under each of the developed alternate fare structures can be calculated for each individual trip. Further, since each trip record includes rider demographics, the aggregate and distributional impact can be estimated for different subpopulations of Metrorail riders.

The average fare by group under each alternate fare structure is not the only measure of fare equity, however. Changes in the number of trips are an important component of fare equity, and a key layer to understand alternate fare scenarios. The ability of riders to travel on Metrorail is not captured solely by measures of average fare; an unobservable aspect of fare equity is the number of trips not taken because the cost is prohibitive. Riders are sensitive to price changes; higher costs will result in less travel. But to gauge the extent of rider responses, and to capture differences across rider segments, it is useful to integrate measures of price sensitivity. Here, the price elasticity for transit fares is used as the measure of price sensitivity. Price elasticity is calculated as the percentage change in demand for a given percentage change in price. In regards transit fares, the fare elasticity is the percentage change in transit journeys given a change in the cost of each trip. The more sensitive a rider is to price changes, the more responsive their travel behavior will be, and the higher the corresponding fare elasticity.

Given that the alternate fare structures were developed to be revenue neutral in aggregate while holding ridership constant, the cost of individual rail trips correspondingly fluctuates with some becoming more expensive and others less so. To capture rider responses to fare changes, elasticity estimates were applied to each trip based on journey and rider characteristics. The source of the elasticity estimates is a study conducted specifically for the Metrorail system by The National Center for Smart Growth at the University of Maryland [61, 62] to develop granular elasticity estimates that are differentiated by journey and rider characteristics. Given the trip and demographic information collated in the trips dataset, the most elastic option for a given trip based on the rider's race, income, access to transit benefits, and the number of bus-rail transfers on a given journey was selected.

After assigning fare elasticities, the detailed characteristics of each trip in the dataset were used to calculate the fares due under each fare structure. Although the price of trips between origin and destination pairs change in divergent ways, not all pairs are equally traveled. To determine system-wide travel volume changes, elasticities were applied to calculate trip changes for each station pair and aggregated across all station pairs and travel periods. Figure 4-5 shows the directional change in price for each trip in the weighted trips dataset. The x-axis groups trips into crow-flies distance-based buckets, the y-axis indicates how many Metrorail trips of that distance occurred, and the coloring indicates how prices would change under each fare structure. All fare structures except scenario 1a, the current fare structure, are displayed. The current fare structure is the baseline, and there would be no change in cost for any trip when compared to itself.

As Figure 4-5 shows, the number of trips affected varies widely based on what aspects of travel are used to determine price. Further, the selection of a fare structure framework can cause the changes to be spread across many trip distances (scenario 1c), or result in benefits and costs accruing more narrowly to trips based solely on the distance (scenarios 3) or time of travel (scenario 1b).

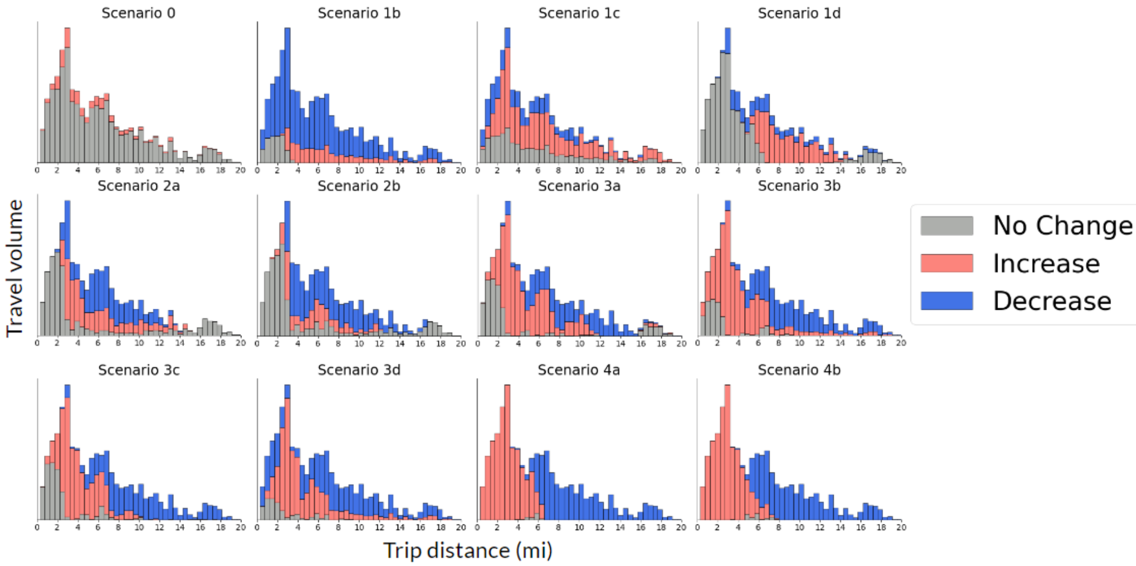

Figure 4-5: Fare Changes by Trip Distance under Alternate Structures

4.4 Findings

4.4.1 Travel Behaviors by Rider Segment

Table 4.3 displays summary statistics for trips grouped by both race and income. While the rider groups of interest to this case study are defined by race and income separately, this table provides useful background when exploring rider segments. The average travel distance (in crow-flies miles) and the ratio of peak travel both influence fares under the current system. So does the directness of each trip, included in the table as travel circuity. This metric captures the indirectness of trips taken on Metrorail are, defined as the ratio of Metrorail track distance to crow-flies distance. A value of 1 indicates that the distance along the rail network is a direct line matching the crow-flies distance; as the circuity measure increases, so too does the detour required to navigate from origin to destination using Metrorail.

As shown in Table 4.3, slightly under half of the trips in the weighted dataset were taken by riders of color. When considering income, low-income riders take just under 14% of Metrorail trips. Most low-income riders are people of color but all low-income riders, regardless of race, are more likely to travel during off-peak periods.

Considering average trip distance and circuity, low-income riders take both the shortest trips as well as the least direct trips. Delving further, low-income riders of color take by far the most circuitous trips in the dataset. Riders of color with incomes above the livable wage take both the longest trips, and much more direct trips as compared to their low-income counterparts. Among white riders, those making above the livable wage take both longer and more direct trips, but the discrepancy is not as large.

		People of Color		**White**	
Weighted Metrics	All riders	Below Livable Wage	Above Livable Wage	Below Livable Wage	Above Livable Wage
Trips	12,129	1,218	4,544	456	5,911
Pct. of total		10.0%	37.5%	3.8%	48.7%
Average distance *(Crow-flies miles)*	6.05	5.56	6.52	5.51	5.83
Travel circuity *(Track miles / crow-flies miles)*	1.29	1.46	1.30	1.33	1.27
Trips by period					
Peak	70.4%	54.6%	74.5%	63.4%	71.1%
Off-peak	29.6%	45.4%	25.5%	36.6%	28.9%
Average fare ($)	3.28	3.08	3.44	3.12	3.20

Table 4.3: Combined Trips Dataset Summary Statistics

4.4.2 Fare Equity Findings Across Race

This section walks through the three measures of equity where riders are segmented based on race. Each chart shown also contains scenario 1a, the current Metrorail structure, as the baseline for comparisons. Figure 4-6 shows the equity benefits

WMATA realized by removing the transfer fee in September 2021. The average fare per journey went down, but it went down much more for riders of color as a higher proportion of their trips involve a transfer between rail and bus. The decrease in average fares from removing the transfer fee almost halved the fare differential between riders of color and white riders. It generated increased travel by all rider groups, with more total trips and a higher proportional increase in Metrorail use for riders of color. However, the results still leave riders of color paying roughly 5% more for each trip under the current structure when compared with white riders.

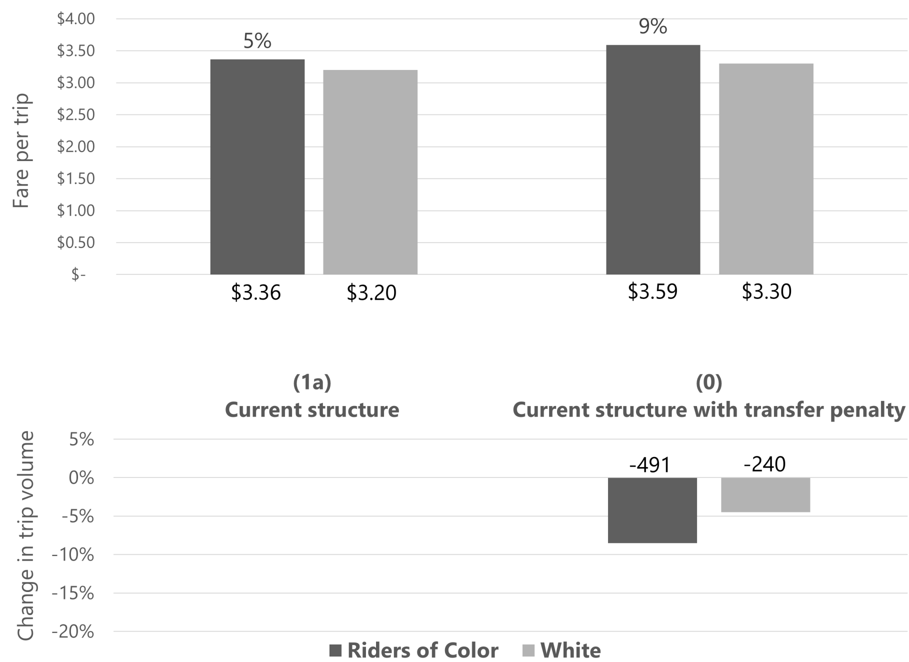

Figure 4-6: Scenarios 0 by Race

Fare structures maintaining the current composite distance basis result in varied average fares, but the prices across segments fluctuate together leaving the price differential unchanged. The only of these scenarios expected to result in a noticeable change in trips is Scenario 1b, which removes peak pricing. The majority of trips, regardless of race and income, occur in the peak period; accordingly, without peak pricing the majority of trips would become cheaper. However, it is important to remember that the converse also occurs; trips in the off-peak, which low-income travelers take at higher frequencies, are saddled with higher fares.

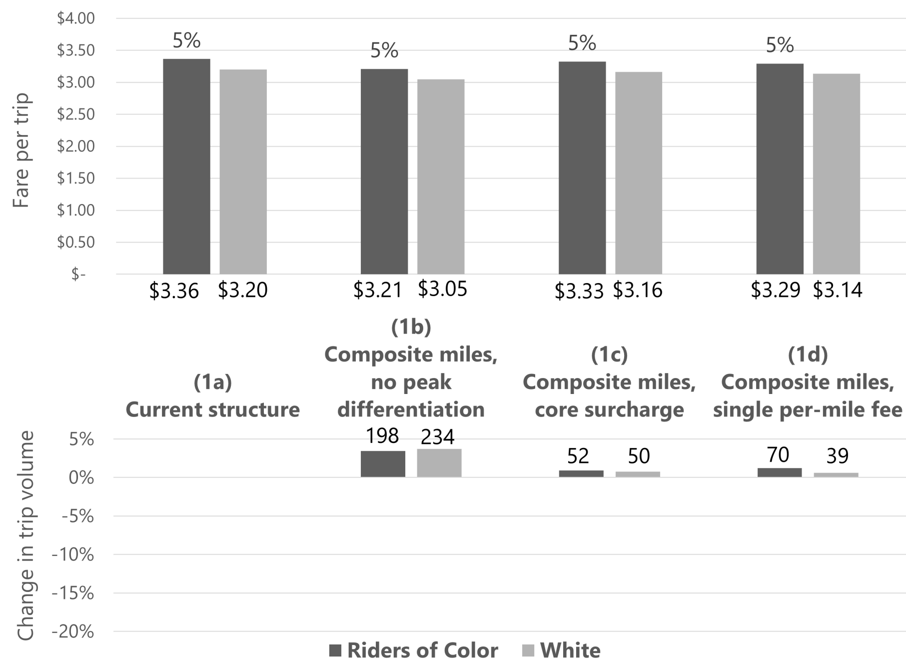

Figure 4-7: Scenarios 1 by Race

Altering the distance basis used to calculate fares results in minimal changes to fare levels and trips taken, but a switch to crow-flies distance as the fare basis provides a slight positive benefit to riders of color. This is intuitive, as these riders have the highest travel circuity, which is penalized by the current composite distance fare basis. Looking at track miles, which comprises the other portion of the composite distance, the results indicate little change from the current fare structure.

Zonal results differ across the chosen implementation used, with zonal difference bases (3a, 3b) equalizing average fares across race. On the other hand, zonal boundary crossings structures (3c, 3d) have no impact on fare differential across groups. Overall, the impact of all zonal structures is to depress the number of rail trips, with white riders expected to decrease travel proportionally more in three of the four implementations.

Flat fares equalize the fare per trip, but lead to an expected drop in trips for all rider groups. This is intuitive; to reach a revenue neutral fare, assuming constant

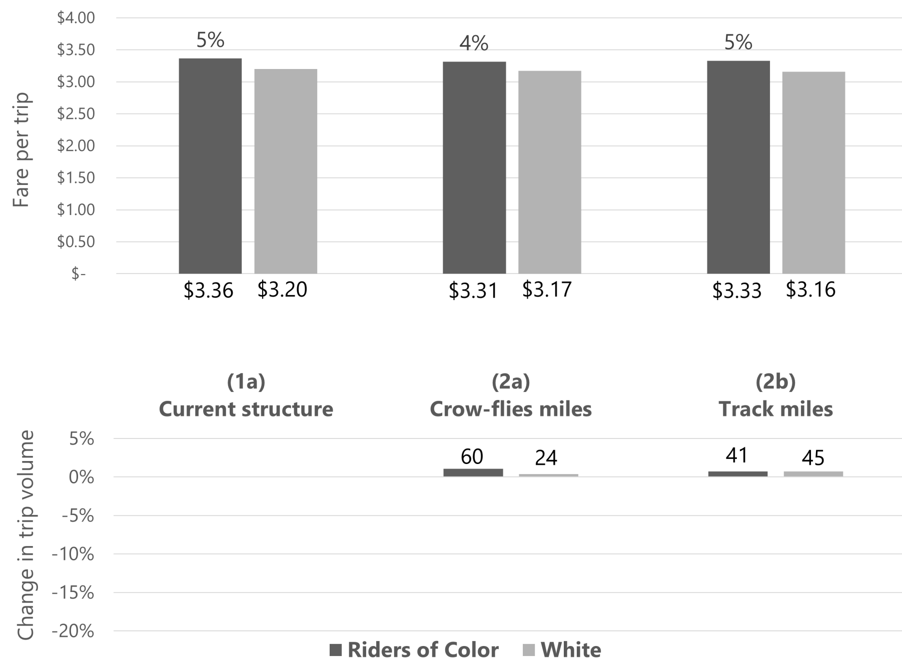

Figure 4-8: Scenarios 2 by Race

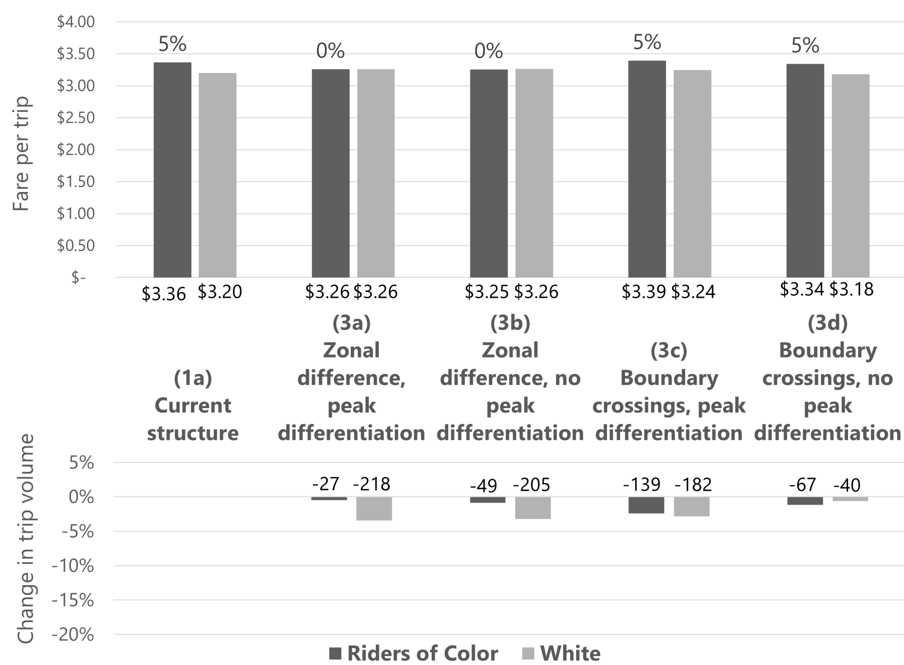

Figure 4-9: Scenarios 3 by Race

ridership, the calculated flat fare must be higher than the current fare for many short-distance trips, whereas longer trips see a price reduction. Given the distribution of trip lengths, there are many short trips negatively impacted by this change with long trips benefitting from subsidized travel. When considering riders based on race, this results in more lost trips by white riders as they take trips that are shorter on average, meaning a higher proportion of trips by white riders would become more expensive.

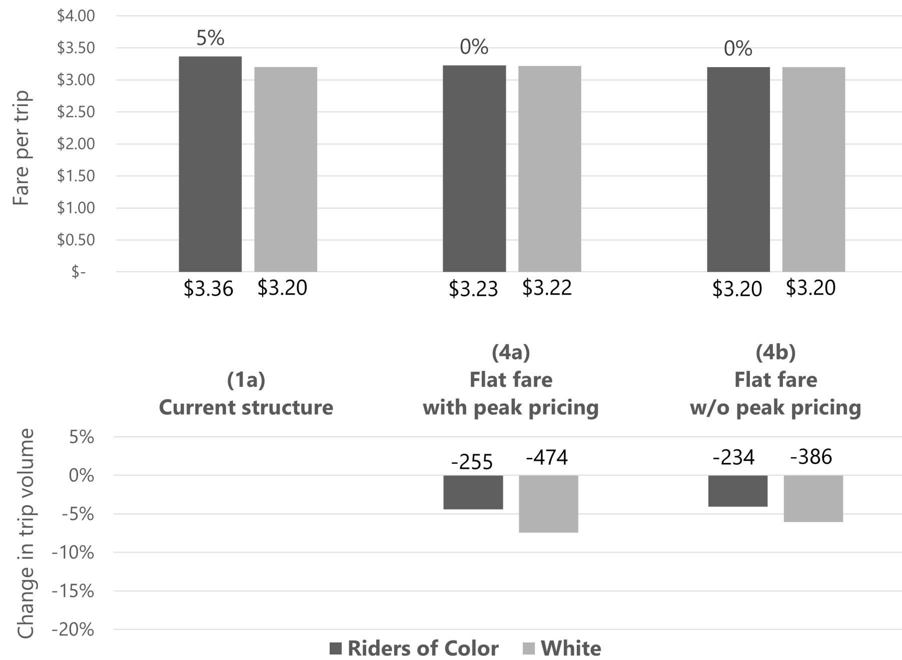

Figure 4-10: Scenarios 4 by Race

4.4.3 Fare Equity Findings Across Income

This section walks through the three measures of equity where riders are segmented based on income. As in the previous section, each chart shown contains scenario 1a, the current Metrorail structure, as the baseline for comparisons. Figure 4-11 shows the equity benefits WMATA realized by removing the transfer fee in September 2021. The average fare per journey went down, and it went down much more for low-income riders who make many more transfers between the rail and bus systems. It generated increased travel by all rider groups, with more total trips and a higher proportional increase in Metrorail use for low-income riders. This results in the current system where low-income riders pay 6% less per trip for each Metrorail journey than their peers making more than the livable wage. It also unlocked many more trips within the WMATA network, as indicated by the expected increase in trips when moving from scenario 0 to scenario 1a.

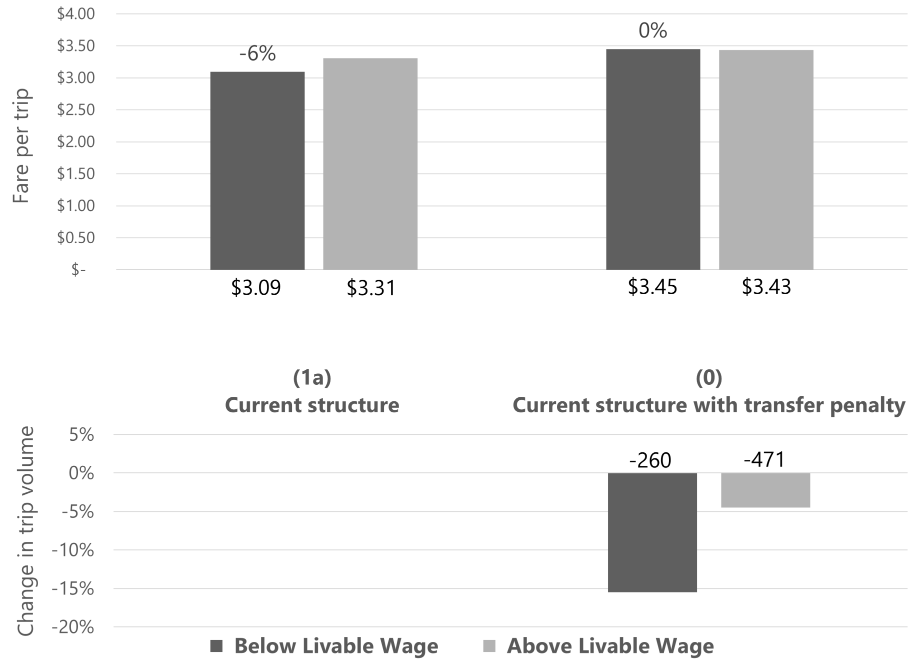

Figure 4-11: Scenarios 0 by Income

Variations on the current composite-distance fare basis show less consistency across structures than when race was used to group riders. First, removal of peak pricing (1b) would reduce prices for all travelers, but is most beneficial to higher-income riders who travel more in the peak and accordingly experience a larger drop in average fare. On the other hand, structure 1d contemplates a single, consolidated per-mile rate coupled with a decreased off-peak per-mile charge. For low-income riders, this is estimated to reduce average fare and boost low-income trips with negligible impacts to higher-income travelers. Finally, a core surcharge (1c) provides slight benefits to low-income riders.

When altering the distance basis used to determine fare, a crow-flies basis (2a) benefits low-income riders, as they take the most indirect trips. This is true for both white riders and riders of color, as low-income riders of each group have more circuitous travel than their higher-income counterparts. Although the results of a

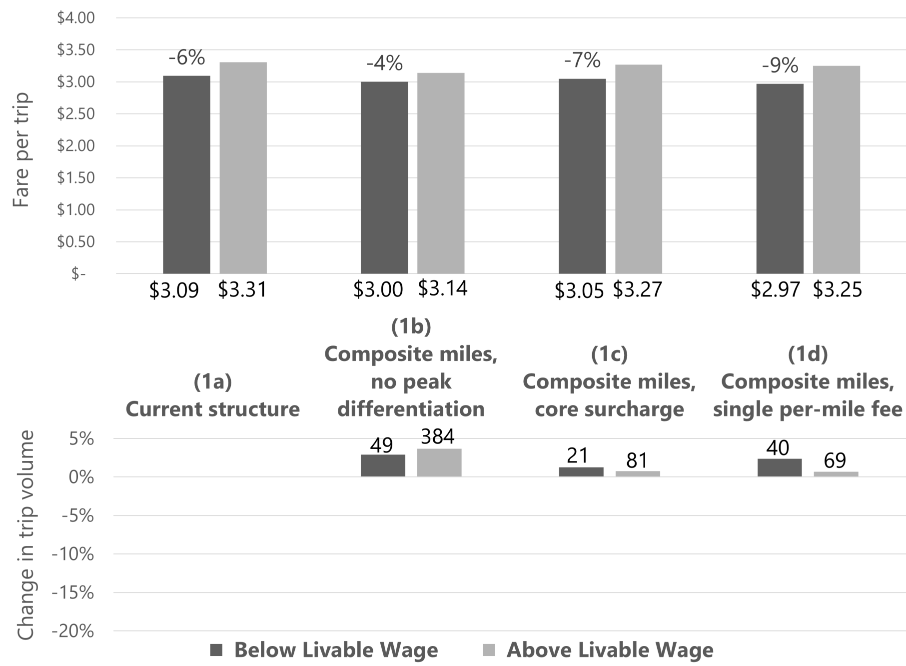

Figure 4-12: Scenarios 1 by Income

switch to a crow-flies miles basis is a net benefit to riders across incomes, the largest benefits accrue to riders below the livable wage. The average price per journey drops, resulting in a larger fare differential compared to higher-income riders, and rail travel volumes increase. In comparison, track miles as a basis (2b) has almost no impact on the fare differential between groups, but it spurs additional travel almost solely for higher-income travelers who take more direct trips on the Metrorail system.

The equity impact of a zonal fare system, when considering rider income, varies based on its implementation. A zonal difference system with peak pricing (3a) provides a slight benefit to low-income riders, but at the expense of a systemwide decrease in trips. Every other zonal implementation performs worse in terms of equity. Removing peak pricing from the zonal difference structure (3b) increases the fare differential between groups, but this is achieved through reduced travel for both segments. Boundary-crossings implementations (3c, 3d) have much larger negative impacts on low-income riders, who would experience higher costs and fewer trips given their indirect travel patterns, which are penalized more heavily when fares are based on the

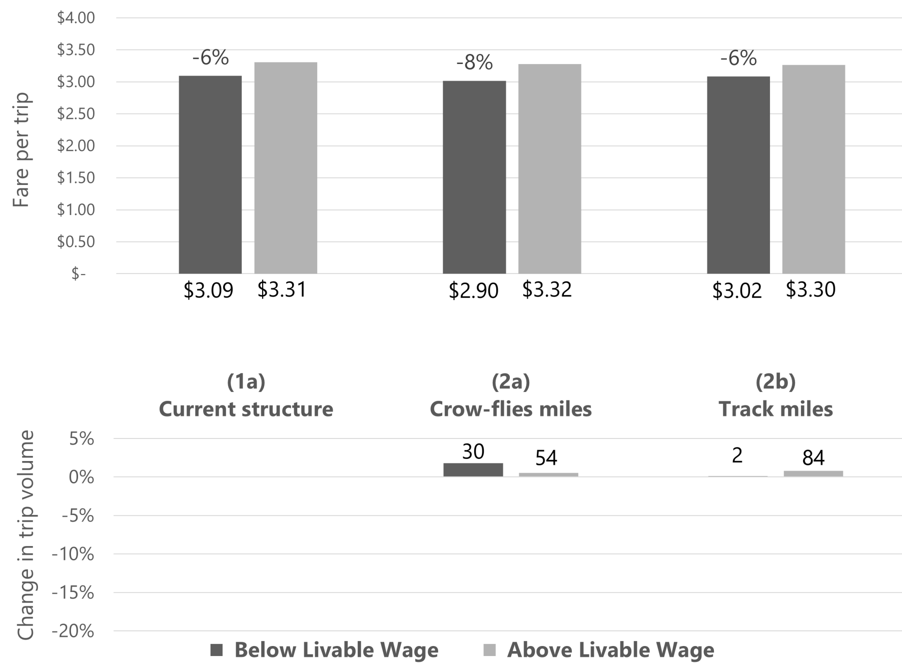

Figure 4-13: Scenarios 2 by Income

number of zones crossed.

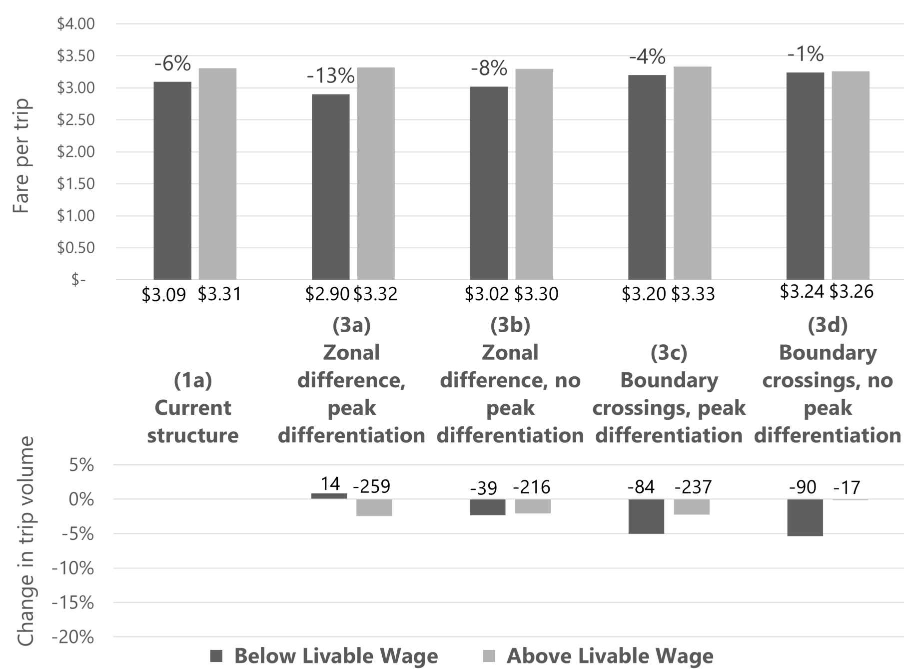

Figure 4-14: Scenarios 3 by Income

Flat fares (4a, 4b) result in the largest expected decrease in travel. As when riders were segmented by race, all groups reduce Metrorail travel. It is low-income riders would pay higher prices and reduce travel to a larger extent than higher-income travelers. A fully flat fare without peak pricing (4b) results in a proportional drop in trip volumes for low-income riders that is over twice as large as observed for riders making above the livable wage.

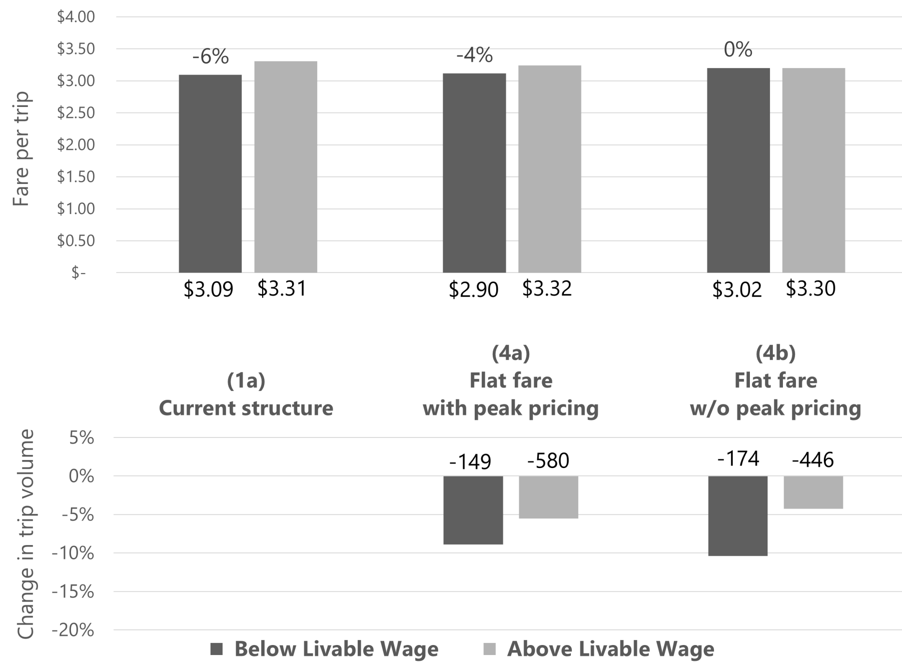

Figure 4-15: Scenarios 4 by Income

4.4.4 Discussion

An interesting result from this analysis concerns the differing fare scenario impacts on rider segments when defined by race and income. These two groups are equally important when considering equity, but differing travel behaviors can result in contrasting recommendations. One clear example can be observed in the disparate impacts of flat fares (scenarios 4); within the Washington, D.C. area this would result in riders of color no longer paying higher per-trip fares, but low-income riders would face the burden of the largest fare increases. Further, a revenue neutral flat fare is expected to decrease travel for all groups, but it is expected to reduce trips taken by white riders to a greater extent than riders of color, and separately reduce trips by low-income riders more than riders earning above the livable wage. While flat fares may by some measures benefit riders of color, they would simultaneously disadvantage low-income riders.

Other investigated fare structures confirm that equity impacts do not always vary in concert across race and income. Consider the four potential zonal fare structures explored: estimates show that when segmenting riders by race there is, at worst, no change in the fare differential and, at best, an equalization of average fares between riders of color and white riders. When evaluating the same zonal implementations based on rider income, however, the impacts range from a reduction in average low-income fares and a decrease in higher income travel (3a) all the way to fare increases and travel decreases almost solely for low-income riders (3d).

Although there are differing impacts on fare equity when considering race and income, there are nevertheless attributes that show similar impacts across fare scenarios. A clear result from this analysis is that peak pricing provides a net benefit to low-income riders without disadvantaging riders of color. When comparing white riders and riders of color, fare implementations with and without peak pricing result in the same fare differential between groups regardless of fare basis. This holds across all scenarios evaluated. On the other hand, across the very same structure types any removal of peak pricing negatively impacts low-income travelers.

Penalizing indirect travel further reduces fare equity. This is most plainly evaluated by altering the distance measure used as a basis for fares (scenarios 1a, 2a, 2b). Whether considering race or income, utilizing crow-flies distance to calculate fare results in more advantageous outcomes for both rider segments. Just as importantly, travel data shows that riders that are both low-income and people of color are substantially more impacted by these penalties than riders belong to any other race and income combination. Switching to a crow-flies basis would provide major benefits to riders belonging to the intersection of the two equity groups considered. Riders belonging to both equity groups are the most disadvantaged by the current measure of distance. Regarding the other considered distance basis, using track miles along the WMATA network provides a slight advantage to white riders and those earning above the livable wage. Switching the current structure to a crow-flies distance basis would not only better align travel cost with the benefits received, but advantages both low-income riders and riders of color.

Similarly, when comparing zonal fare structures, implementing a zonal fare structure based on the difference between origin and destination zones provides the most benefit to low-income and riders of color. Basing the zonal fare on the number of zonal boundaries crossed, even with a lowered incremental cost for each additional crossing after the first four, results in worse outcomes for low-income riders and riders of color. Riders pay to travel from origin to destination, as quickly and safely as possible; charging them for taking circuitous trips penalizes them monetarily for something beyond their control, after already requiring more time to travel an indirect route.

The results highlight an additional opportunity to reduce fare complexity, increase fare legibility, and simultaneously provide a benefit to low-income Metrorail users; consolidation of per-mile rates. Charging a single per-mile rate for all distances between the minimum and maximum fare, and combining this change with a single, lowered off-peak per-mile rate subsidized by increased peak revenues would provide a benefit to low-income riders with minimal impact on peak travelers.

The policy tradeoffs considered in this case study illuminate which elements of travel, when incorporated into fare policy, can enhance equity in the Washington, DC metro area. The results report weighted average fares and aggregated trip volumes under each fare scenario after identifying traveler demographics using a methodology that can be applied to additional transit networks. However, the exact elements of fare structures that improve or decrease fare equity are context-dependent and sensitive to such factors as the specific demographic residential choice, development patterns, gentrification, and travel patterns.

Chapter 5

Fare Product Design and Pricing

5.1 Introduction

Travel of all stripes saw a dramatic reorganization over the course of the COVID-19 pandemic; mass transit systems felt these changes acutely. The pandemic caused major reductions in all travel, some forms of which has rebounded more quickly and some which has not. Public transit falls into the latter category. Commute trips have historically been the backbone of most public transport systems in the US, with a 2017 report from the American Public Transportation Association finding that 50% of transit users regularly traveled on their system 5 days a week. Across all travel, 49% of transit trips were for a work commute [63]. To summarize, in the years leading up to the pandemic, commuting was the primary purpose for just as much transit travel as all other trip purposes combined. Stay-at-home orders affected all travel, but ongoing remote work and hybrid in-office schedules continue to result in lower volumes of transit use on a given workday. This represents a large chunk of trips for which demand remains missing.

Beyond the impact on aggregate transit ridership levels, changes in commute patterns have scrambled the calculus on how calendar-based pass products should be designed and priced. The last pre-COVID fare changes for the CTA went into effect in January 2018. At that time, single-ride fares for bus and rail were $2.25 and $2.50, respectively, and up to the present those rates have not changed. The prices for all pass options effective in 2018 remained unchanged until promotional discounts were offered on 1-day, 3-day, and 7-day passes during the summer of 2021. These discounts were permanently adopted in November 2021 alongside a decrease in the cost of the 30-day pass [4]. Table 5.1 below shows pass prices before COVID, current pass prices, and the pass multiple for rail in both periods. The pass "multiple" represents the rider's breakeven point between pay-per-use fares and a pass product. It is calculated as the number of pay-per-use trips required to reach the pass price. Figure 5-1 shows the daily cost of the pass offerings.

Fare Product	**Pre-COVID Cost ($)**	**Pass Multiple**	**Current Cost ($)**	**Pass Multiple**
1-day	10	4	5	2
3-day	20	8	15	6
7-day	28	11.2	20	8
7-day (CTA + Pace)	33	13.2	25	10
30-day	105	42	75	30

Table 5.1: CTA Pass Price Changes

These price reductions were a positive attempt by the CTA to respond to the changing utility of their pass products. The fare for a single-ride has not been changed for either the bus or rail systems; this choice by the CTA makes pass products more desirable when compared to pya-per-use fares. In particular, the steep per-day price drops on shorter passes sought to induce additional ridership with minimal fare revenue losses. This is most apparent when viewing the 1-day pass, now priced at 2 rail trips, or a single round-trip.

The CTA showed that they were open to reconsidering the price of their fare products; the next step is to reevaluate the fare products offered. The case study

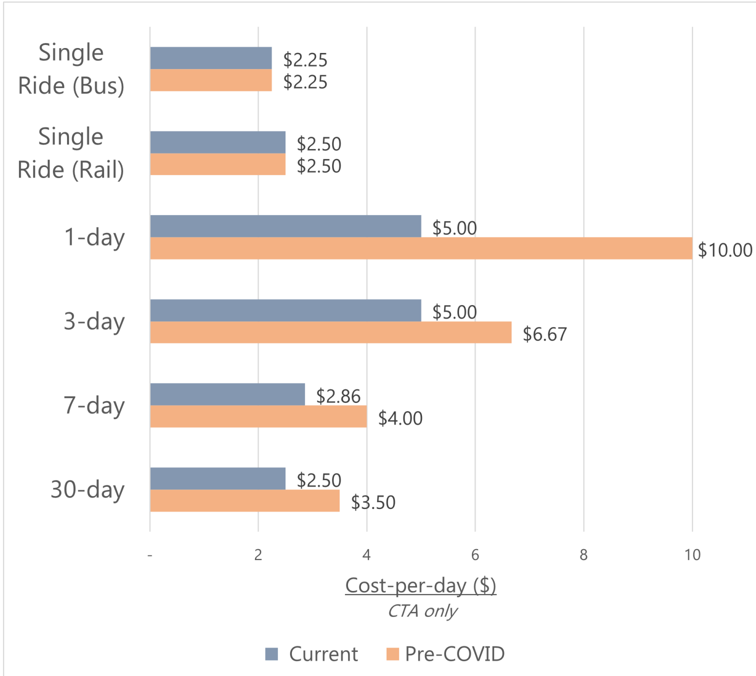

Figure 5-1: CTA Fare Product Costs per Day

presented in this chapter seeks to do just that. A method is presented to identify the benefits of fare products to CTA ridership, and that process is used to understand how relative fare product utility has changed over time given shifting travel. New, nontraditional fare products are then designed and introduced into the process, resulting in a fare product choice for each rider in multiple hypothetical scenarios. This illustrates which fare products would see the most switching, in turn allowing approximations of revenue impacts to the agency. For the rider, the ability to travel more should result in additional trips, something transit agencies are eager to generate. This methodological process allows agencies to experiment with hypothetical fare products and identify promising options for further study or potential pilot programs.

5.2 Research Goals

The case study presented in this chapter uses travel records from the CTA in Chicago, IL to explore fare product designs that encourage transit use and reflect post-COVID travel patterns. The overall idea is explored through two distinct but intertwined analytical steps used to evaluate two aspects of fare product design.

- **First**, to what extent did simply lowering the price of all CTA transit passes in Chicago counteract any drops in their utility to transit users? Did the price drops, in conjunction with new travel patterns, change the relative attractiveness of each pass offering?
- **Second**, what aspects of pass product design can be altered to match current travel patterns. Passes normally confer an unlimited amount of travel for a set number of consecutive calendar days. How do altering these two dimensions (unlimited travel; consecutive calendar days) increase or decrease the attractiveness of a product? Are the findings consistent across potential price points?

Considering the first research question, it is intuitive to think that dropping pass prices would boost their attractiveness and persuade some users to reconsider their fare product choice. When looking at the pass multiples in Table 5.1 it is apparent that the 30-day pass, which previously required 42 rail rides to reach cost parity with pay-per-use ('PPU') travel, would not be attractive to very many riders post-COVID given the large drop in commute trips. For a pre-pandemic worker reporting to the office five days a week, the 30-day pass price would have made some economic sense. For that same employee working a hybrid schedule post-COVID, and working on-site in the office two days a week, that same 30-day pass would no longer make any sense. The equivalent post-COVID pass multiple would need to be on the order of 16 trips, not 42. Further, the riders still traveling frequently enough post-COVID to make the 30-day pass a smart purchase can plausibly be expected to overlap with the riders least able to pay the large upfront costs of that pass. Clearly, the pass price reductions made passes more attractive; this case study illustrates the number of riders that stood to benefit from the reduced prices.

Each pass product became more attractive to riders by the sheer fact that they became less expensive, and it takes fewer pay-per-use trips for the break-even point to be reached for all options. However, the price reductions for each pass product were not uniform. The largest price drop occurred for the 1-day pass, which was a 50% decrease; the smallest price drop occurred for the joint CTA/Pace 7-day pass which saw a 24% decrease. Given price reductions of differing amounts, new pass multiples, and less consistent travel, the relative utility of each product is likely to have changed. In this chapter, these changes in the relative utility across fare products are appraised. Agencies need to be able to understand if their pricing decisions might result in their desired outcome. In the case of the 2021 pass price reductions, the CTA's goal was to draw riders away from pay-per-use travel and onto passes. Evaluating rider travel patterns and varying fare products can together shed light on the efficacy of such actions.

That is where the second research question investigated by this case study comes squarely into play. How can existing pass products be modified to better match post-pandemic travel? The methodology developed in this chapter is able to analyze relative changes in the attractiveness of fare products as price points change. But that methodology is not restricted to considering fare products already in existence; it can just as easily be used to investigate hypothetical offerings.

To appraise novel fare products (also referred to in this thesis as "flex-passes") alongside existing pass and single-ride products, the parameters defining the flex-pass must be established. Parameters could be made to include time of day restrictions (e.g., "only valid on weekends"), rider demographic requirements (e.g., student or university passes), or various other constraints on use. The investigation here is focused on all riders traveling on the CTA post-COVID, so the changes for the hypothetical fare products will occur primarily across three characteristics:

- **The number of trips conferred.** Current pass offerings at the CTA confer an unlimited number of trips for an individual rider for the duration of their validity.

- **The period of validity.** Traditional fare products are valid for a set amount of continuous time. This can be measured in hours, as with single-ride tickets which can be used to transfer within the system for up two hours after activation, or in days, such as the 30-day pass valid for 30 consecutive calendar days.

- **The price.** This can also be reframed as the discount on each trip in exchange for an up-front payment, which in turn assumes a commitment to ride in a customer's mind.

This case study develops and experiments with flex-pass designs that do not require usage on consecutive calendar days, or that confer a fixed number of journeys on the CTA system instead of unlimited travel on a fixed number of days. These flex-passes are then analyzed individually alongside existing fare products to gauge how many riders could benefit, and how riders should shift from the existing products to the considered flex-pass. Repeating the process with different flex-passes at varied price points can be used to test robustness and price points before any products are actually introduced.

5.3 Methodology

This section develops the methodology used to investigate the relative utility of fare product offerings from a given transit operator. The methodology set forth in this chapter uses actual transit travel records on the CTA network gathered from agency AFC data. In order to estimate the relative attractiveness of a particular product to the entire population of system users, it is insightful to understand how many riders should choose that product from the entire set of options. The use of AFC-based data allows determination of the preferred fare product for each rider in each time period, demonstrating the extent to which travel behavior shifts change how useful each fare product is to ridership at large. The fare products that are gaining and losing share due to these changes can then be seen. Similarly, evaluating different groupings of fare products using AFC data from a single time period can isolate the impacts of price points and usage restrictions on the relative utility of all fare products.

5.3.1 Parameters for Inclusion in Analysis

To begin this analysis, the first step was to assemble the data for analysis. AFC records are available at the level of individual riders for all account-based Ventra travel. The large drop in transit ridership over the course of the pandemic provided an obvious opportunity for product reevaluation. However, not all of the riders using the CTA regularly before the onset of the COVID-19 pandemic have returned to regular CTA use. The subset of riders active on the CTA leading up to COVID, and also active when society began reopening, provide the foundation of this analysis. These individuals were still traveling, and their travel patterns changed between the two periods. The best corresponding fare products for these riders can be quantified using AFC records.

Travel pattern changes over time were observed and evaluated using transit use records from October 2019 and October 2021. October is in general well-suited for transit data benchmarking, as it has no major holidays and is not peak season for vacationing. The weather is neither freezing cold, as in winter, nor particularly hot, as in summer. October can normally be considered a generic, average month and is therefore useful for comparisons across time. October 2019 was selected as a representative month shortly before the onset of COVID-19; October 2021 was selected as a representative post-COVID travel month. By October 2021 vaccines had been widely available, infections levels had receded, and COVID-caused impacts on work travel had begun to solidify. Additionally, the CTA's pass price reductions were largely in effect by October 2021, allowing exploration of impacts from lowered prices on utility to riders.

After selecting two time periods for comparison, one prior to COVID and the other afterwards, the next step is to establish thresholds for inclusion of individual riders. One clear requirement for inclusion is the existence of identifiable travel on the CTA in both time periods. There may be many reasons that a Ventra account has CTA travel records for only one of the months. Explanations could include receiving a new Ventra card or opening a new account, a cessation in CTA use, or a switch between limited-use and account-based fare products. Regardless of the cause, riders only active in one of the months are excluded from the analysis. To measure shifts in fare product choice across riders, a direct comparison dictates the same riders are observed in each period.

Filtering all Ventra accounts down to those active on the CTA in both periods is only the first step. The focus of this analysis are the fare products offered by the CTA, and the changing utility of those products; for that reason, very infrequent riders can be excluded. Riders that travel very infrequently have one clear fare product that they should choose: pay-per-use. The analysis conducted here would not result in insightful outcomes for these riders. For that reason, all riders included in this case study traveled on the CTA a minimum of three calendar days in each of the two months.

After creating a dataset of CTA travel by riders meeting the inclusion criteria, it can be clearly observed that the frequency and quantity of CTA trips by these riders changed between October 2019 and October 2021. Figure 5-2 shows that on the days these riders used the CTA network, they took fewer trips. As shown, there is an obvious jump in the number of riders taking a single journey on the CTA post-COVID, with corresponding drops in two-, three-, and four-trip travel days. Similarly,

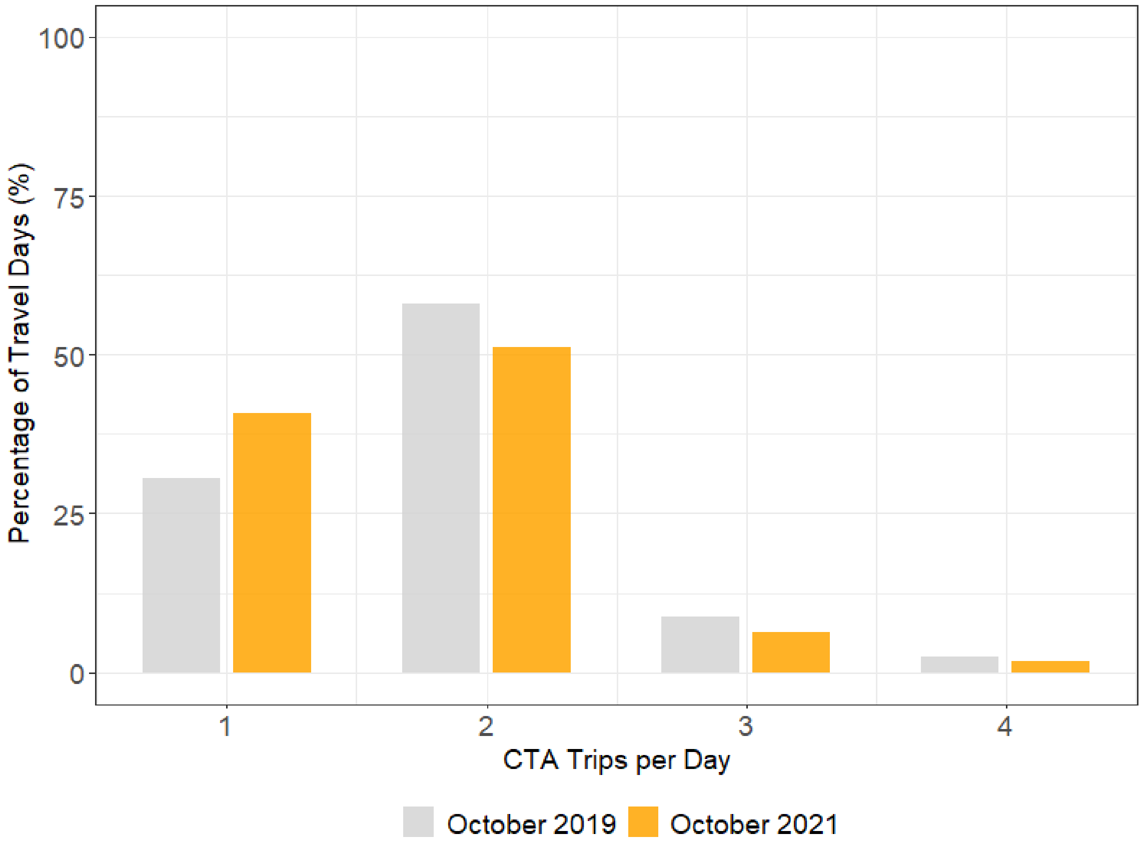

Figure 5-2: Trips per Day by CTA Travelers

Figure 5-3 shows that these same riders used the CTA on fewer days per month in 2021. In 2019, the distribution of CTA travel days for these riders was bi-modal, with many riders traveling less than five days in the month but even more riders traveling 20-25 days out of the month. By 2021, this distribution had completely shifted with an inverse relationship between the number of travel days per month and the number of Ventra accounts. The clear majority of riders were traveling no more than ten days per month on the CTA.

5.3.2 Changing Fare Product Choices

During the two-year period between October 2019 and October 2021 changes clearly occurred in the transit travel patterns of CTA users. On the whole, travelers used

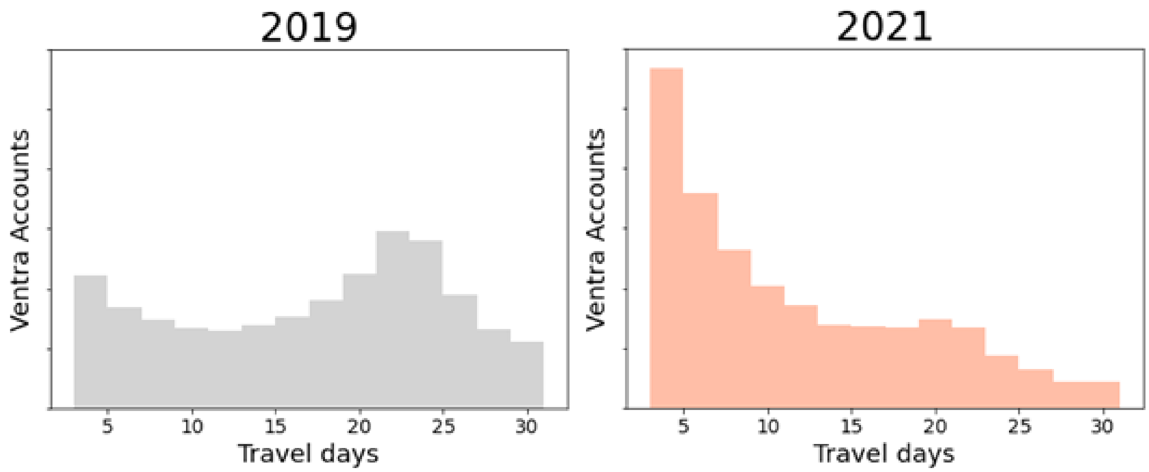

Figure 5-3: Changes in CTA Travel Frequency

the system less often and when they did use transit, they took fewer trips. The exact same fare products were available during both time periods.

Assuming that there are riders who seek to optimize their travel costs by spending the minimum possible amount on transit travel while still being able to complete all of their trips, it stands to reason that there should be travelers that buy multiple fare products to match their travel. The number of riders selectively determining a fare product purchase, observed where mixing and matching occur, could be expected to have increased over the course of the pandemic, as travel frequency became more sporadic and future travel expectations became murkier. This assumption is not borne out by AFC ridership records, however, as shown in Figure 5-4. The percentage of riders that regularly traveled on more than one fare product type remained largely stable from pre-COVID through the early stages of the COVID-19 pandemic. The fare promotions introduced by the CTA for the summer of 2021 corresponded to a slight increase in the number of users that mixed fare products, and when those reductions were made permanent alongside price drops on additional fare products there was another small jump in the volume of riders choosing to mix fare products.

The number of Ventra accounts using multiple fare products increased concurrently with fare promotions introduced by the CTA for the summer of 2021, and this

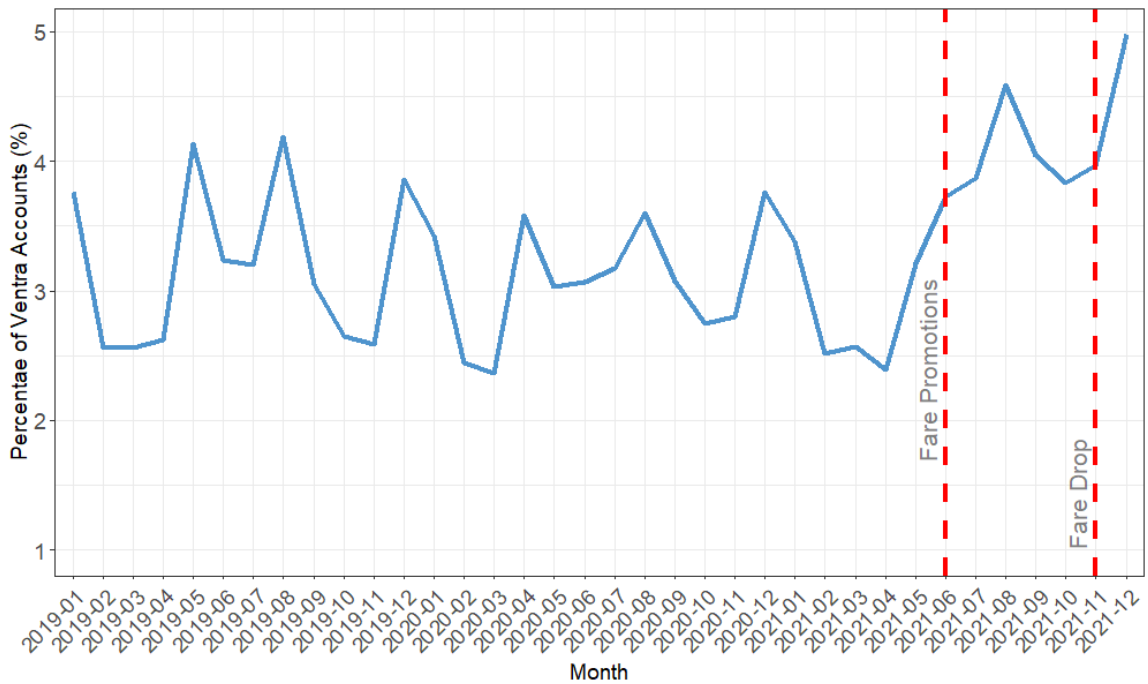

Figure 5-4: Ventra Accounts using Multiple Fare Products

remained the case when those pricing promotions were made permanent alongside additional pass price reductions in the fall of that year. While the number of Ventra accounts that utilize multiple fare offerings is small, the percentage of transit trips taken by those riders is higher. As shown in Figure 5-5, in December 2021 8% of CTA trips were taken by a multi-product rider, although these riders made up around 5% of all Ventra accounts. This suggests that these users were higher-frequency transit riders taking the most advantage of the fare changes. Viewed differently, more than 9 out of 10 trips on the CTA were made by users loyal to a single fare product. For most riders, that single fare product of choice was pay-per-use single-ride tickets. This is unfortunate given the presumed goal of shifting riders away from pay-per-use to pass products through fare restructuring.

The Venn Diagrams in Figure 5-6 are used to illustrate the dominance of pay-per-use single-ride tickets as the fare product of choice both before and after the CTA's pass price changes. These charts use three illustrative fare products: pay-per-

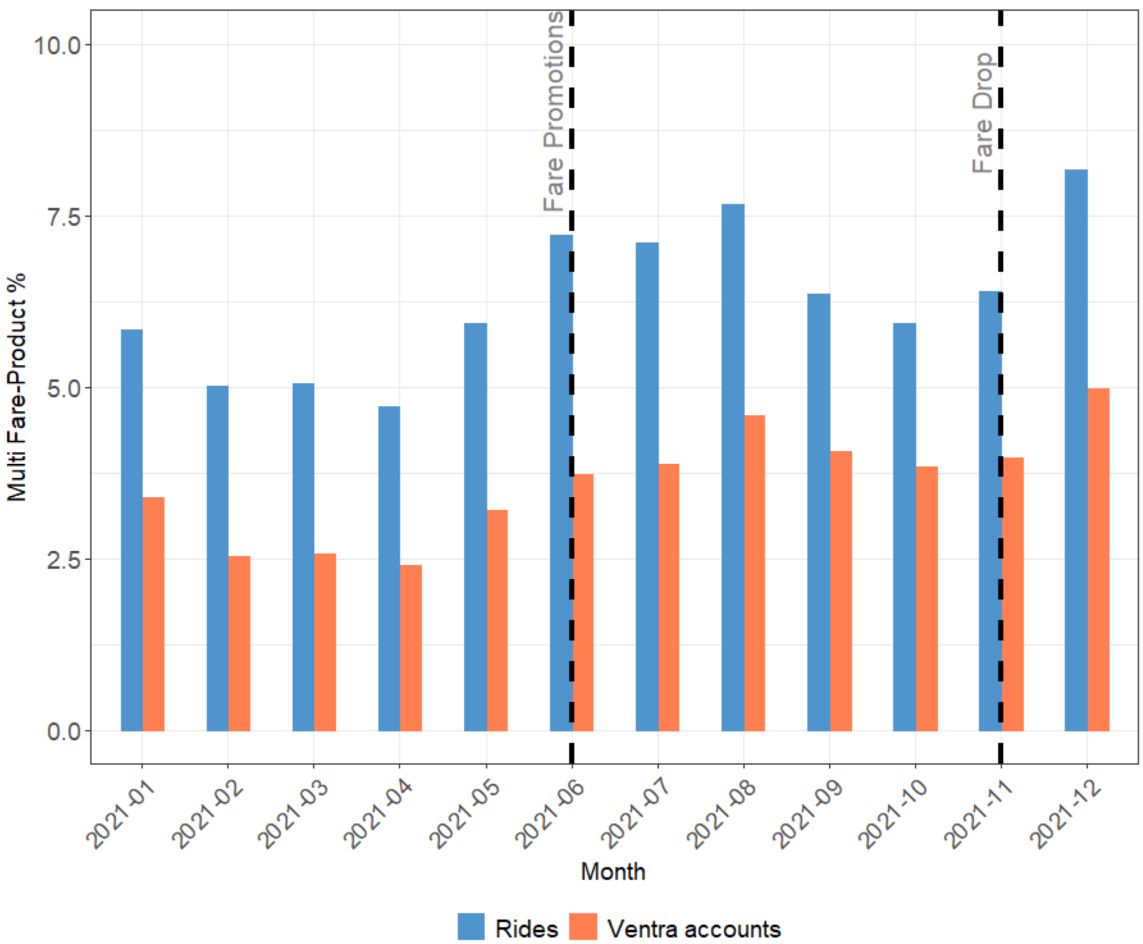

Figure 5-5: Multi-Product Accounts and Multi-Product Trips

use, the 1-day pass, and the 7-day pass. April 2021 saw some riders that used all of these passes in combination the month before the promotional pricing, although that segment of riders was very small. It was also comprised mainly of pass users who occasionally rode pay-per-use. In comparison, there is a much larger cohort of riders that only rode on pay-per-use fares. The price reductions on pass products led to an increase in the proportion of riders that traveled on the 1-day and 7-day passes, as well as more riders using multiple products, as shown for October 2021. But once again, the dominant fare product choice was pay-per-use. These results suggest that the steps taken by the CTA to make pass products more attractive were successful. The promotional pricing introduced in spring 2021 was marketed as part of a post-COIVD campaign to invite riders back to transit [64]. The lower price of the 1-day, 3-day, and 7-day passes was a key component of this campaign, and the CTA advertised it as “More Fun, Less Fare”. The higher usage of short-duration passes suggests the CTA was moderately successful in marketing the price changes, and a number of riders changed fare payment selections in response to these lowered prices. Riders traveling more than twice per day should choose one-day passes as opposed to pay-per-use, indicating there should be significantly more mixing of the payment modes after the reductions. Nevertheless, while the data suggests that price discounts spurred more product mixing, there are nevertheless still many riders to persuade in the quest to encourage pass product adoption by pay-per-use travelers. Additional marketing promotion campaigns may be justified.

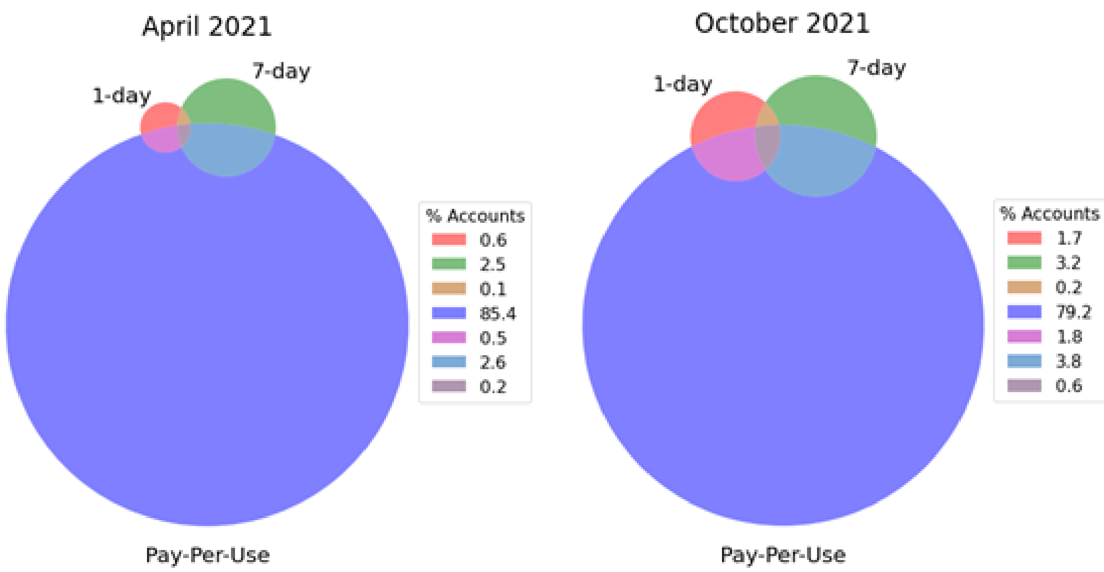

Figure 5-6: Fare Product Mixing

AFC data further shows that the number of riders that had returned to the CTA system by October 2021 varied depending on the primary fare product used before the onset of COVID-19, as shown in Figure 5-7. 75% of Ventra account-based 7-day pass users active in October 2019 had no identifiable transit use two years later. It is the same story for 1-day pass and pay-per-use riders. Compare this to the 50% of

30-day pass users who had not returned to the CTA by the same point. This seems somewhat counterintuitive, as the literature indicates weekly pass users are more likely to be low income and/or transit dependent, both of which are associated with higher levels of continued transit use during the COVID-19 pandemic. One potential explanation for this finding is that some lower-income riders may purchase limited-use passes and paper tickets from ticket vending machines (TVMs) or other physical sales channels that are not associated with a Ventra account. Continued travel on physical media would not appear in the AFC data, which could cause customer churn to appear higher among 7-day pass users. In comparison, 30-day passes are not available through TVMs, and therefore users are more easily associated with Ventra account identifiers. Nevertheless, there is no data that suggests 7-day pass users that purchase tickets from TVMs and users of account-based 7-day passes are meaningfully different in significant ways.

Regardless of pre-pandemic fare product choice, a large proportion of pass users had switched to pay-per-use travel. Of the 7-day pass users that remained on the system, just over half kept using the 7-day pass as their primary product in 2021. Compare this to 30-day pass users who switched to pay-per-use by a ratio of over two to one. For many riders, their fare product choice understandably changed as societal travel patterns shifted. In many cases, this shift was towards pay-per-use travel which requires the smallest upfront payment and provides the most flexibility for riders uncertain about future travel. The 7-day pass users kept choosing their pre-pandemic fare products at the highest rates, possibly because past research show these riders are more likely to be transit dependent. The literature shows in many instances the reason the 7-day pass is chosen is because it is a better deal than pay-per-use for frequent riders, but the cost is more affordable whereas the 30-day pass may be priced out of reach for lower-income and marginalized populations.

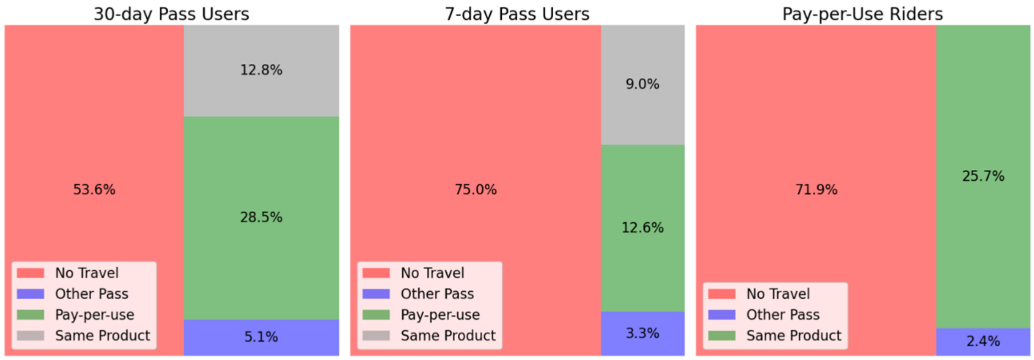

Figure 5-7: Fare Product Choice by Pre-Pandemic Product Group

5.3.3 Creating a Model of Fare Product Choice

AFC records collected from Chicago indicate clear changes in the fare product decisions made by riders over the course of the pandemic. This product selection was further affected by the 2021 pass pricing changes. To isolate the effect of COVID-induced travel pattern changes from the effect of price discounts on the different attractiveness of offered fare products, a process was developed to mimic selection of the combinations of fare products that would accommodate travel taken on the CTA at the lowest cost to the rider. Proper estimation of these bundles of products relied on a set of behavioral assumptions. The primary assumptions are rooted in economic logic. First, CTA travel is beneficial to each rider, so at a single price point the product that confers the most travel is strictly preferred. Second, consumers will want to choose a combination of products that meet travel needs at the lowest possible cost. Third, the transaction costs for all pass products are uniform, and there is no cost to users from using multiple fare products. In practice only some riders actually do mix-and-match from among the available fare products. To understand how offerings match rider behaviors, as opposed to how riders choose a fare product, no aversion to product switching was imposed in the model. The methodology developed here relies on an additional non-economic assumption, which is that riders generally know what their future transit travel will be. Within the context of fare product evaluation as developed in this case study, that is both a necessary assumption and more insightful than alternative constructions.

One additional consideration that is not captured by this methodology is the inherent value of flexibility. One clear benefit of traveling on single-use tickets is that future travel does not need to be known, or even considered. Riders simply pay when they need to ride. While this is contra the CTA goal of inducing ridership by moving users onto pass products, it nevertheless illustrates that riders value flexibility. This aspect is taken into account later when designing and evaluating hypothetical flex-pass offerings. It is important to note that the results will be, if anything, conservative given the rider perception of flexibility is not considered when generating bundles of fare products for each user's travel behavior.

Determining a preferred bundle of fare products is done at the level of individual riders; these individual choices can then be aggregated to understand high-level product utility across the entire CTA ridership. To establish preferred fare products for each Ventra account, actual transit trips taken during the evaluation period were ordered chronologically and evaluated on a daily basis. The totality of the individual's travel is used to determine fare product choices in sequence, as a rider would. Having knowledge of the rider's actual CTA travel, the pay-per-use cost of that travel is calculated, broken down by day. Pay-per-use travel requires the lowest upfront payment; it therefore serves as the reference cost against which other fare products are compared to determine if they are advantageous. Travel records are then evaluated against pass products beginning with the longest duration passes and continuing to other pass products in descending order of duration. The more travel days a pass is valid for, the larger the per-day discount that pass offers. By way of example, when traveling for 30 days, a 30-day pass is cheaper than 30 1-day passes; the 30-day pass is therefore the preferred product for those travel days. When a fare product covers travel at a lower price, it is added to the rider's bundle and all travel that would be covered by that product is then disregarded. The remaining trips are again evaluated against the pay-per-use cost to determine what the next product, if any, in the sequence should be[1]. This continues until all remaining travel should be paid for with single-trip pay-per-use fares.

The process is then repeated with the longest duration pass from the previous iteration excluded. Iterating in this way is necessary to properly account for any gaps in transit use. To give an example, for a rider traveling 15 days on the CTA out of 30, and assuming each of those days was a round-trip, the 30-day pass would be selected as the recommended pass by this looping evaluation method. However, if the majority of those 15 travel days were taken over two 7-day periods with a break in between, a combination of 7-day and 1-day passes could confer the necessary transit access and save the traveler money. The analytical method developed here would create both of these potential rider bundles, and then determine which provides travel at the lowest cost. In case of a tie in cost, the bundle conferring the most potential travel is selected.

This process is repeated for each Ventra account meeting the travel thresholds. A preferred bundle of fare products is output for each of these riders. This method is first applied to pre-COVID and post-COVID travel to illustrate how certain fare products became more or less useful to riders over time. It is then applied to hypothetical pass products to gauge which riders which most benefit from these products, and to what extent.

5.3.4 Designing Novel Flex Pass Products for Evaluation

The procedure described above is applied using a suite of fare products and a set of rider transit travel as inputs; altering either will change the results. This case study

[1]The cost of a 1-day pass for the CTA is $5, and the cost for a 3-day pass is $15. For analytical purposes, the 3-day pass can be viewed as three 1-day passes. There is no added benefit to purchasing a 3-day pass instead of three 1-day passes, so later iterations consider these two passes jointly.

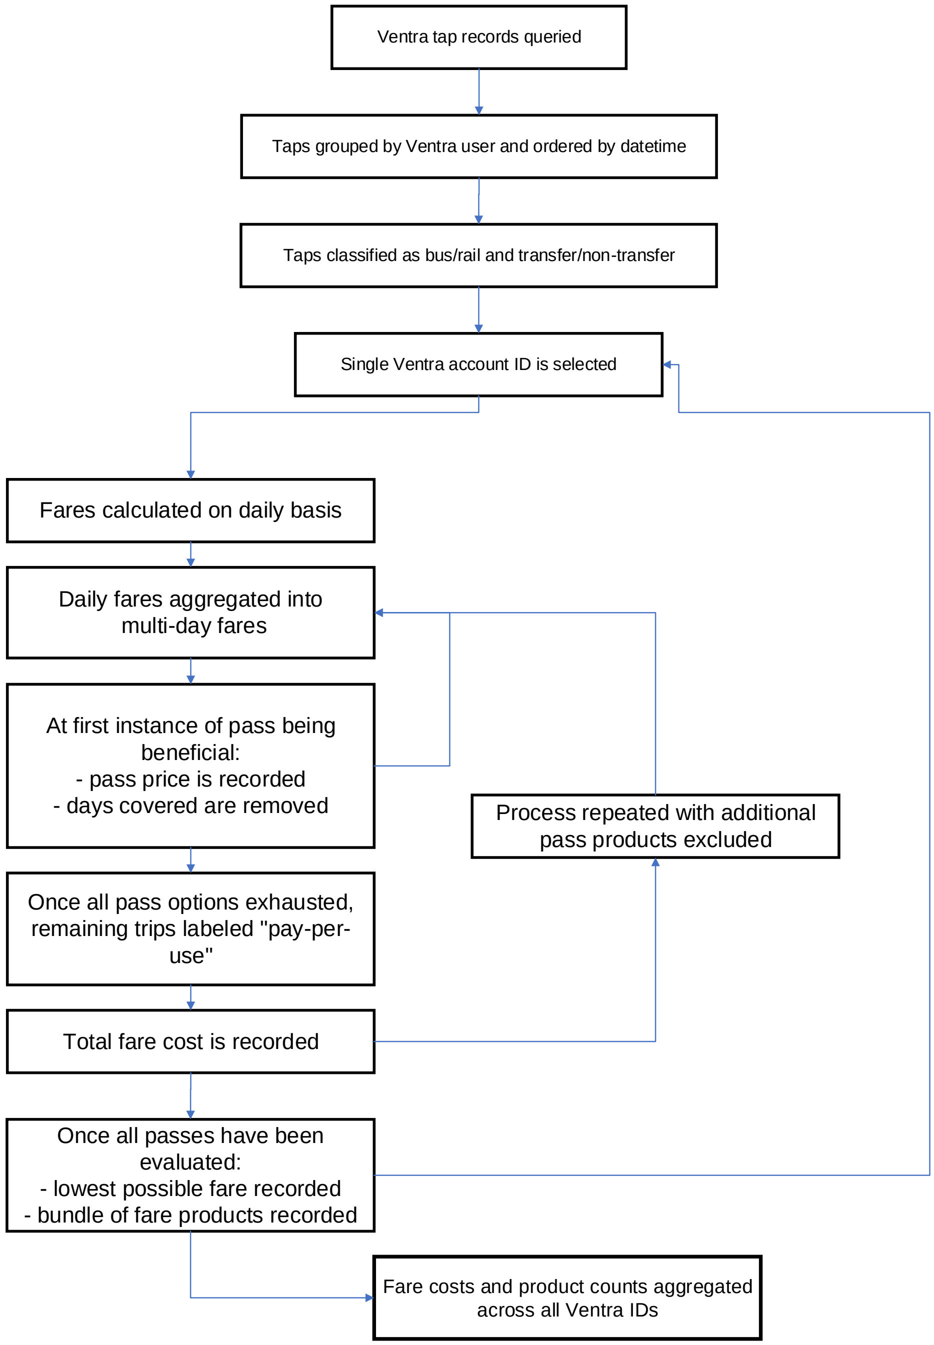

Figure 5-8: Schematic of Preferred Bundle Evaluation Procedure

uses that fact to experiment with hypothetical fare products. Using rider travel patterns from a post-COVID time period, in this case October 2021, products with novel designs can be introduced into the suite of available choices. To evaluate the benefits a certain flex-pass would provide to CTA users, the parameters governing its use must be developed.

Analysis of changes in travel frequency and volume, and simultaneous shifts in fare product choice, help identify two key ways that traditional pass offerings now match the travel of fewer riders. Those two clear changes are:

- **Changes in regularity or frequency of CTA travel.** The CTA's pass products confer an unlimited number of trips on a set number of consecutive calendar days. This design made sense when travel was an everyday, or nearly everyday, occurrence, but that is not the case for many after the COVID pandemic. These individuals may have a regular travel pattern that has become less frequent, or they may travel inconsistently but somewhat frequently. For either rider behavior, longer-duration pass products are unlikely to be the best option.

- **Reductions in number of trips taken on a given travel day.** The riders analyzed in this chapter take fewer transit trips on days when they did use public transportation in Chicago. Pass products are priced using assumptions about the number of trips on each travel day. Since commuting was a major purpose of transit trips, the default assumption was that most riders would ride twice, completing a round-trip. This case study revisits this assumption in designing novel flex-passes.

Flex-pass options that better align with these two changes were considered, and ultimately two flex-pass designs were tested. These two designs were determined in consultation with the CTA, ensuring that they aligned with CTA agency priorities and could potentially be implemented given existing policy and technological constraints.

The first flex-pass considered offered unlimited travel on a set number of CTA travel days, but those travel days did not need to be on consecutive calendar days. The pass can be seen as similar to a bundle of 1-day passes with the stipulation that they must be used until the bundle is fully used; other fare products, such as pay-per-use, cannot be used while this flex-pass is valid. The final design offered seven travel days, making the existing 7-day pass the obvious comparison point. The traditional 7-consecutive-day pass would remain, allowing riders to choose between both. However, given the value of additional flexibility with a non-consecutive day pass, a price premium is proposed. The exact pricing premium was determined through experimentation. I refer to this first hypothetical fare product as the "7-day flex-pass".

The second product design was centered on purchasing upfront a set number of CTA journeys as opposed to a set number of travel days. Riders are free to use as few or as many trips as they desire on any given calendar day, but there is additional flexibility in that one-way travel can benefit from this new product. Such a product provides a per-ride discount in exchange for an upfront purchase of a set amount of travel, providing the agency revenue at the start and creating an impetus for the buyer to ride, as no further transactions or costs are required until the journeys are all taken. The pricing for this product is set in the window between the per-ride single-use cost and existing passes based on round-trip travel. Ideally, frequent round-use travel should continue to be pushed towards traditional pass products and not undercut by this offering. Multiple permutations of CTA trips and per-ride discounts were evaluated to point towards the combination that would result in the outcomes that best approximate agency goals. This second flex-pass type is referred to as the "Trip Pass".

Metropolitan Region	Chicago, IL			
Transit Agency	Chicago Transit Authority (CTA)			
Product Category	Pay-per-use	Pass	Flex-pass	Commuter Rail Add-On
Fare Product	Bus Fare | Rail Fare	1-Day | 3-Day | 7-Day | 30-Day	7-Day flex-pass | Trip-Pass	Link-Up / Regional Connect Pass

Figure 5-9: Hierarchy of CTA Pass Products

5.4 Findings

The analytical process developed in the previous section is applied to AFC records to generate comparisons of two types. First, cross-time evaluation of existing CTA fare products and pricing is used to highlight the separate impacts of both COVID and pass price reductions on the relative utility of the CTA's fare offerings. Second, post-COVID travel patterns are used to asses novel, non-traditional flex-pass products in order to shed light on pass structure aspects that complement current travel behaviors. The results from both of these two applications are reported and discussed in this section.

5.4.1 Rider Benefits from Previous Pass Price Reductions

As previously laid out, following the onset of the COVID pandemic the CTA has incorporated changes to their fare product offerings ranging from new pass prices to more complete multi-agency integration (discussed more fully in the next chapter). Lowered pass prices occurred in a few steps, beginning with promotional pricing in early summer 2021 which was then made permanent that fall. Those pass prices have remained constant up to the time of writing, but in February 2023 all CTA passes also became jointly valid on Pace with no surcharge [65]. This will, in effect, lower the cost of joint 7-day passes gave CTA riders access to Pace at no extra charge. While recognizing that fact, the data in this case study are analyzed using the pricing levels adopted in Fall 2021 as shown in Table 5.1.

Figure 5-10 shows the primary fare product chosen by each Ventra account in the dataset for both months evaluated. What is clear is that the majority of riders in both periods chose to pay for transit on a trip-by-trip basis. Most of the shifting between these time periods was towards pay-per-use travel, with the largest shift in absolute terms coming out of the 30-day pass. This shift towards pay-per-use fares is expected, given less frequent and less consistent transit use among many riders post-COVID.

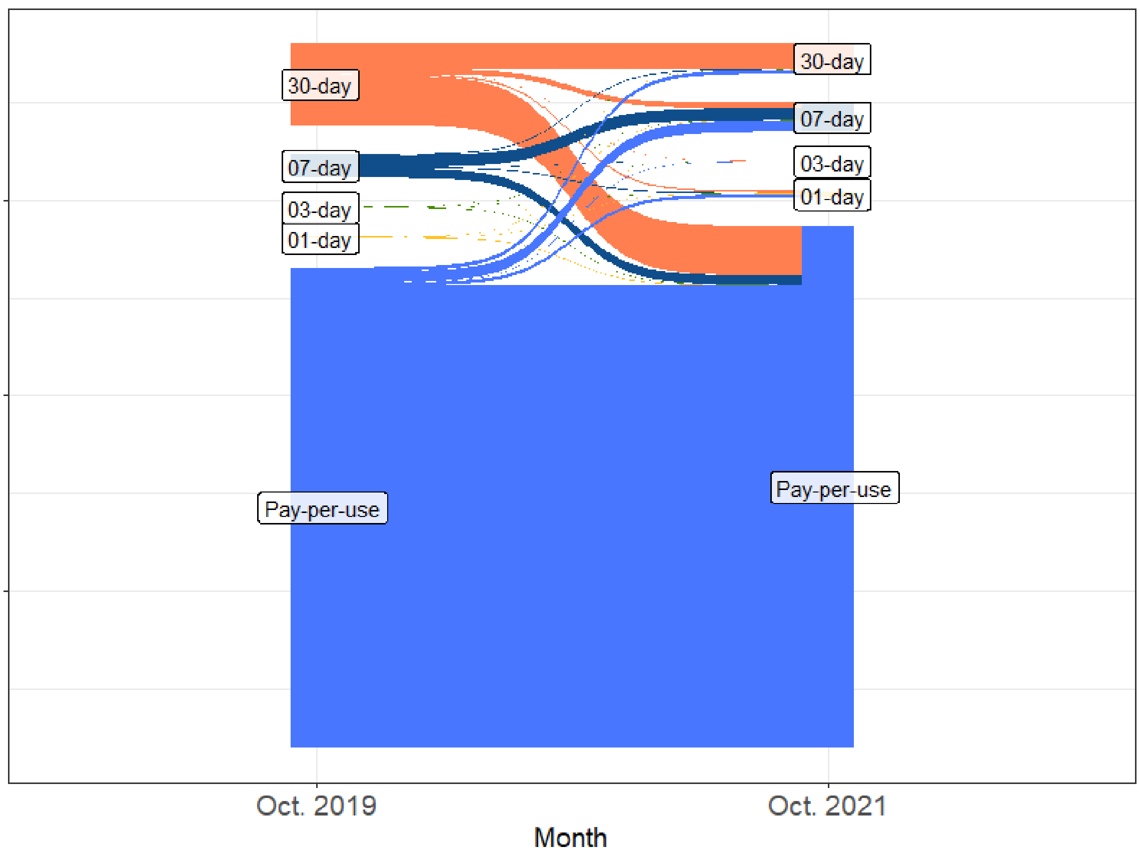

Figure 5-10: Actual Fare Product Choice by Ventra Account, October 2019 & October 2021

While Figure 5-10 shows what fare product CTA riders actually chose for travel, it does not imply that these were the optimal choices for each rider. As a part of the evaluation process, each rider is matched to the fare product(s) that comprise the best economic option. When the analyzed records are used to match the recommended fare products in October 2019, many of the riders are reassigned to different fare products. Figure 5-11 shows how these shifts would occur between actual product choice and best product choice. This reassignment should not be surprising; there are many potential reasons for inefficient fare product choice ranging from personal circumstance to mistaken expectations about upcoming transit use to comfort provided by habitual use of a particular fare option.

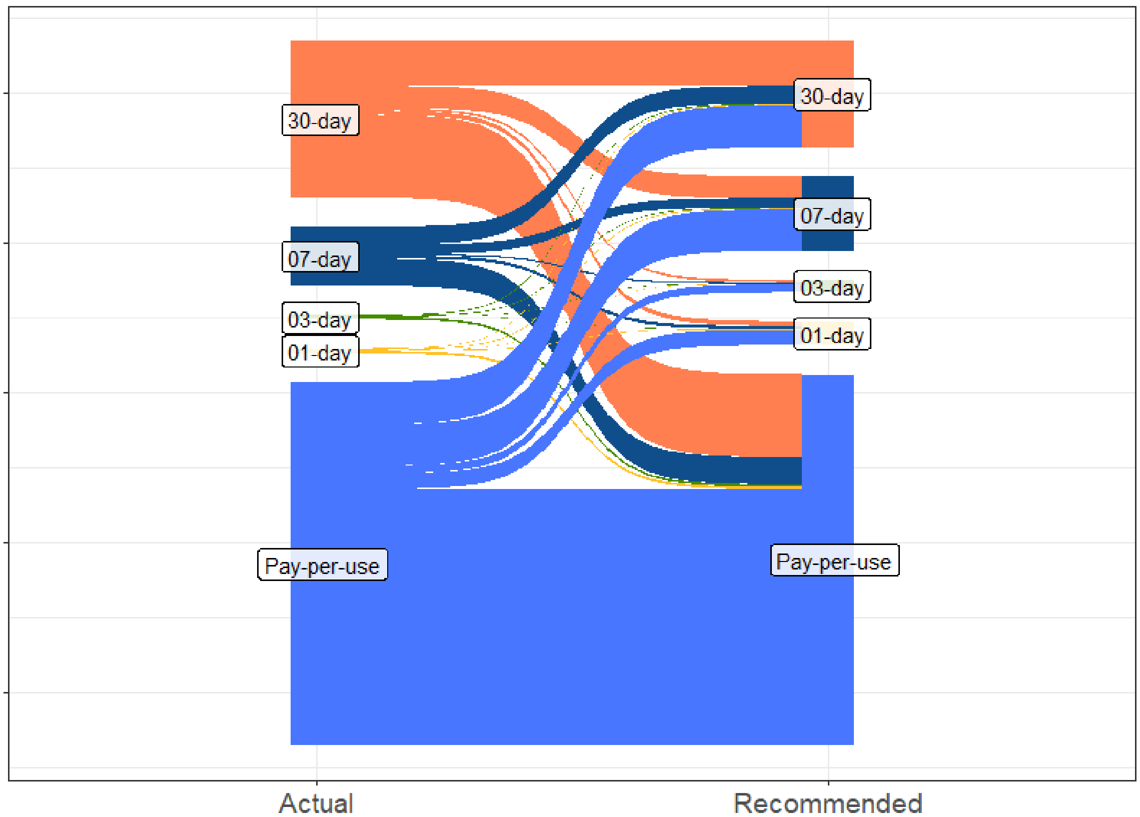

Figure 5-11: Actual and Recommended Fare Product Choice, October 2019

One key takeaway from this exercise is that optimal fare product selection does not result in any significant decrease in the number of pay-per-use riders. As shown, many 30-day pass users should have been using a different fare product. Most of the shifting should have occurred into the pay-per-use option, with some riders moving to 7-day pass products. Conversely, a significant number of pay-per-use riders could have benefitted from pass products, mainly the 30-day and 7-day offerings.

Repeating this process to determine the preferred fare products for each Ventra account using post-COVID travel shows that the recommended mixture of fare products has shifted much more dramatically towards pay-per-use fares. Nearly 75% of all trips taken by these riders should have been taken using single-ride pay-per-use tickets as opposed to any of the CTA's pass offerings. The largest loser due to COVID, in terms of competitiveness, is the 30-day pass. This is illustrated clearly in Figure 5-12. The y-axis measures how many rider travel days should have been taken on each of the fare products. Travel days is used to adjust for the different validity periods of all transit products. Riders utilizing multiple fare products have their travel assigned based on the travel days conferred. As an example, a rider that purchases a 1-day pass and a 7-day pass would have seven travel days assigned to the 7-day pass and one to the 1-day pass. In other words, $\frac{7}{8}$ of that rider's travel would appear under the 7-day pass.

Figure 5-12 adds another layer to this picture: the measured impacts of reduced pass prices on fare product attractiveness. Recommended fare product choice after lowering pass prices show that even given post-COVID travel patterns, each pass product was returned to a level of attractiveness at least equal to their position before the onset of COVID. Most fare products were actually made more attractive than they were in October 2019. This is most clear with the 1-day pass, which was not an attractive option for travel in 2019, but was made useful for many travelers with a new price. The travel patterns observed in October 2021, alongside fare changes, brought the amount of travel for which pay-per-use is recommended down substantially from its pre-COVID levels.

Looking at the actual fares paid by riders, and the amount they could have saved by switching to the recommended fare products, an average monthly savings can be calculated. Table 5.2 shows this amount for the most commonly used pre-pandemic fare products, alongside the number of riders for whom there was at least a $5 dif-

Figure 5-12: Impacts of COVID and Pass Pricing on Relative Fare Product Utility

ference between actual fares and the recommended fares. The number of affected accounts is shown as as both an absolute value and a percentage of the respective rider segment. As can be seen, 30-day pass users would have saved the most from switching to other products. There are also many more 30-day pass users that should have chosen other products, given the number of riders in this segment. Riders that used the 7-day pass could save nearly as much as 30-day riders on average. However, there were many fewer of these riders and a higher proportion of them could have benefitted from the recommended fare product. Some of these 7-day pass users would have benefitted from pay-per-use travel, while other should have purchased a 30-day pass. Those riders that could have benefitted from a 30-day pass might belong to the group of riders for whom upfront costs of a 30-day pass are too great, and for that reason they chose their second-best option. This table shows that counterintuitively it is pay-per-use riders who, as a group, stood to save the least from pre-pandemic product shifting. This is at odds with the CTA's preference to stimulate pass uptake and encourage transit use. This finding largely boils down to pass multiples. A 30-day pass required 42 trips to be more effective than pay-per-use travel. Given this high multiple, there were few riders that stood to benefit substantially from an upfront purchase, even pre-pandemic. The greater risk riders faced was purchasing a pass and not using it to its full potential. The number of riders that took substantially more than 10 transit trips each week was very low, which is necessary to make the 30-day pass a vehicle for substantial savings over pay-per-use fares.

Updating this analysis and applying to October 2021 travel with reduced prices, and looking at the potential savings to rider groups, the range of potential savings is much narrower across all fare products. The number of riders that would benefit from product switching is also much narrower, although 7-day pass users remain the group making fare product decisions least similar to their recommended products. These potential savings for 2021 are shown alongside the 2019 values in Table. Pay-per-use riders in the post-COVID period, with new pass prices, stand to save much more on average than they did pre-COVID. On the other hand, 7-day and 30-day pass users have smaller losses from suboptimal fare product selection.

Another outcome of the lower pass prices is its effect on recommended product mixing over a one-month travel period. Before pass prices were lowered, roughly three out of four riders could have purchased their travel at the lowest price using a single fare product, as shown in Figure 5-13. After pass prices were lowered, the number of riders for whom mixing and matching multiple fare products to meet their travel goes up significantly. The number of unique products in each purchase bundle is more than one for nearly all riders. Riders that should use two products more than doubles, with an even larger proportional increase in the number of riders that should be using three different fare products. These results indicate that pass price

Table 5.2: Potential Savings by Rider Fare Produce Choice

	Fare Product Choice		
	30-Day Pass	7-Day[a] Pass	Frequent[b] Pay-Per-Use
October 2019			
Accounts Overpaying by More than $5	12,683	6,758	4,927
% of Fare Product Users	*70.6%*	*95.8%*	*36.0%*
Average Potential Monthly Savings ($)	21	20	7
October 2022			
Accounts Overpaying by More than $5	5,609	5,429	5,646
% of Fare Product Users	*54.3%*	*89.1%*	*66.5%*
Average Potential Monthly Savings ($)	16	15	14

[a]Minimum two 7-day passes; no 30-day pass usage
[b]Fare spending equivalent to at least three 7-day passes; no 30-day pass usage

changes should have the effect of encouraging exploration of multiple fare products while concurrently providing an opportunity for savings to discerning riders.

5.4.2 Estimated Fare Product Choice with Novel Fare Products

Exploring the CTA's existing fare offerings in light of both travel shifts and pricing changes was insightful; the sequential nature of the changes allows the impacts of each aspect to be isolated. The changing utility of each pass product can be apportioned to either travel shifts or pricing changes, which together result in the utility experienced by CTA ridership at large. The next step is to use the same procedure to experiment with hypothetical fare product designs. The novel flex-passes explored were not assumed to replace any of the currently existing fare products; they were additions to the set of offerings. The CTA was specifically interested in considering additional products that enhance rider options. This was not an attempt to substitute for any of the products already used by current riders. Adding in new fare products to the analytical procedure has the benefit of illustrating which existing fare products would

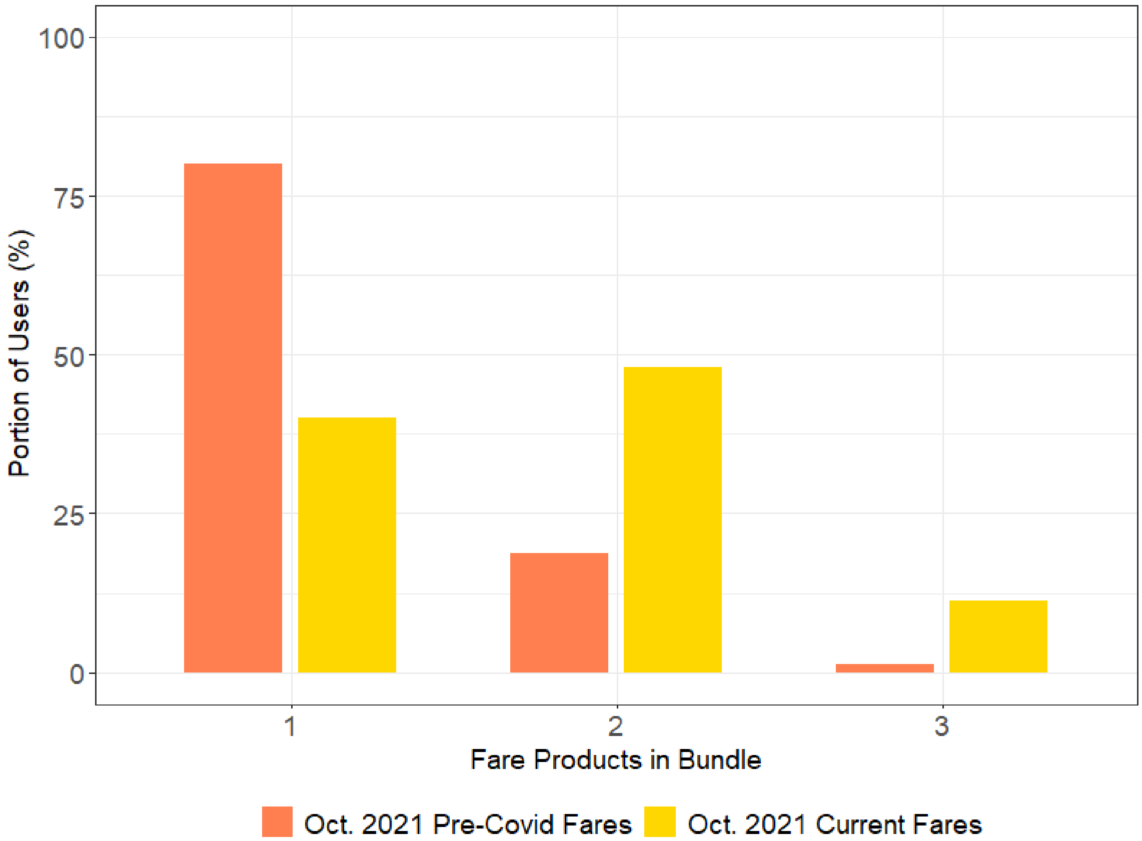

Figure 5-13: Recommended Fare Product Mixing

lose trips to any hypothetical fare products. If a fare product were fully removed, the riders using that fare product would all have to be reassigned; by simply adding a flex-pass only the trips that are more economically served are reassigned, indicating expected net benefits of any new flex-pass.

Both of the hypothetical flex-passes were evaluated individually for the post-COVID travel month of October 2021, and the results differed in interesting ways. The 7-day flex-pass was analyzed first. This product provides seven days of travel on the CTA for $25. That travel can occur on non-consecutive calendar days, but must occur within 21 calendar days of first use. Shifts in the recommended fare product choice are shown in Figure 5-14. This chart uses the same travel records from October 2021 to evaluate three pricing regimes: first, product choice with pre-COVID pricing; second, using the CTA's adopted price decreases; and third, with the hypothetical 7-day flex-pass introduced as a payment option. Evaluating the travel using pre-COVID fare product prices can help show whether the new flex-pass provides additional improvements over the CTA's current pricing, or whether any of the fare product groups may revert towards previous levels. Results indicate that the 7-day flex-pass would siphon trips from all existing fare products. Pay-per-use trips go down marginally, with 1-day and 30-day pass purchases also seeing a slight decrease. The largest shifting occurs from the existing 7-day pass, however, which would lose more than half of its trips to the new flex-pass. Travel on the 7-day pass would actually account for a smaller percentage of travel days than before the pass price reductions in 2021. These findings show that the flexible 7-day pass might make sense as a replacement for the existing 7-day pass, but as an additional product it does little to move riders away from pay-per-use and onto transit passes.

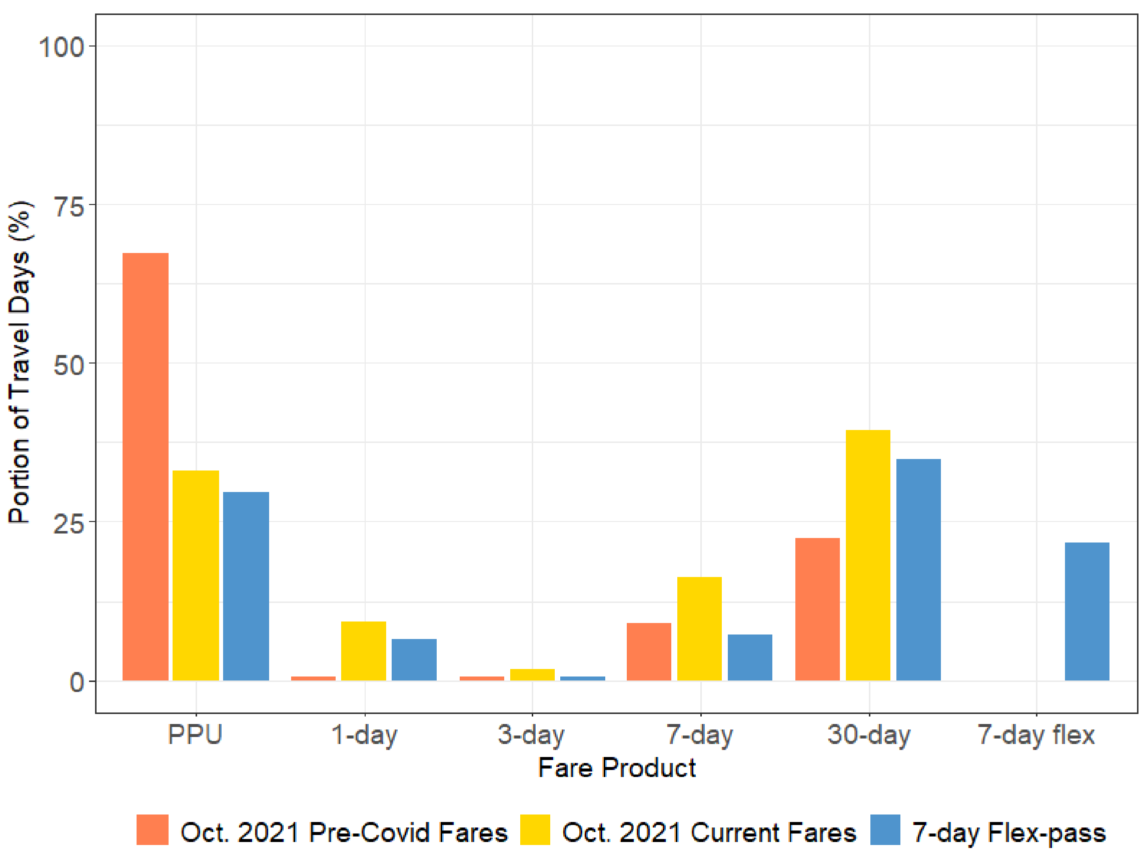

Figure 5-14: 7-Day Flex-pass Results

The second flex-pass design offers a fixed number of CTA journeys for a set price. The first iteration considered 10 trips for $17.50, or $1.75 per trip. Each trip includes systemwide transfers and there is no expiration date. As with the 7-day flex-pass, the trip-pass must be used for all CTA trips after it is activated until it is completely disbursed. This constraint on use was implemented for this analysis, but agencies could choose to implement such a product without requiring consecutive use. Analyzing the trip-pass with 10 trips at this price-point ($17.50) shows the largest shifts occurring out of pay-per-use travel and into the trip-pass. This is a positive result, as the goal is to design a product that appeals primarily to riders that otherwise would be traveling on single-ride pay-per-use fares. This can be seen in Figure 5-15. Pass products, namely the 7-day and 30-day offerings, see a slight decrease in selection. The 30-day pass sees a very small drop, whereas for the 7-day pass it is somewhat larger, but usage remains above where it stood prior to the 2021 pass pricing changes.

After analyzing both of these flex-pass designs, a case can be made that a bundled trip pass would be more effective than the 7-day flex pass in attracting pay-per-use riders while minimizing the impact on traditional pass product purchasers. Figure 5-16 below shows a direct comparison between the results for each flex-pass. The trip-pass was better for more travel, but only slightly. Nevertheless, it was successful in pulling substantially more travel away from pay-per-use fares, with lower travel taken from the pass products. The only pass option that lost more riders to the trip-pass then the 7-day flex-pass option were current 30-day pass users. However, total shifts from this pass are small, and only marginally different between the two flex-passes.

After selecting the trip-pass as the more promising flex-pass option, the fare product evaluation process was repeated while introducing different combinations of trips and prices. Each variation of the trip-pass offered an even number of journeys for ease of communication since many riders commonly think in terms of round trips.

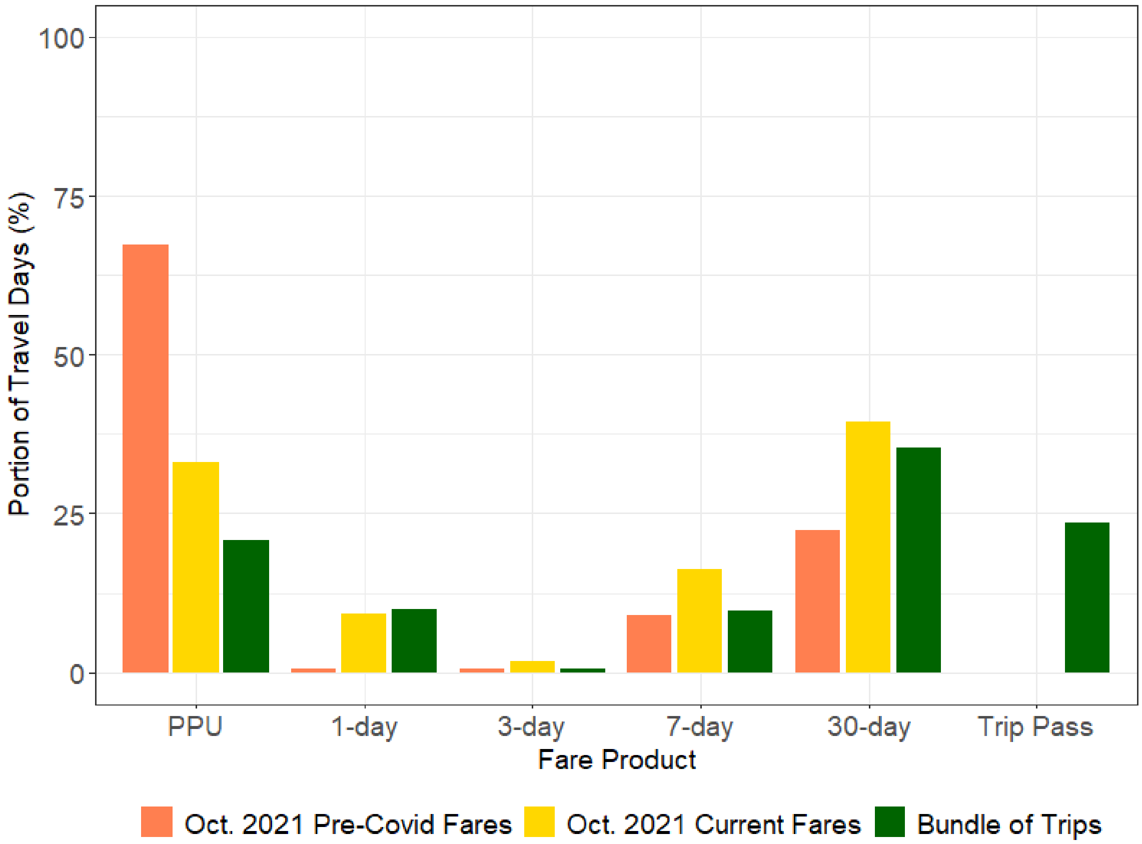

Figure 5-15: Bundle of Trips Flex-pass Results

Price points were set at increments of $5 with one exception where the price was set at a multiple of $2.50. The decision to predominately price in $5 increments conforms to CTA policy for pass products; this is done for ease of communication and the fact that certain payment channels require exact change. Figure 5-17 shows the results from iterating through the analytical procedure with trip-pass bundles of between 10 and 20 trips priced between $15 and $35. Having a clear goal in mind beforehand, namely to identify a fare product design that appeals primarily to pay-per-use travel with small impacts on pass product usage, allows these results to help ascertain the sweet spot for per-trip pricing and volume of trips conferred. A trip-pass that provides riders 10 start-to-end journeys on the CTA network for $20 performs best in terms of reducing pay-per-use travel while simultaneously supporting continued use of existing pass products.

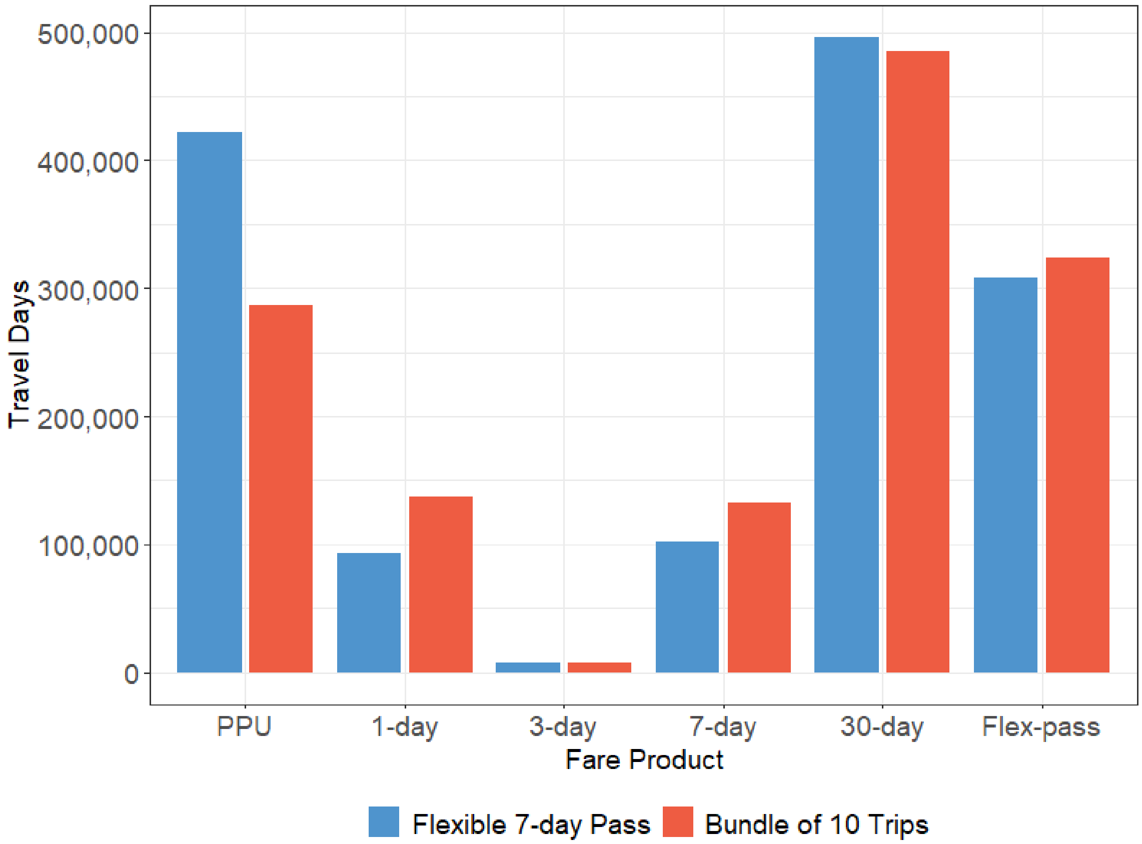

Figure 5-16: Preferred Fare Product Choice by Flex-pass Type

A combination of 10 trips for $20 has one added benefit to the transit agency: it can be framed as a complement to the existing 7-day pass which is also priced at $20. Using pre-COVID assumptions of ridership, that 7-day pass offers 5 round trip commutes, with evening and weekend travel included for free. This hypothetical flex-pass offering with 10 trips likewise could be used for 5 round trip commutes, but without the perk of traveling nights and weekends for free. Instead, it has the perk that those transit journeys can be used whenever the pass holder needs to travel, with freedom as far as travel day and little risk of buyer's remorse.

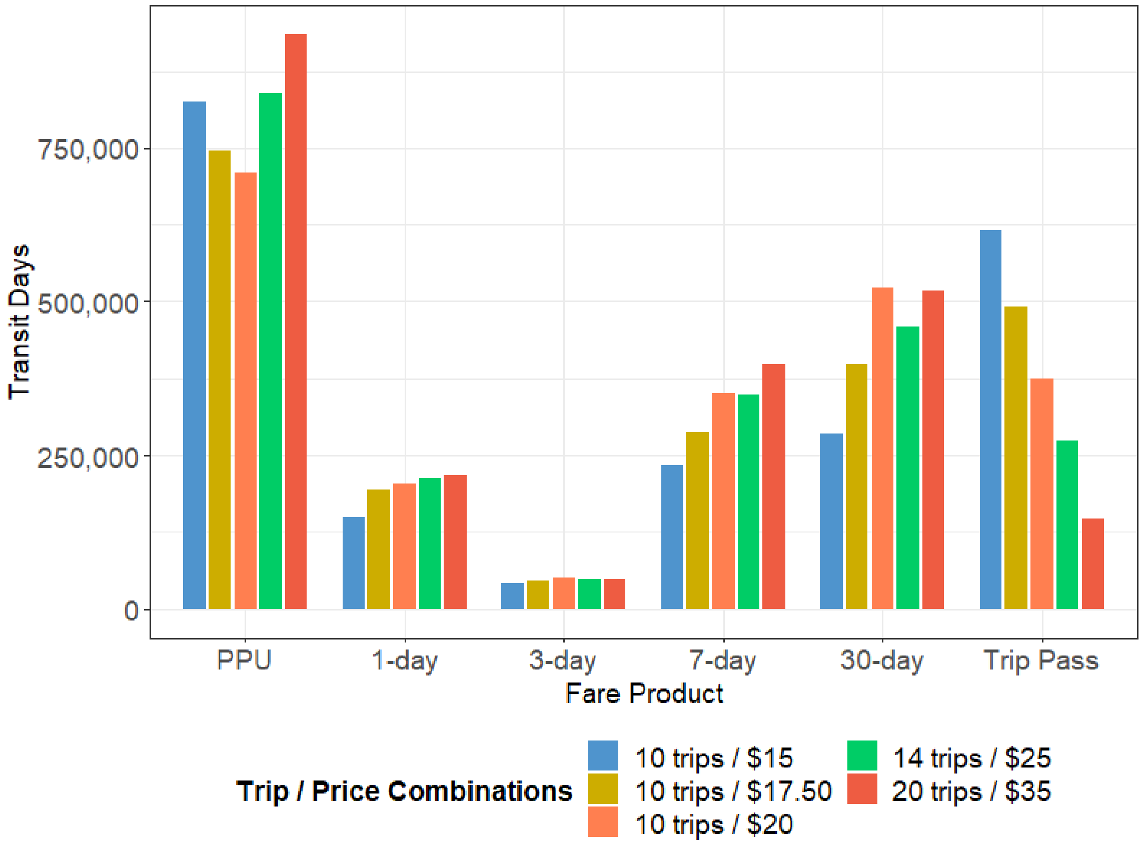

Figure 5-17: Results for Multiple Flex-pass Permutations

5.4.3 Potential Agency Impacts from Hypothetical Fare Product Switching

The analytical procedure presented in this case study is oriented towards questions of how well-designed fare products can be made to align with actual transit use. However, looking at actual fare choices made by CTA users shows that riders often purchase products that are economically suboptimal for their travel. This is most clearly demonstrated in Table 5.2 above showing the potential monetary savings for an average rider by choosing the fare products recommended.

There are many riders that do not choose an optimal fare product for their realized travel. This was shown in Figure 5-11. Nevertheless, the procedure developed in this chapter and the optimal fare product combinations it recommends can be used to approximate potential fare revenue impacts. Such approximations provide additional information to the agency as it considers piloting or introducing flex-pass offerings.

To extrapolate to the fare revenue impacts, some assumptions must be made about the relationship between actual and recommended fare product choice. The first assumption is that riders using transit fewer than five days per month will not switch to a new flex-pass offering. This is because the multiple for these passes occurs at 5 travel days. An existing product will always be at least as beneficial as the new flex-pass for these riders. The second assumption is that potential shifts to the new flex-pass product will occur in the same ratio as the (optimal) recommended fare product use. As an example, if recommended fare product choice has 20% of travel on the 7-day pass shifting to a new flex-pass, then 20% of actual 7-day pass users are assumed to shift to that new fare product. Third, any new flex-pass offering is equivalent to a set number of existing pass products as defined in Table 5.3. The number of flex-passes that would need to be sold to equate each current pass is a key part of the revenue change calculation.

	Flex-Pass	
Existing Pass	Bundle of Trips	7-Day Flex
1-Day Pass	5	5
3-Day Pass	2	2
7-Day Pass	1	1.5
30-Day Pass	$\frac{1}{3}$	$\frac{1}{3}$

Table 5.3: Pass to Flex-Pass Conversions

Using these assumptions, approximation of the fare revenue impacts of the 7-day flex-pass and the trip-pass offering 10 trips for $20 shows that the costs to the CTA would likely be small. The net change in fare revenues for October 2021 would have totaled a loss of about $310,000 if the 7-day flex-pass was offered. The estimated net change in fare revenues with the trip-pass introduced would be a loss of $50,000 dollars. These decreases in fare revenue are quite low considering the benefits they provide to the riding public, as well as the magnitude of CTA fare revenues. Lowering fare revenues may be a difficult choice for an agency to make, but the benefits to travelers are substantial.

Additionally, the fare revenue impact estimates are likely conservative in that the better performing of the two hypothetical offerings, the trip-pass, may well spur increased uptake as travelers use up their pass and decide to purchase another. The per-trip savings are small, but available to all bus and rail riders. The price-point is commensurate with the current 7-day pass. The flex-pass structure seems straightforward and easy to understand. Higher adoption rates may spur additional travel, and tapping into demand for more travel would bring in currently unrealized fare revenues, ultimately shrinking the amount of fare revenue lost.

5.4.4 Discussion

The methodology developed in this case study and applied to transit travelers riding on the CTA evaluates the utility of existing pass products in different travel environments. At a given price point the relative attractiveness of all products can be ordered, showing which products are most useful for the transit-riding public, and how effectively they accommodate travel patterns in aggregate. This analysis is focused squarely on the utility of product offerings, not on questions of how users choose a fare product. Those questions are important to the success of any new product; a well-designed product that sees little adoption is not achieving the outcomes for which it was introduced. Accordingly, analytical evaluation of marketing practices related to potential fare products would be a helpful step.

A more complete understanding of how CTA riders choose their fare product, and why so few of them mix and match from among the offerings, would be a useful extension of the case study presented here. Many riders make suboptimal fare product choices, and understanding why can help inform fare policy in productive ways. Are these results caused by the design of fare products (e.g., high upfront costs for passes), marketing (e.g., riders are not aware of their options), sales channel constraints (e.g., transactions are difficult to complete), or something else. In particular, further research on the value riders ascribe to flexibility, and the corresponding costs of uncertainty, would be highly beneficial. It is easy to purchase a monthly pass and not worry again about paying for transit; that is exactly what research on price salience would suggest. However, that can result in overspending, as observed among many pass users in October 2019. Similarly, pay-per-use riders may decide against a transit pass because they do not know what their travel requirements will be, and they do not need to try and estimate those patterns if they accept they can just pay for each ride. Pay-per-use is the easiest product to understand: the rider pays for each trip they need. No further consideration needed. For this reason, some agencies have chosen to implement fare capping, which is considered more extensively in 6. The analysis framework developed by this chapter can be extended to various fare capping implementations. However, understanding the perception riders have around fare simplicity (exemplified by pay-per-use) would help design products that benefit riders and predict adoption rates. Fare simplicity is understood to be important, but the way it is appraised by transit users as they choose a fare product remains unquantified.

Chapter 6

Opportunities and Drawbacks from Single-Day Fare Capping

6.1 Introduction

Chapter 5 explored how fare restructuring and pandemic-induced travel pattern changes impacted the relative utility of Chicago Transit Authority fare offerings. The chapter then proceeded to design, develop, and evaluate additional hypothetical fare products. These hypothetical fare products were referred to as "flex-passes", and they had more flexibility regarding around their usage constraints. This appendix considers a separate approach for advancing another rider-friendly fare policy: fare capping.

Fare capping is a fare system that does not require each rider to preemptively decide between purchasing a transit pass and paying individually for each journey, instead allowing riders to "earn" transit passes based on the number of trips they take. With fare capping, a rider simply taps their AFC card at the turnstile or on the farebox each time they travel and once they have spent an amount equivalent to the cost of a transit pass, that transit pass is conferred and the rider is not responsible for paying any further fares until the end of the period. To give a simple illustration, if a bus ride costs \$2.50 and a monthly pass costs \$50, the rider will pay for each trip until they have spent \$50 on transit, at which point the fare system would recognize the rider as traveling using the monthly pass. No matter how many trips a riders takes in a month they will never pay more than \$50.

Some people are drawn to fare capping because it has the potential to benefit riders by eliminating the opportunity for any overpayments. Nevertheless, there is a difference of opinion among two camps regarding the appeal of fare capping. One opinion belongs to those who think that the surest way to boost transit ridership is to increase the uptake of transit passes. Traditional transit passes provide unlimited travel, meaning the marginal cost of each trip is \$0 once an upfront commitment to transit is made. The downside to fare capping, according to this philosophy, is that the marginal cost for each journey is unchanged until the cap is reached. This creates a disincentive to transit travel that is absent from pass products, potentially lowering overall transit use. Conversely, the second camp believes the upfront commitment of a transit pass is a major barrier to entry. This group believes that if riders could gradually "earn" a pass through just by traveling, the public would be both more apt to travel and more apt to generate additional transit trips after the fare cap has been reach. In addition, riders would not need to break out a calculator and a crystal ball before traveling to determine exactly how often they will be using transit and which product is the best bet.

Awareness of fare capping as a fare policy has increased over the last number of years. Forms of fare capping have been implemented at a handful of agencies with some key differences. One of the aspects that sees the widest variation between implementations is the duration or number of products considered for capping. There is no uniform approach. For instance, some agencies apply caps on a weekly basis and others on a monthly basis. Some agencies offer a mixture of pass products and capping options, others offer only fare capping. Table 6.1 shows how a subset of agencies have chosen to institute fare caps.

Agency	**Metro Region**	**Capped Pass Products**			
		Day	Week	Month	Year
AC Transit	Oakland, CA	X	X	X	
IndyGo	Indianapolis, IN	X	X		
Transport for London	London, UK	X	X		
TriMet	Portland, OR	X		X	
Valley Transportation Authority	San Jose, CA	X			
MTA	New York City, NY		X		
Mass Transit District	Champaign-Urbana, IL			X	X

Table 6.1: Fare Capping Implementations

6.2 Research Goals

Given differences in implementations, this chapter explores fare capping in the context of the CTA. Given the CTA's existing fare products and ridership, alongside agency priorities, it is worth evaluating whether some forms of fare capping could jointly benefit both riders and the CTA. The CTA's price reductions for passes in Fall 2021, which did not include corresponding reductions in single-ride pay-per-use fares, indicate that the agency is interested in encouraging pass use among their customers. The goal of increased pass use is to hopefully generate increased transit travel volumes, as with passes the average cost per trip is reduced each time a trip is taken and, simultaneously, the marginal cost per trip is always set at $0. Even after incentivizing pass use, the CTA continues to see high usage levels of single-ride pay-per-use fares.

Consider the cost of a 1-day pass, which the CTA reduced to $5.00. This cost is equivalent to taking two rail journeys. All rail riders taking at least one roundtrip have no reason not to purchase a 1-day pass. At this reduced price the uptake and usage of 1-day passes increased, but pay-per-use nevertheless remained the dominant choice for fare payment. When considering the CTA's clear preference to incentivize pass use alongside the fare products currently offered, the utility of single-day fare capping becomes an interesting implementation option. Single-day fare capping would, in essence, support the same goals as the current 1-day pass offering, but without requiring a decision to purchase upfront, and without any corresponding chance of buyer's remorse. Further, single-day fare capping would not devalue any of the longer-duration passes offered, each of which would remain much more beneficial for frequent CTA travelers. The hierarchy of pass products would remain unchanged, continuing to support the CTA's desire to incentivize pass purchases.

The high usage of pay-per-use travel, and the competitive pricing of the 1-day pass, raises the question of how often CTA users take exactly two trips per day. Figure 6-1 shows the percentage of riders doing this for the six-month period beginning in January 2022 and continuing through June 2022. The red column represents each of the individual travel days for which the fare was either a 1-day pass or single-ride pay-per-use. The smaller yellow column alongside shows the number of all travel days on which the user took exactly two separate CTA journeys, but no more and no less. As the figure demonstrates, the percentage of travel days composed of exactly two trips is quite stable over this period, always falling between 40% and 42%. These riders already taking two trips would be targeted by a single-day fare capping scheme. Many of them already pay fares equivalent to the existing day pass, so setting the cap at the same price point would not reduce the fare revenues these riders are already generating. If these two-trip riders were induced to take another trip at no additional cost, all of the extra ridership would be a win for the CTA without losing corresponding fare revenues. Potential revenue losses would only come from riders that are already taking more than two trips on the CTA system, but they compose a much smaller portion of the ridership.

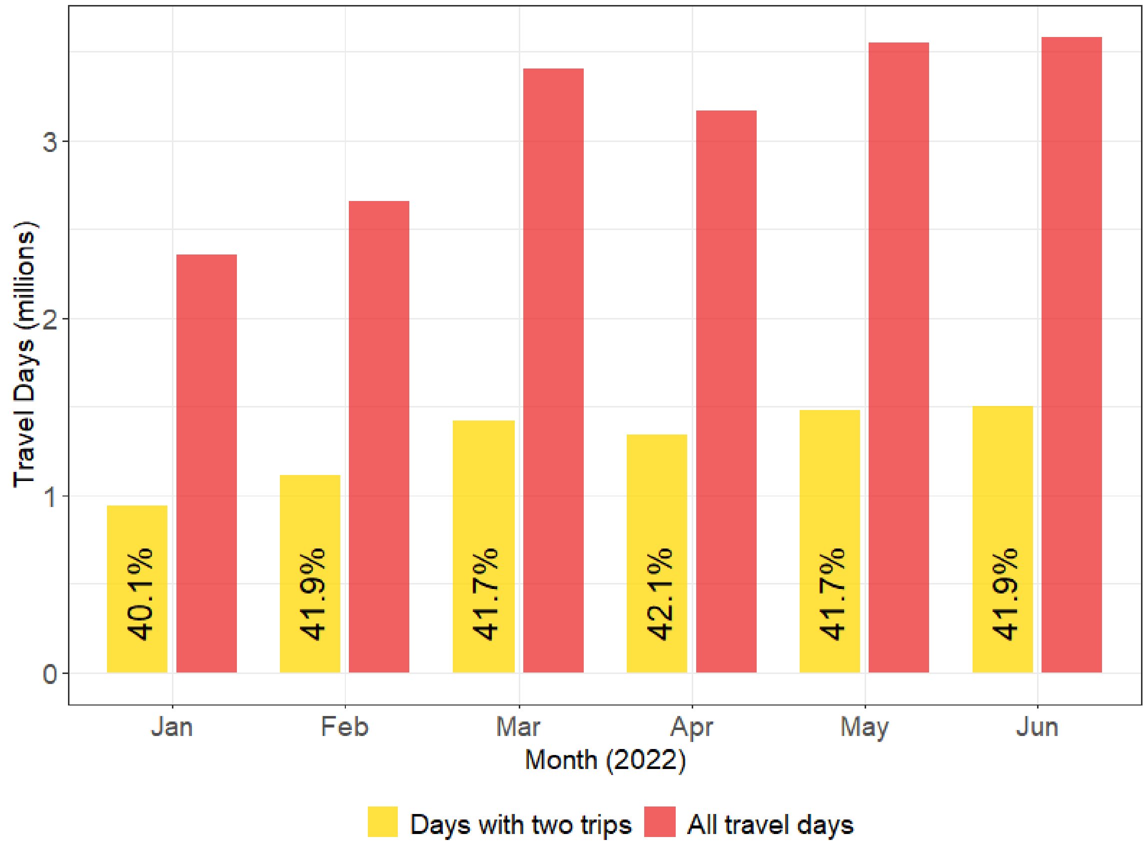

Figure 6-1: Percentage of CTA Travel Days with Two Trips

6.3 Methodology

6.3.1 Accounts Identified for Analysis

To further evaluate the potential impacts, the first step is to better understand the riders that would be affected. This is done using the AFC data available for CTA travel. This information includes the fare product used for travel, from which the appropriate rider segments can be identified for analysis. It also includes identifiers for the first trip in a journey as determined for fare payment purposes, alongside identifiers for the bus stops and rail stations visited by each Ventra account.

With single-day fare capping, affected riders would be pay-per-use travelers and riders that currently purchase the existing 1-day pass. The number of riders belonging to each of these two groups is not equal, however. Figure 6-2 shows the number of trips taken with pay-per-use fares and 1-day passes between January 2022 and June 2022. The vast majority of the affected riders are those paying for each trip individually. 1-day pass purchasers form a small minority in comparison. The data shown in Figure 6-2 only includes trips taken with associated Ventra accounts, as payment by some identifiable card would be necessary to track fare payments and apply fare caps. Cash fares and limited-use paper tickets would not be usable for capping purposes.

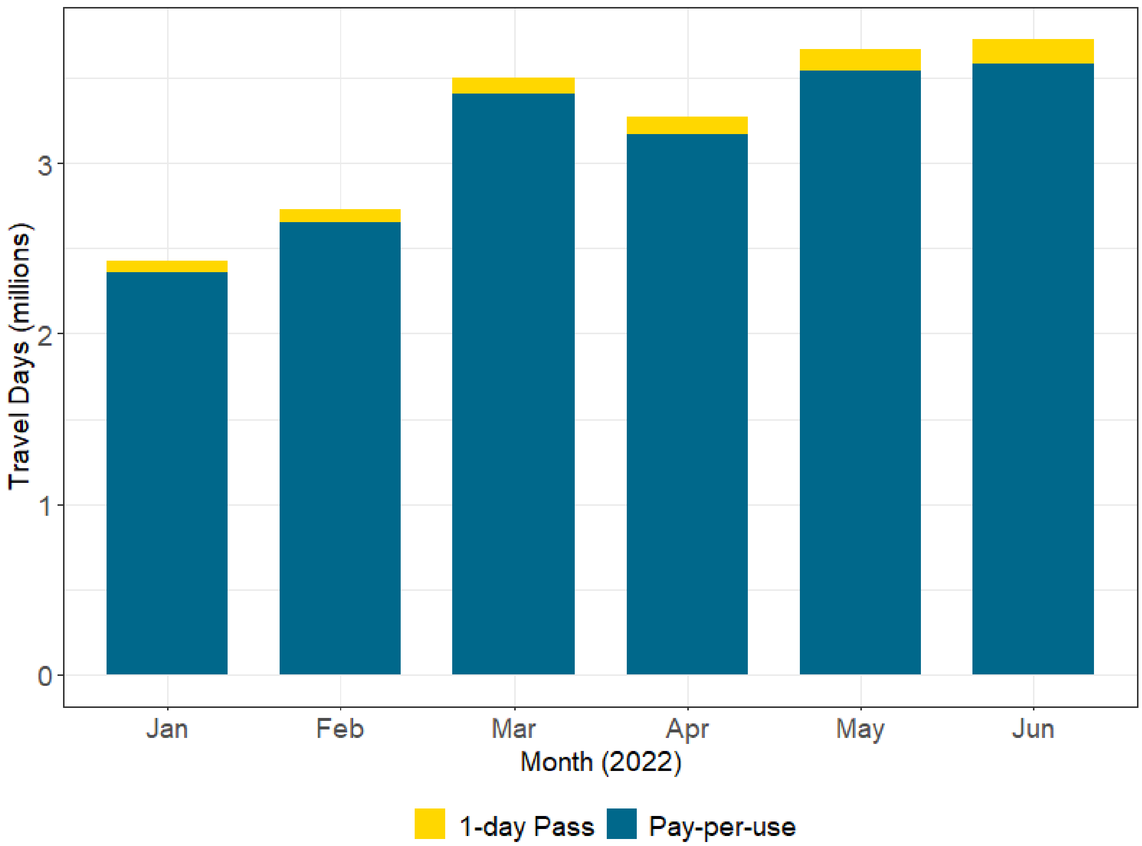

Figure 6-2: Travel by Pay-Per-Use and 1-Day Pass

Fare capping prevents riders from overpaying for their travel. Pay-per-use riders would never pay more than the cost of a 1-day pass. 1-day pass purchasers would not have to worry about buying a pass and then not using its full value. Riders of both types benefit, but seen through a different lens these rider savings can also be viewed as fare revenue losses to the transit agency. We can measure what the agency revenue loss from single-day fare capping would have been and categorize the loss based on whether it is from capping pay-per-use fares or from preventing 1-day pass overpays.

6.3.2 Transit Value Usage Differences

Even if the absolute levels of overpayments are known, it is important to understand their spread across riders. The existing 1-day pass products offered by the CTA can be purchased in two ways. First, the pass can be purchased by a Ventra account holder and loaded onto an AFC card for use. Alternatively, customers can purchase a paper ticket that is valid for one day. This can be done with cash or a bank card, and this limited-use pass is not associated with any Ventra account. Riders purchasing 1-day passes through these two different channels are likely dissimilar in certain dimensions. Limited-use paper passes may be more frequently purchased by members of disadvantaged groups using cash who lack access to a Ventra card. Paper-ticket based passes could also be indicative of infrequent CTA users or tourists that do not have any need for a reusable Ventra card. Figure 6-3 shows the distributions of CTA travel value used by customers on these two 1-day pass media. The figure shows that Ventra account-based 1-day passes are frequently underutilized by their purchasers. Slightly over 50% of the passes purchased with an associated Ventra account did not take CTA trips totaling at least $5.00. Conversely, 1-day passes that were sold as limited-use paper tickets were much more likely to reach at least $5.00 worth of transit travel. Only around 20% of the paper passes were used for less than their full value, compared to over 50% of account-based 1-day passes. 1-day paper tickets also witnessed more riders traveling well beyond the $5.00 value. Nevertheless, riders having Ventra accounts would benefit from fare capping, as overpayment through purchasing 1-day passes would no longer be possible. On the other hand, paper-ticket users would not be impacted by implementation of a daily fare cap, as only identifiable Ventra accounts are subject to the cap.

Looking at the distributions in Figure 6-3, one may wonder how so many CTA

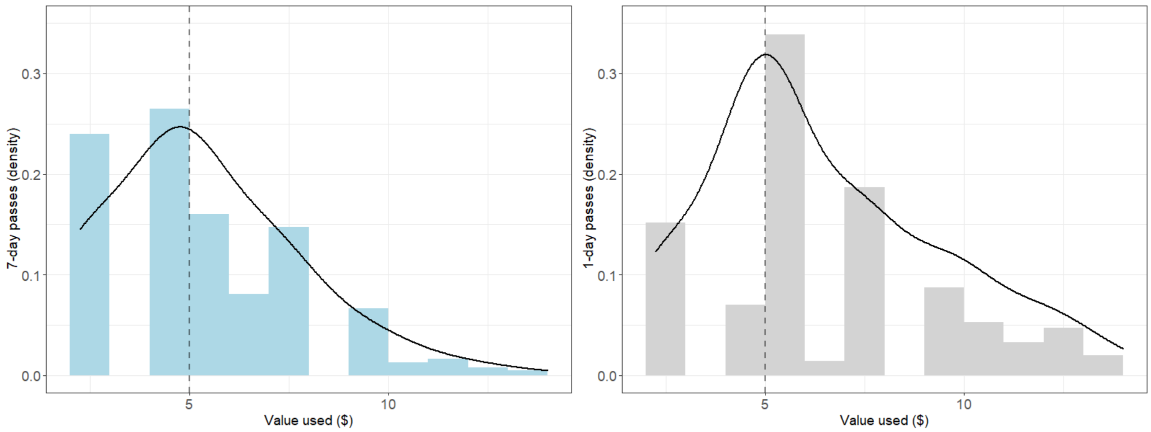

(a) Account-based (b) Ticket-based

Figure 6-3: Transit Value Usage Distributions of 1-Day Pass Users

users managed to spend between $4.00 and $5.00 on transit travel over the course of a single calendar day. Among account-based passes, this transit-value bucket accounts for over 25% of all 1-day passes. To fall into this bucket full-fare riders, who together form the vast majority of all CTA users, must either take two bus trips, or one bus trip and one rail trip. This highlights a specific area of difficulty for the existing 1-day pass; the pass is only an optimal choice for riders that either take more than two CTA journeys in one day or that take two journeys both originating on the rail network. For all other CTA two-journey users, it is actually less expensive to purchase individual single-ride fares.

6.4 Findings

6.4.1 Fare Revenue Impacts

Figure 6-4 shows CTA fare revenues generated from pay-per-use fares and 1-day passes alongside the fare revenues calculated assuming a $5.00 daily cap had been instituted. In this figure, a $5.00 cap is used because that is the price at which the 1-day pass has been set, allowing this hypothetical single-day cap to be a perfect substitute for the existing pass. As the figure shows, the decreases in fare revenue would be relatively small each month. In no month is the difference in fare revenues between scenarios more than 5.4%. The numbers displayed in the chart show only actual travel, with no adjustments or assumptions about induced demand from a fare cap. The numbers also indicate, however, that revenue neutrality could be achieved if 14% of travelers that took a single CTA trip one day decided to take a second. Revenue neutrality would also be achieved if 35% of the Ventra accounts active took one more paid trip on another day.

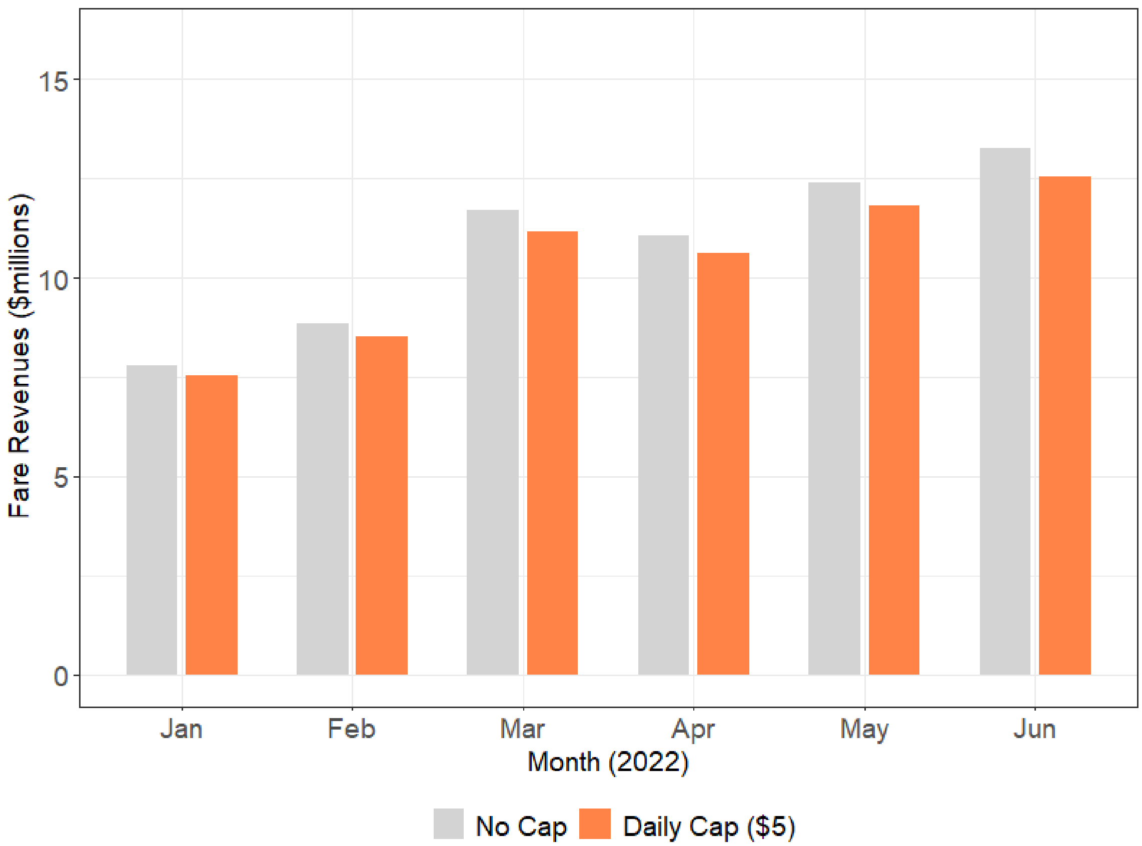

Figure 6-4: Fare Revenues with $5 Daily Cap

Beyond the total savings to riders, the number of individual Ventra accounts that stand to benefit from fare capping is a necessary consideration. It is important to understand whether the benefits are accruing only to a small number of high-frequency riders that should already be using pass products, or whether the benefits would accrue in a more widespread way across Chicago's transit-riding public. Figure 6-5 shows the subset of pay-per-use riders that would reach a $5.00 daily cap, as well as how frequently those riders would reach the cap. As can be seen, about 40% of accounts would infrequently, if ever, reach the cap. Another 50% or so would reach the cap no more than one out of five transit days. The remaining 10% of accounts are those that more frequently spend at least $5.00 on pay-per-use fares during a single calendar day. These numbers suggest the benefits of daily fare capping would accrue to many Ventra accounts, and the fare caps would not have been regularly effectuated for the majority of riders. All of the riders who spend more than $5 per day on pay-per-use fares should have purchased a pass, and all riders who spent exactly $5 could have purchased a pass and achieved the option of additional travel.

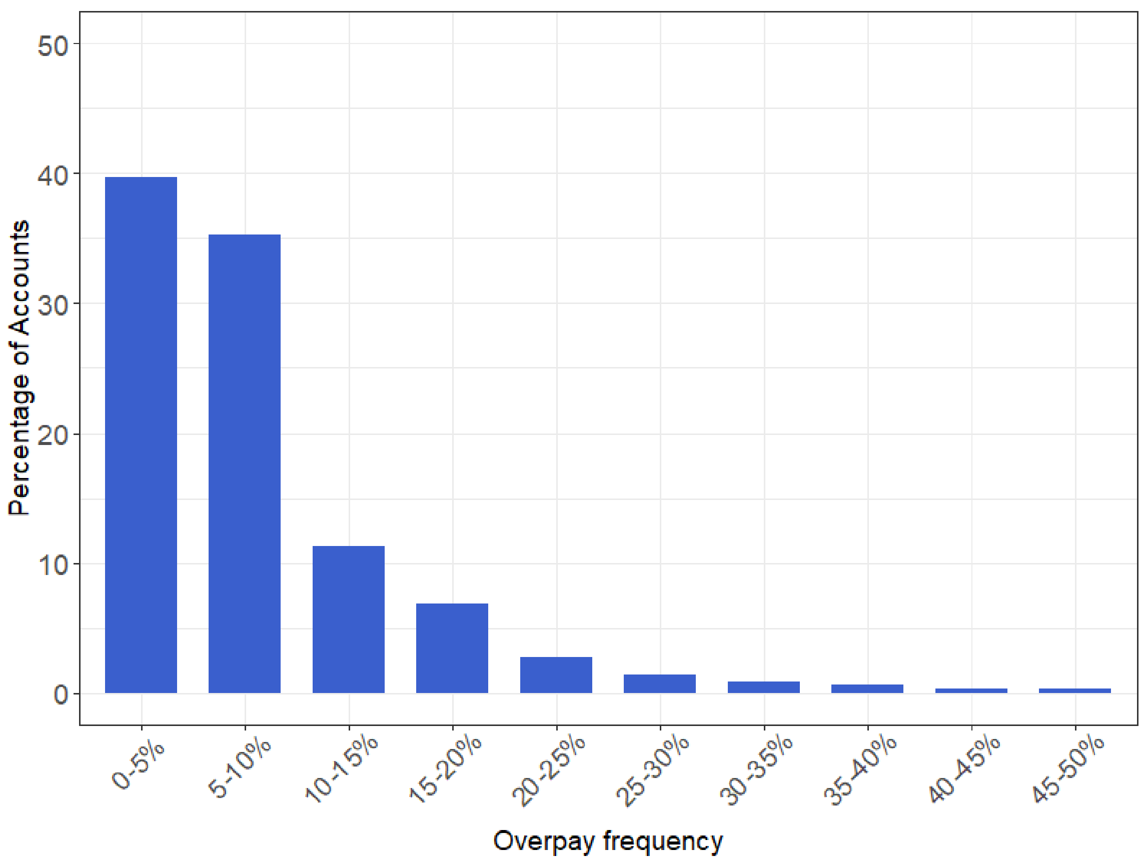

Figure 6-5: Overpayment Frequencies of Pay-Per-Use Accounts

Since pay-per-use overpayment occurs periodically among many Ventra accounts, it is worth considering whether this is because riders are truly unaware of their future travel needs or whether riders avoid considering their travel. Figure 6-6 contains only the subset of Ventra accounts that have purchased both 1-day passes and pay-per-use fares between January 2022 and June 2022. To be included, the riders must have traveled on at least twelve calendar days, and they must have purchased at least six 1-day passes. The figure shows the frequency of fare overpayments among this subset. Here, overpayments are defined as either pay-per-use fares totaling more than $5 on a single calendar day, or less than $5 worth of travel on a 1-day pass. One would expect these riders have a lower percentage of overpayments when compared to the much larger cohort of riders shown in Figure 6-5. This hypothesis turns out to be partially true, but not completely. The riders that have mixed-and-matched products did overpay less often, with about under 35% managing to overpay no more than 5% of the time. However, this subset of riders still overpay regularly; 20% choose the suboptimal fare product on at least 1 out of every 10 travel days. While these overpayments are less frequent than observed among the entire ridership, overpayments are still occurring regularly.

Taken together, these two figures indicate that even with a 1-day pass offering that costs only $5.00, a single-day fare capping product would benefit many riders. Ideally, it would also convince riders who have taken two trips to take another, as there would be no further cost to utilizing transit for the rest of the day. As demonstrated by the subset of riders who clearly tried to discern the best fare product for individual travel days, certainty around upcoming transit travel is difficult to achieve. The large majority of travel days already occur using pay-per-use as opposed to passes valid for a single day, as shown in Figure 6-2. Providing these riders an incentive to keep traveling should encourage additional trips. An upside of doing so, as found in Figure 6-4, it that impacts on monthly fare revenues would be small. The savings would accrue over many riders, not solely those frequent riders that consistently overspend

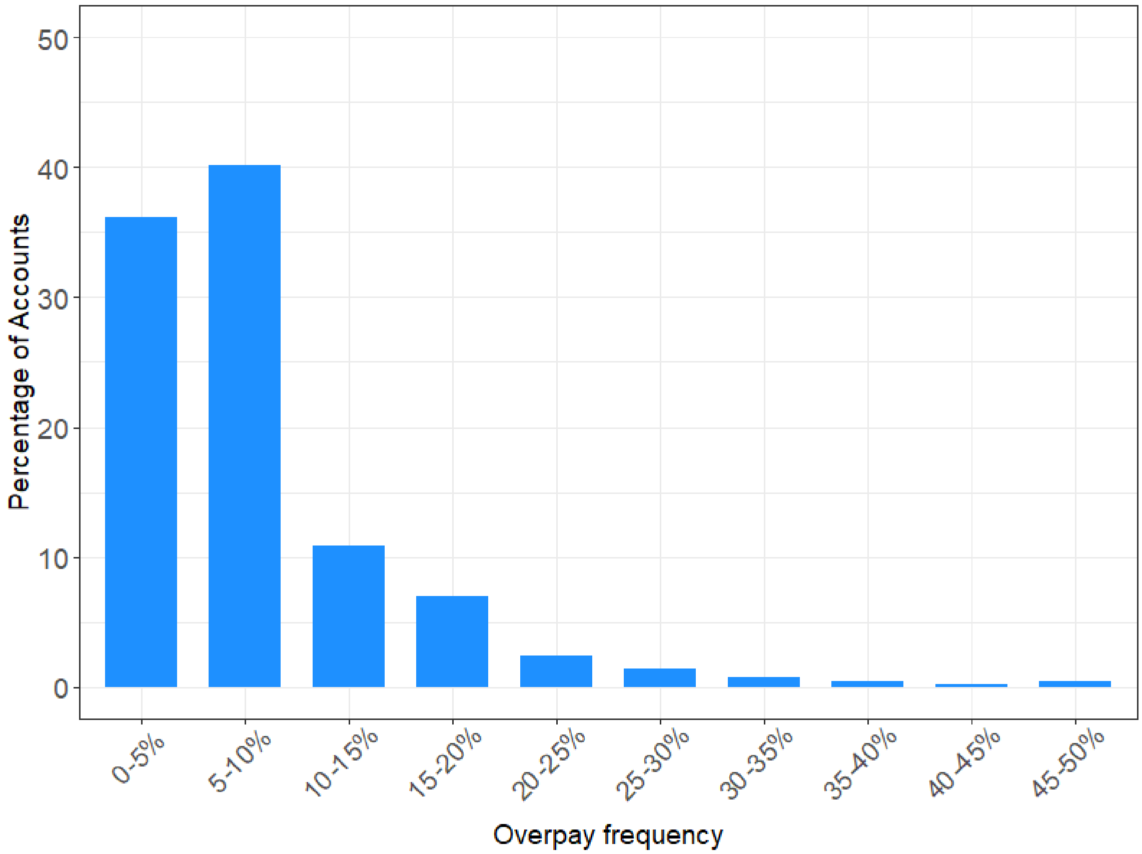

Figure 6-6: Overpayment Frequencies of Fare Product Mixing Accounts

by making the suboptimal choice of paying for each trip individually.

6.4.2 Consideration of Other Caps

An additional consideration concerns the level at which the fare cap is set. Given the different pricing for bus and rail services, not all two-trip days on the CTA will cost \$5.00. Two-trip days could in fact cost as little as \$4.50 if the rider only uses the bus. To understand how many two-trip riders would achieve a \$5.00 cap, we can explore the combination of CTA services used by the public. Figure 6-7 shows that less than 50% of travel days composed of exactly two trips would be subject to a \$5.00 daily cap. The remainder of the travelers would be able to take a heavily discounted third trip, but a small fare would still be required before hitting \$5.00 in transit fares for the day. This suggests that perhaps a different daily-fare cap should be considered: a two-trip cap. Capping fares after the first two trips would expand the benefits to all Ventra accounts regardless of whether they are bus riders or rail riders. Everyone would have access to continued travel for the remainder of the day at zero marginal cost. Any additional costs to the CTA would likely be small, as full-fare riders would be saving at most $0.50 compared to the $5.00 cap.

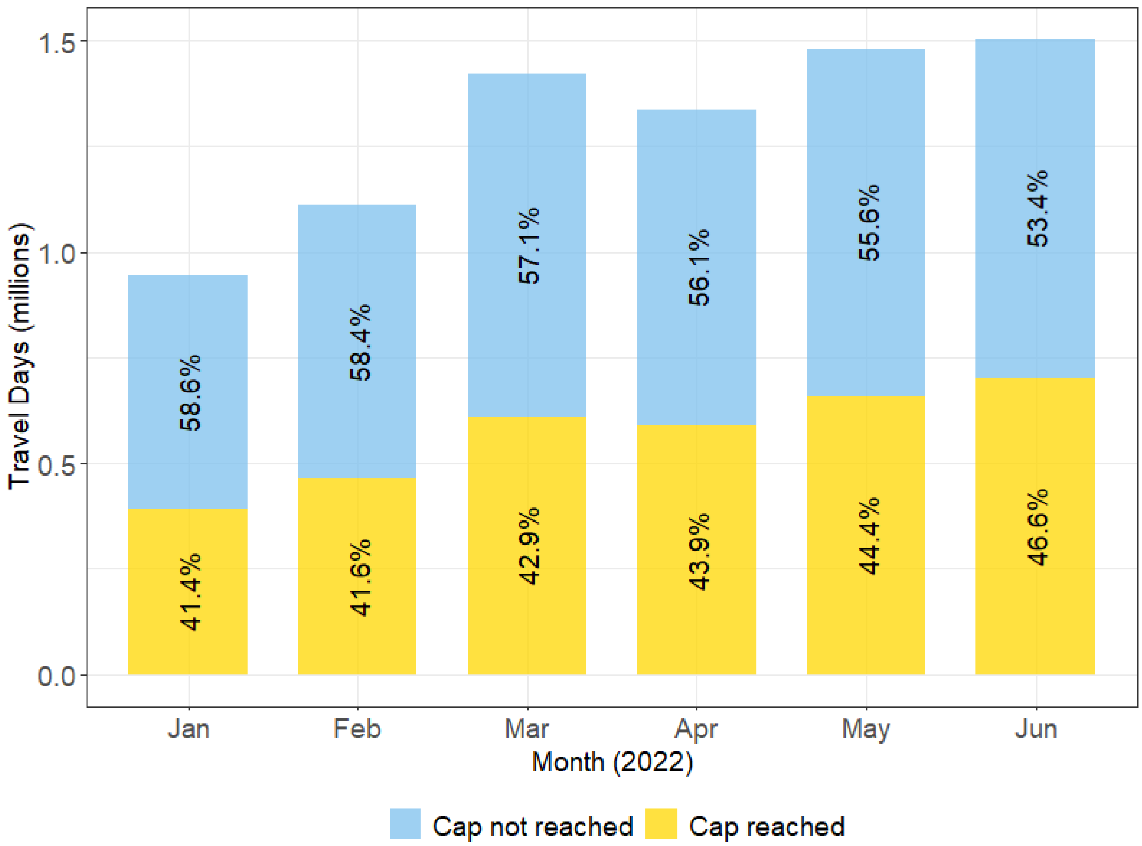

Figure 6-7: Percentage of Two-Trip Days Reaching $5 Fare Cap

The benefits of a two-trip cap as opposed to a $5.00 cap become even more compelling when riders paying reduced per-trip fares are considered. Reduced fare riders, seniors, and students currently have a discounted per-trip fare. Of the CTA's pass offerings, only the 30-day pass is available at a lower price to reduced fare riders. The cost of 1-day, 3-day, and 7-day passes are the same across all fare classes and rider groups regardless of how much each pays for a single journey. This runs counter to the CTA's goal of encouraging pass adoption, as the incentives are not consistent across rider populations. Consider a reduced fare rider, for whom rail fares are $1.25; the 1-day pass is only beneficial if they take at least 4 rail trips. This is equivalent to a $10.00 price point for the 1-day pass among full-fare riders. As a reminder, $10.00 is precisely the previous 1-day pass price before the CTA's fall 2021 pass price reductions. A single-day fare cap that is triggered after two trips would not only extend the benefits of a 1-day pass-like product to bus riders, it would also provide a 1-day pass-like offering to riders regardless of fare class. It also removes difficulties on the agency side regarding how to price passes and account for each of the various single-ride fares charged by the CTA. The rider decision and agency communication processes would both be simplified, as would the legibility and equity of the fare system. The CTA's message to riders could become that monthly and weekly passes are still available and beneficial for frequent riders, but no matter what each rider will not ever need to pay for more than two transit journeys on a single day.

Any additional fare revenue losses incurred by using a two-ride cap instead of a $5.00 cap for single-day travel is expected to be small, as for the vast majority of riders two trips costs between $4.50 and $5.00. There is a much smaller portion of riders that are paying reduced fares and for whom two CTA journeys would total less than $3.00. Figure 6-8 confirms this is the case. Fare revenues collected from pay-per-use fares and account-based 1-day passes are shown alongside the fare revenues that would have been paid by these riders had there been a $5.00 single-day fare cap, as in the previous Figure 6-4. One additional column has been added to this figure that shows the anticipated fare revenues under a single-day cap that is set at two CTA journeys. Additional fare revenue loss caused by moving from a $5.00 cap to a two-ride cap is minimal. This supports the position that a single-day fare cap, if implemented, should be reached based on the number of journeys taken, not a set dollar value. The benefits to riders are clear, and the costs to the CTA are marginal.

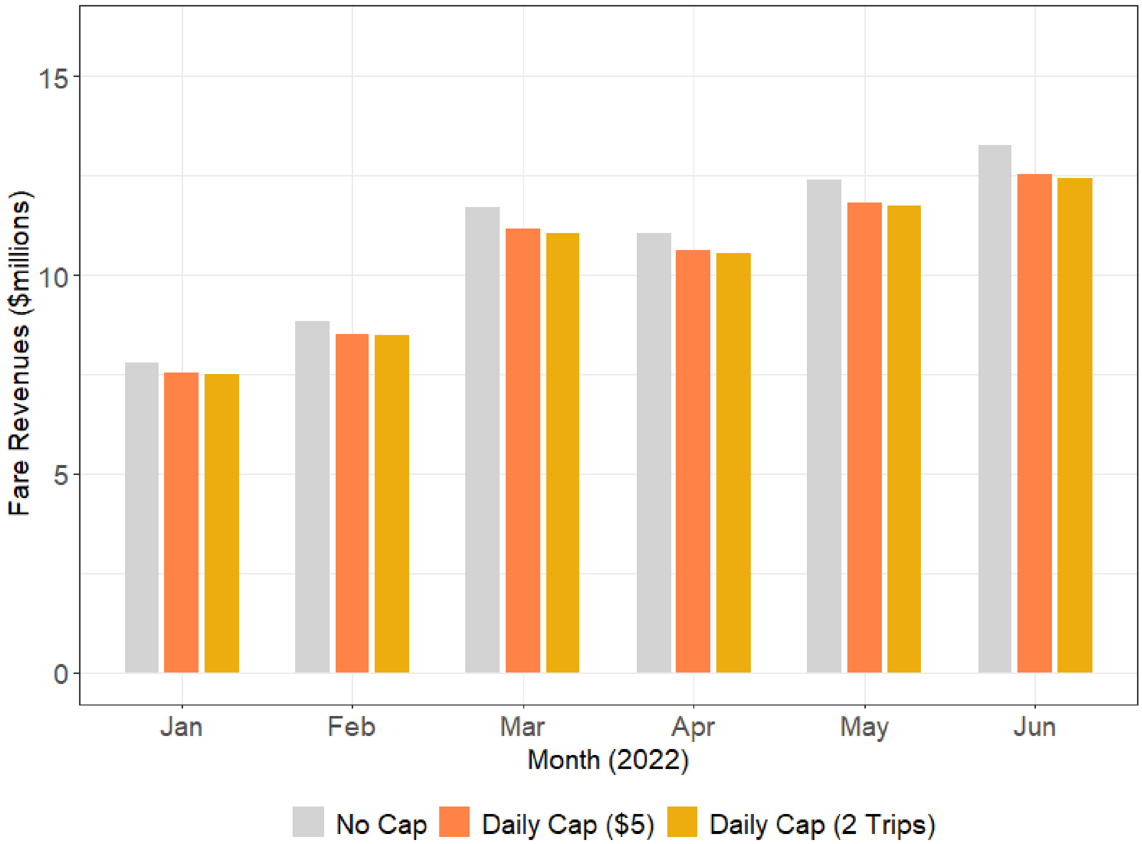

Figure 6-8: Fare Revenues under Different Daily Caps

6.4.3 Geospatial Impacts

While exploring the benefits of fare capping to understand how the benefits are distributed, there are multiple ways to group the beneficiaries. The number of riders who would benefit, and the frequency they would do so, is captured in Figures 6-5 and 6-6. However, the benefits of single-day fare capping can also be considered geospatially, allowing identification of the Chicago neighborhoods that would be most affected. Chicago is known for racial and economic stratification, with higher-income and whiter residents tending to live to the north of the Loop and lower-income and minority populations generally clustered south and west. This can be seen in the racial dot map of Chicago shown in Figure 6-9 [66]. This figure uses data from the 2020 US Census to illustrate where residents of different races live within the Chicago city limits, with each dot representing 100 residents. While race and income information is not known for all Ventra accounts, the locations at which each transit journey begins is known. Given the stark demographic differences across Chicago, geographic information on transit journey locations can be used as a proxy to evaluate potential differential impacts across race. It is important to consider whether single-day fare capping would predominately benefit certain neighborhoods, and by extension certain demographic groups, more dramatically than others.

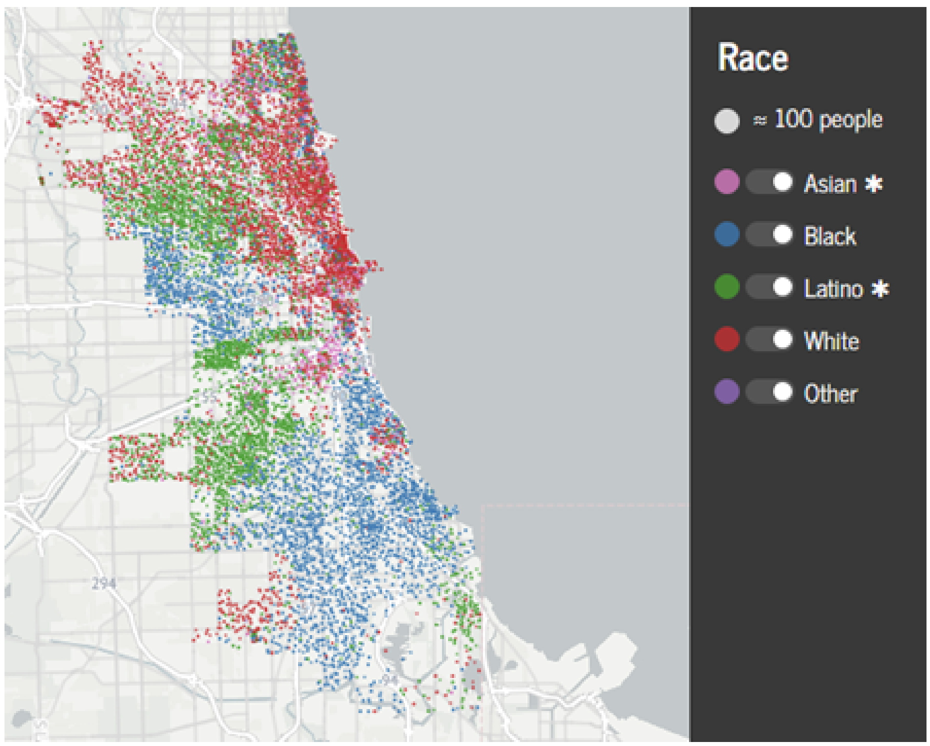

Figure 6-9: Racial Dot Map of Chicago from 2020 Census

The AFC records used to analyze ridership and fare product use in Chicago include information on the bus stop or train station at which a fare card was used. AFC records include information on whether each tap is the start of a new journey or just a leg in an ongoing trip. Combining this information with the calculations on Ventra cards reaching a single-day fare cap supports a mapping of the locations frequented by capped riders. Figure 6-10 shows a map of Chicago that illustrates the journey origins for capped Ventra accounts. The cap used in Figure 6-10 is a two-trip cap, not a fare cap set at a specific dollar amount. Even so, most of the riders who would reach the cap and benefit are taking trips in the northern parts of the city, with many riders also benefitting along the Pink and Orange lines. However, the South Side and sections of the West Side see many fewer trips that would be impacted by a single-day cap. This is the case even though a two-trip cap is used, negating any differences due to whether the rider is traveling on rail or bus. A natural follow-up question is why this is the case.

Figure 6-11 illustrates a likely reason for this discrepancy, and gets to the heart of the issue. As detailed previously in this chapter, a fare cap can only be applied if all travel can be ascribed to a specific traveler. An identifiable transit card is necessary, which in the case of Chicago is assumed to mean a corresponding Ventra identity. Figure 6-11 shows, in gray, the stations where a higher percentage of travelers would reach a two-trip daily cap. It also shows, in purple, the stations and stops where more than 30% of trips are not paid for with a Ventra account. At these locations many more riders are paying with cash or limited-use paper tickets, making the trips ineligible for any benefits from fare capping. The areas with lower Ventra account usage are clustered on the South and West sides in historically disadvantaged communities. Improving Ventra access in these areas would be necessary to ensure any fare capping policies could see a more equitable implementation. Further Ventra uptake in these neighborhoods would be a useful goal in and of itself, even without fare capping. There are benefits provided by Ventra that are not accessible to these riders. Ventra adtoption could also improve the ridership data available to the CTA for planning purposes.

The geographic differences in fare capping benefits are demonstrated at the aggregate level, but Figure 6-12 shows these inequalities are less pronounced among reduced fare transit users. Nevertheless, even among this subset of riders, the northern neighborhoods continue to show the densest benefits from fare capping although

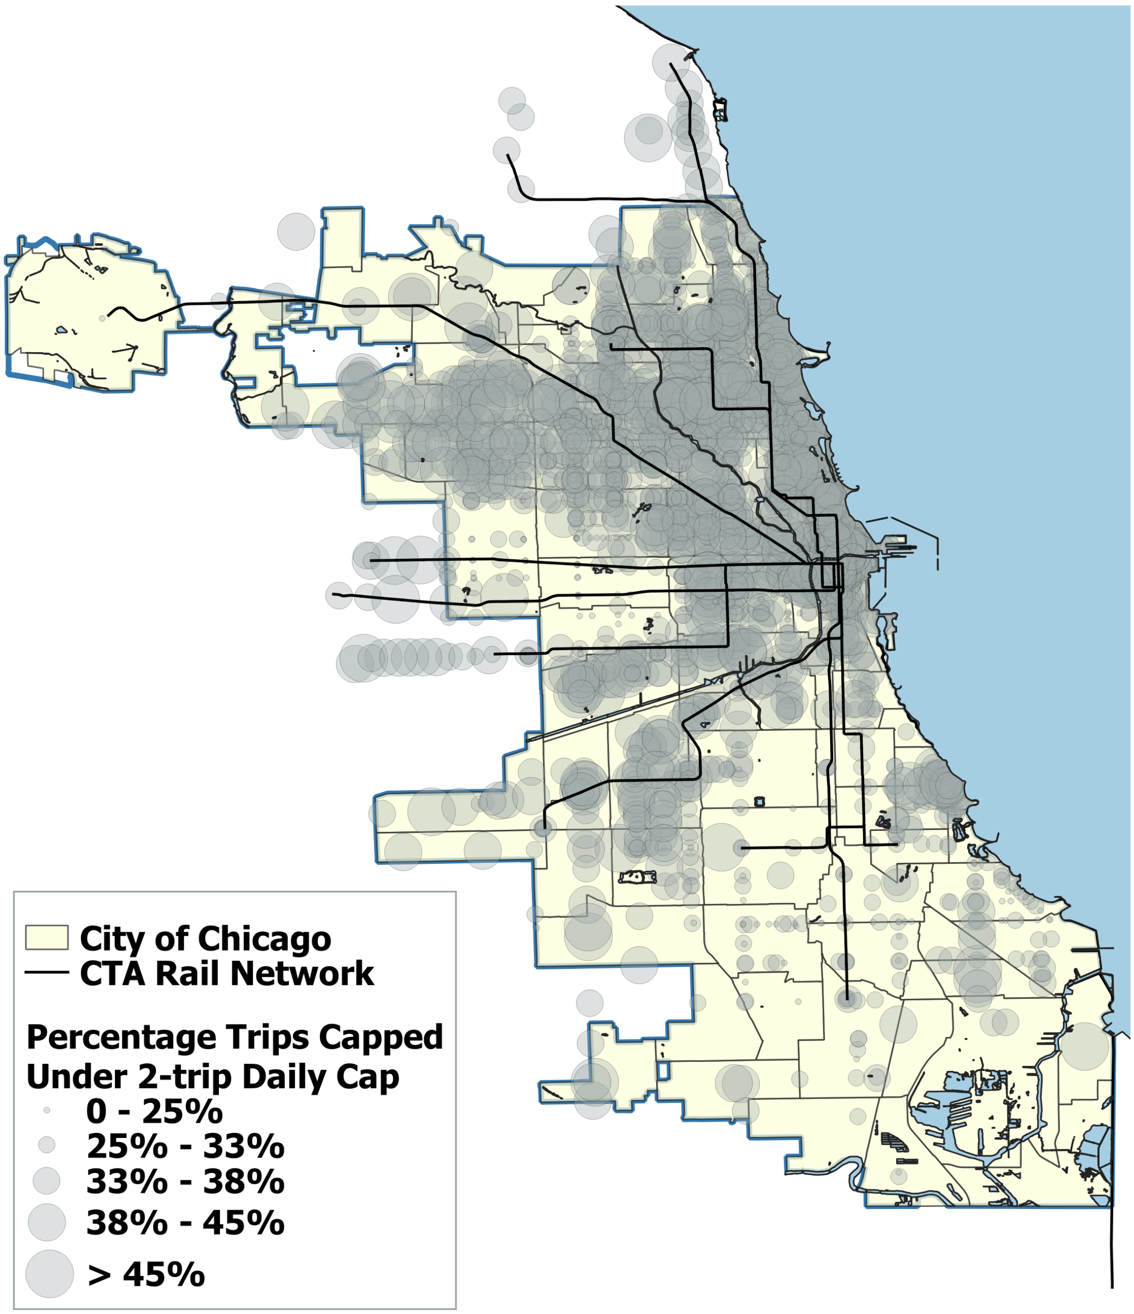

Figure 6-10: Stop Locations Used by All Riders Reaching Fare Cap

many trips to the south and west would also benefit. One key takeaway from this map is that there are many more reduced fare riders using Ventra account-based travel in disadvantaged neighborhoods compared to full-fare riders. This suggests that reduced-fare riders are much more likely to use a Ventra card regardless of where in the city they travel. The reasons for this are intuitive; Ventra cards confer the

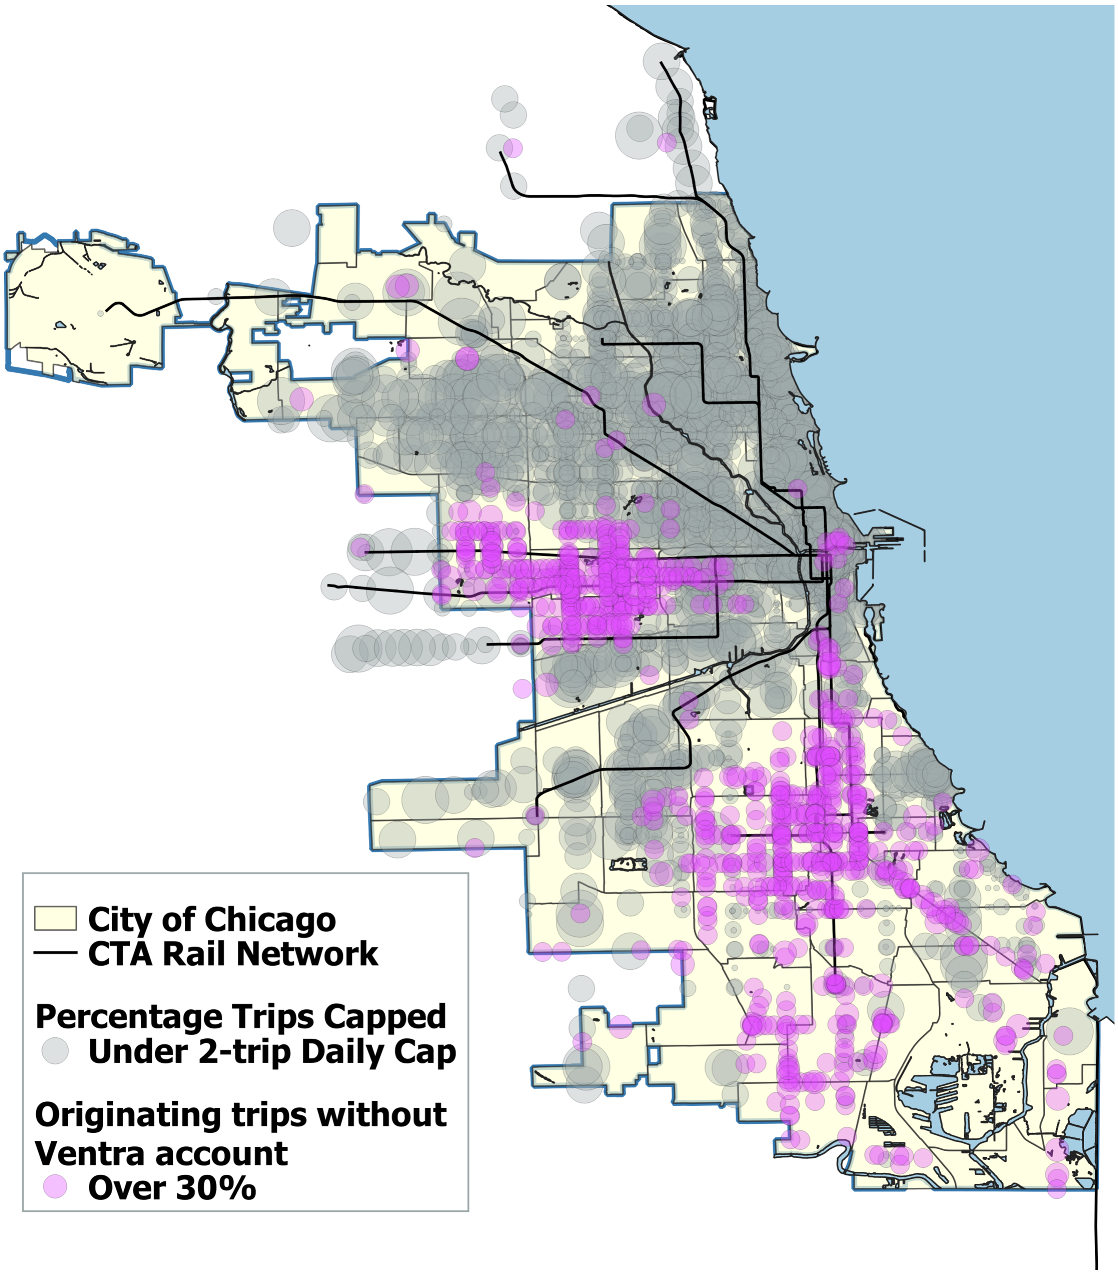

Figure 6-11: Fare Capping and Non-Ventra Payment Locations

ability to ride the system for lower costs. To access this discount, one must have a Ventra card. Further, a much higher percentage of all trips taken by reduced-fare riders would benefit from a two-trip single-day fare cap than the percentage of full-fare riders that would be affected. The legends shown in Figures 6-10 and 6-12 show that the magnitude of riders benefitting varies significantly. Many rail stations and bus stops see more than 50% of trips by reduced-fare riders that would be impacted by a two-trip fare cap. This map highlights the potential of a fare cap to provide benefits to all riders, including those for whom the 1-day pass multiple is highest. Many reduced-fare riders are traveling using pay-per-use single-ride tickets across the entirety of the CTA's service area because it is less expensive to take three or four transit trips than to buy a 1-day pass. Even without fare capping, reduced fare riders deserve pass products that are priced at a reduced level. They may be a small portion of all CTA users, but incentivizing pass usage and transit travel is just as important.

6.4.4 Discussion

This chapter considers questions around one potential approach to fare capping in detail. The implementation evaluated would only apply fare caps to single-day travel on the CTA. A fare capping implementation that only extends to single-day travel was selected because it leaves maintains incentives for pass adoption among frequent riders. After selecting a fare capping implementation strategy, the riders who might be induced to take additional trips under a single-day cap are identified, conservative estimates of fare revenue loss associated with such a switch are generated and found to be small, and the distribution of benefits from fare capping across ridership are illustrated.

From this exercise, three key findings are identified. First, should fare capping be implemented, it should be done a per-ride basis as opposed to a set dollar value. The fare differences between travel modes and rider classes complicate pricing, but focusing on capping at a given number of trips would ensure the benefits are spread among all riders. Second, Ventra adoption is much lower in certain neighborhoods within Chicago. The CTA should push to increase Ventra card uptake in these areas, which would ensure riders have access to all of the benefits available through Ventra, and simultaneously provide better travel pattern data for the CTA to use. Finally,

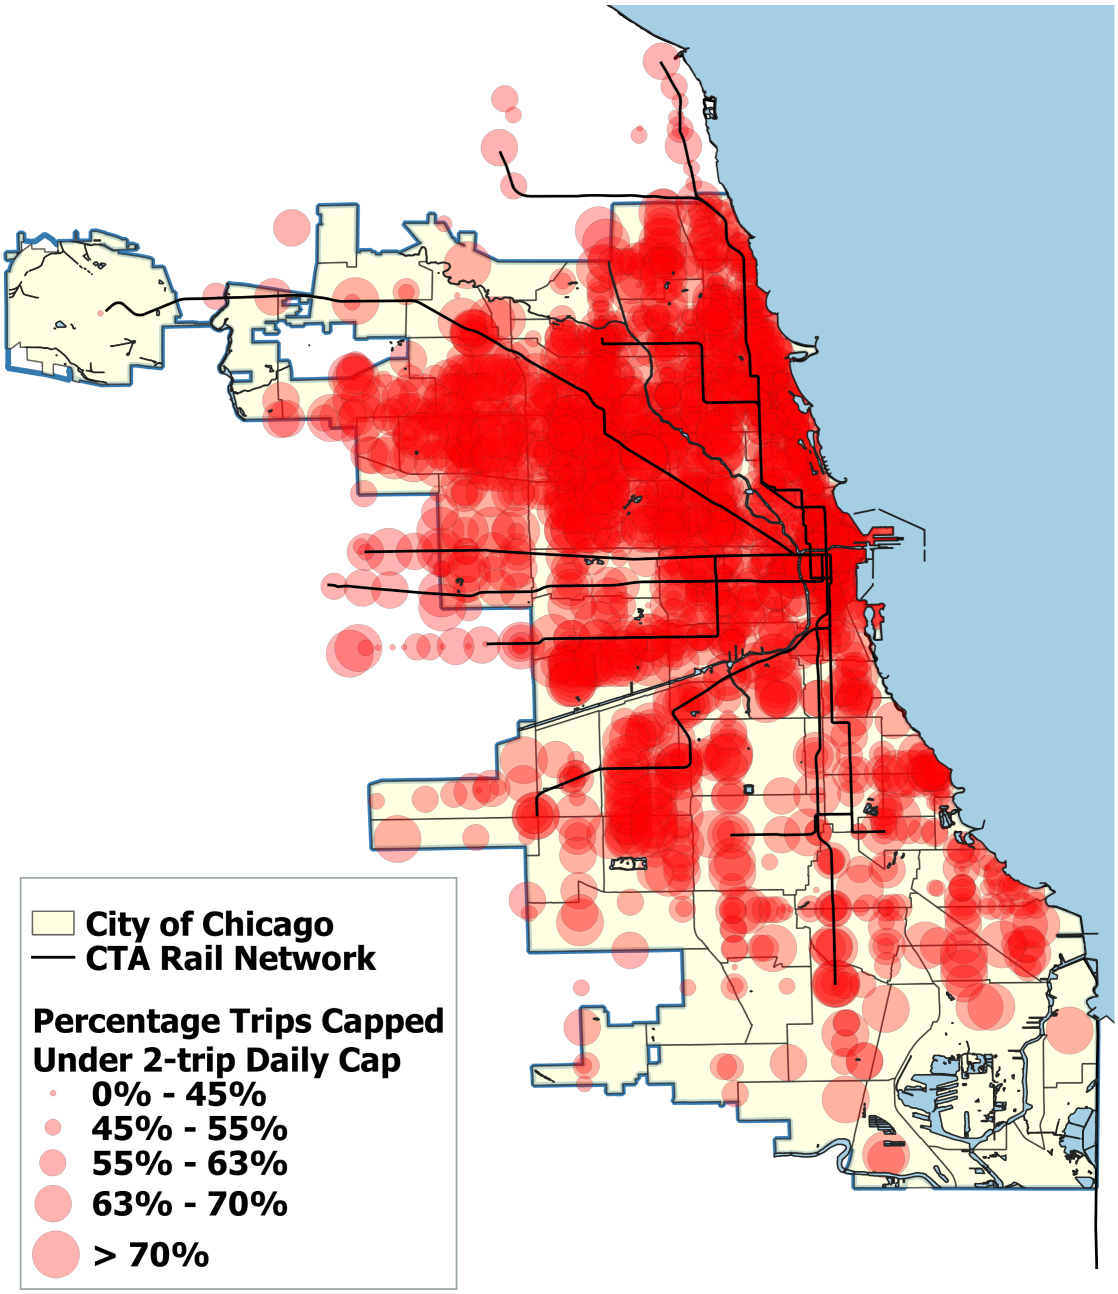

Figure 6-12: Stop Locations Used by Reduced-Fare Riders Reaching Fare Cap

single-day fare capping with the cap set at a specific number of trips would negate the need to offer existing passes at multiple price points. Nevertheless, if single-day fare capping is not a strategy the CTA is interested in, discounted 1-day, 3-day, and 7-day passes for reduced-fare riders must be considered. These riders are traveling on pay-per-use at much higher rates because they do not have access to short-duration reduced-fare passes. The CTA reduced pass prices to incentivize their use; these incentives should be extended to all riders.

Chapter 7

Regional Fare Integration

7.1 Introduction

Earlier chapters explored fare products and fare equity; this chapter considers another important aspect of fares: multi-agency integration. Multi-agency fare integration can provide benefits to both riders and transit agencies. In metropolitan areas with more than one public transportation provider, differing fare structures and fare products can discourage journeys that traverse multiple agencies. Riders can become captive of a single system, even if improved journeys would be possible by utilizing multiple agencies. The historical position of agencies may have been that they had more to lose than gain, in terms of lower ridership and decreased revenues, by cooperating with other providers. However, our post-pandemic transit environment provides an impetus to reevaluate this position. Transit ridership decreased substantially due to COVID-19 stay-home directives and public health emergency orders. Although those orders are no longer in effect, and in many aspects society has returned to its pre-pandemic patterns, transit use remains well below pre-COVID levels. Agencies that had hoped to weather the storm of COVID and wait until riders returned are still waiting. Mass transit providers hope to grow the size of the transit-riding public;

simply defending against the defection of any existing riders is not enough for today's agencies. Transit providers looking to grow ridership through collaboration stand to increase boardings as transit users end up with improved accessibility options. The benefits to riders from such arrangements should be apparent; the chance to access more transit at a given price is a benefit in itself. More efficient trips are available, as are more travel options with fewer constraints.

This chapter presents a case study exploring the opportunities for closer regional integration in fares between separate public transportation providers in the Chicago metropolitan area. The three regional agencies provide urban rapid transit (the CTA), regional commuter rail (Metra), and bus service in suburban areas (Pace). All three of these agencies are under the oversight of the Regional Transportation Authority ('RTA'), which apportions funding and guides region-wide policy in addition to conducting oversight.

Over the past several years, the CTA and Pace have been working towards closer integration of fares so that riders can traverse both systems seamlessly. The CTA pass pricing changes made permanent in fall 2021 (discussed in chapter 5) applied to both joint passes as well as CTA-only products. At that point the CTA and Pace offered two joint passes: a 7-day pass and a 30-day pass. The CTA and Pace both also offered a separate 7-day option valid only on their own networks. The CTA only offers one 30-day pass option, and it is jointly valid on Pace. Conversely, Pace offers a 30-day pass that is valid only on its network, and that option is $60 as opposed to $75 for the joint-agency 30-day pass. The CTA's 1-day and 3-day pass offerings were traditionally only valid for travel on the CTA system. As of writing, those 1-day and 3-day passes have just been made valid on Pace, alongside the consolidation of the 7-day pass into a single jointly-valid pass [65]. This marks substantial progress in fare integration as all CTA passes are now recognized on Pace, and nearly all Pace passes are recognized by the CTA.

Integrated fares between the CTA and Metra have been less common. Until July 2022 there was an optional add-on product available to commuter rail monthly pass holders that provided CTA travel with restrictions. This add-on was called the Link-Up and it offered CTA travel on weekdays during the morning and evening peak, alongside unlimited Pace travel. Additionally, Metra riders could purchase the Pace PlusBus pass for unlimited Pace travel with no CTA validity. These two passes cost $55 and $30, respectively. Both add-on options were replaced in July 2022 by the Regional Connect Pass, which provided Metra monthly pass users unlimited travel across the CTA and Pace on all days and at all times for $30. This change was made alongside deep discounts to the Metra monthly pass, which was priced at an even $100 regardless of the number of zones traveled. To illustrate the deep discount this represents for some riders, previously the most expensive pass had been $275. This equates to savings for these Metra monthly pass users of up to 64%.

All three transit agencies use the same vendor to track their fare product sales and rider boardings, but the agencies do not all record transactions uniformly. CTA and Pace generate AFC records when a rider uses a Ventra card or limited-use ticket to tap onto a bus at the farebox, or onto rapid rail by tapping at a turnstile. Metra does not use validators to charge for travel; instead riders can purchase trips from conductors onboard or through the Ventra smartphone app. Certain Metra products, such as the Regional Connect Pass add-on, are only available to purchase through the smartphone app. Correspondingly, smartphone ticket sales and usage constitute the large majority of Metra travel. Riders launch their ticket and show it to the conductor; this process creates a usage record in the AFC database. That record is associated with a time and often a latitude and longitude, although no station, vehicle, or ticket validating machine are recorded. This means time of travel, fare product, and use frequency information is available for commuter rail riders, but origins, destinations, and transfer points are not specifically identified.

Recent changes have been made to pricing for products with single-agency validity, as well as to add-on products that provide travel on regional partners. Changes to the add-on offerings were not limited to purchase price; usage restrictions were also loosened, providing more freedom to travel. In this chapter the impacts of these changes are first observed in terms of ridership, revenue, and customer fare product choice. Second, fare elasticities are calculated for multi-agency and commuter-rail riders post-COVID. Third, the potential market size for additional integrated fare products is discussed, including a prioritization of existing fare product users to target. The results of this case study illustrate the potential gains for both rider and agency from regional collaboration and suggest clear avenues to further improve rider accessibility.

7.2 Research Goals

The purpose of the case study presented in this chapter is to explore the benefits that might be realized should Chicago-area public transportation providers work to further improve integrated fare offerings. The specific research questions explored in this chapter are:

- **How have riders responded to recent improvements in multi-agency fare integration?** The replacement of the Link-Up and PlusBus passes with the Regional Connect Pass ('RCP') in July 2022 represents a major advance in regional integration. The RCP provides a single add-on product conferring unlimited travel across all of three transit providers. These improvements were coupled with lower prices.

- **What is the potential market size for multi-agency fare product?** Each agency uses the Ventra platform for ticket sales and fare payment, but the data retained for analysis by each agency is not uniform. Differences across transit agencies are due to both different policy choices and differing fare validation needs for different transit services.

- **Which fare products should be prioritized for integration?** The Regional Connect Pass is an undeniable improvement over the Link-Up pass, but it is only available to Metra riders that purchase a monthly commuter rail pass. Continued progress on regional integration will require that riders using other fare products have access to multi-agency travel. Assuming that progress continues incrementally, the next fare products to be targeted for improvement need to be identified.

Taken together, these three questions provide both an understanding of the benefits already realized through improvements in integration as well as a guide to future strategic decision-making.

7.2.1 Rider Responses to Recent Improvements in Integration

The Regional Connect Pass, introduced in the summer of 2022, furnishes a real-world experiment that can be used to analyze transit use and fare product choice. The Ventra platform contains information on fare product validation used to access all three transit agencies. AFC records are associated with a specific Ventra account, represented using an anonymized identification variable. These AFC records include information on the time and location of the fare tap, recorded transfer points, and the fare product that was used for travel.

From these records, changes in travel patterns at the level of individual riders can be observed. In this case, riders that have used the RCP can be identified. Any changes in CTA travel can be directly detected. Observable changes can include changes in travel frequency, travel regularity, different time of day for trips, or use of new transit networks. Additionally, since CTA travel records include unique identifiers for Ventra accounts, other pass products purchased by these same accounts can be observed over time. The availability of all of these details helps indicate what aspects of the RCP have appealed to users. Is it the lowered price? The freedom to travel at all days and times? Unfettered access to the CTA? Investigating rider travel using available AFC data clarifies which riders benefit from improved integration, how they have responded, and the extent of changes in their public transit use. These questions can then in turn inform future enhancements while also corroborating any successes to date.

7.2.2 Potential Market for Further Fare Integration

Quantifying changes in transit travel and understanding shifts in fare product choice both provide powerful insights on regional fares. Regional integration of transit fares has clearly been improved in Chicago, but the to-date improvements are concentrated among transit riders that purchase a Metra monthly pass. To support extending similar add-on products beyond only monthly pass users, it is beneficial to know what the potential demand might be. Not all Metra monthly pass purchasers have opted to purchase the RCP; in fact, only a small portion of them have. In October, only around 5% of Metra monthly passes chose to add on the RCP. Nevertheless, because of the introduction of the RCP, adoption and usage of an integrated option can be clearly observed and monitored among a clearly defined subset of riders: monthly pass users.

Prior to January 2023, the number of journeys and the volume of each fare product sold were identifiable through Ventra AFC records, but the anonymized transit account identifiers were not uniform between CTA/Pace and Metra travel. This necessitated that a link between the data be identified or, if necessary, constructed. Understanding CTA travel patterns and Metra travel patterns does not provide information on joint-use of the systems. Figure 7-1 illustrates the crux of the issue.

The market size for integrated passes could be small, as in diagram A; it could be large, as in diagram B; or it could fall somewhere in between. Further, the extent of the intersection between CTA riders and Metra riders could plausibly differ based on the rider type, which is partially discernable through the fare products used for travel.

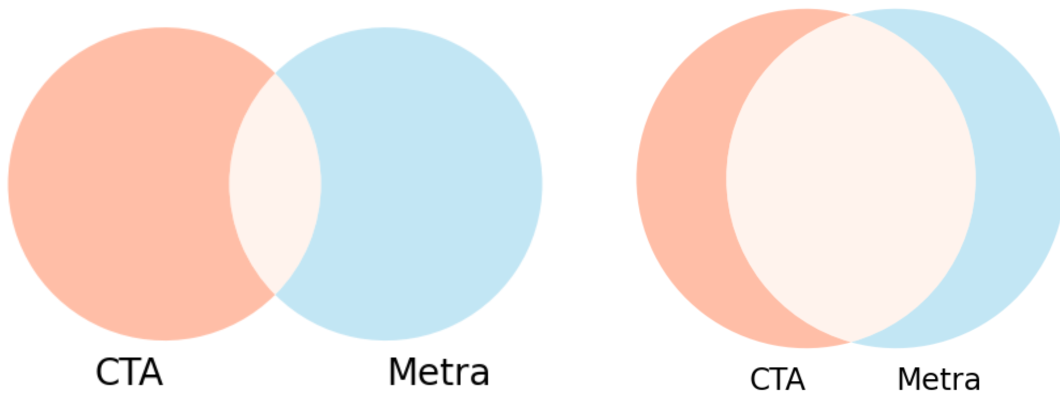

(a) (b)

Figure 7-1: Illustrative Diagram of Multi-System Ridership

A common data field for account identification was found based on the payment method for fare products. Each payment method used to purchase fares has an anonymized, unique identifier for each system. These anonymized identifiers are not consistent across agencies, meaning that a rider purchasing fare products from each agency would have multiple identifiers corresponding to the same payment method. This is the legacy of decisions made at each agency to use separate identifiers for transit travel, possibly to ensure that while data was available through a common vendor, the analysis of each agency's data would remain decoupled and proprietary. In mid-2022, the agencies began populating a new field that allows the various anonymized payment method identifiers to be linked to each other; a visual representation is available in Figure 7-2.

A desire to tap all potential sources of ridership, especially given that travel volumes have still not fully recovered post-COVID, has likely encouraged more collabora-

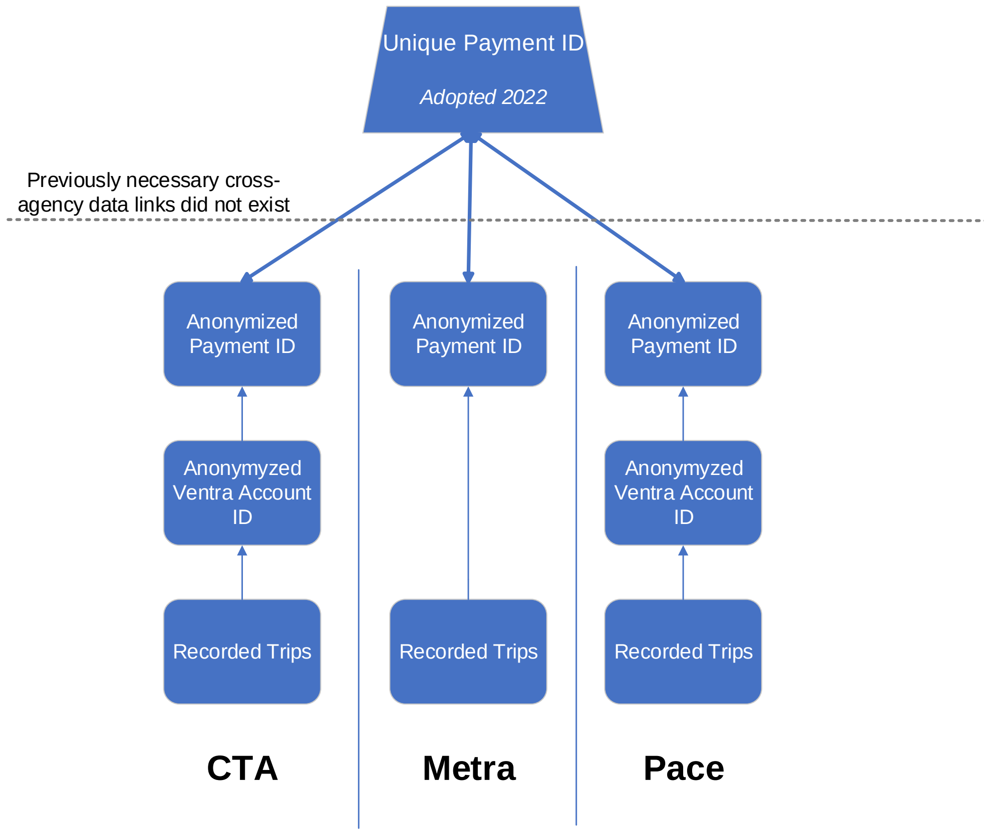

Figure 7-2: Schematic of Cross-Agency AFC Record Keeping

tive approaches among the agencies. A rising tide lifts all ships, and more individuals using transit to move from one point to another presents each agency additional opportunities to increase boardings by offering complementary services.

Having constructed the ability to link AFC records across all three of the regional public transport providers, understanding the current market-size for multi-agency products is central to determining how many multi-agency riders exist. Once these multi-agency riders are identified, the fare products used for travel on each transit system, the frequency of travel on each system, and the number of travel days where a single system is used all orient decision-makers as they weigh the risks and benefits of expanded fare integration. The benefits from increased integration to riders are obvious; ridership gains to each agency should be the major motivation to agencies.

There are no drawbacks to riders from more multi-agency travel opportunities.

7.2.3 Fare Products to Prioritize

The case study presented in this chapter first analyzes travel responses to already introduced improvements to regional travel; it then assesses the potential market-size for continued fare integration. The next step is to assess the fare product choices of existing multi-agency travelers to prioritize add-on products that provide multi-agency access. Full integration is not considered an option in the immediate future; improved integration is likely to first take the form of fare product add-ons, exemplified by the Regional Connect Pass, or new products that are priced at a premium in exchange for cross-agency travel.

Some of the fare products offered by each agency share commonalities, but there are also pass products that are unique to specific agencies. The CTA, Pace, and Metra all offer a monthly pass option, and a pay-per-use option, but only Metra offers a ticket that is valid for ten rides. Given the number of fare products offered, and the variations in those offerings, the question of how to combine add-on offerings becomes pertinent. A Metra rider with a monthly pass can purchase the RCP, valid for CTA and Pace travel during the same month. That same Metra rider cannot currently purchase any other CTA fare product at a discounted price. That is to say that current add-on integration is only an option for alike long-period passes. But where like passes do not exist, what is the solution to encourage integration? Should all riders be able to purchase any fare product from another transportation provider at a discount, introducing the possibility of mixing and matching fare products across agencies? Should agencies that do not have a corollary to a rider's original fare product introduce an equivalent offering? There are three ways to structure fare integration which are illustrated below using the Metra 10-ride ticket as an example:

- Should the CTA be offering a discounted 10-ride ticket to Metra riders utilizing the Metra 10-ride ticket?

- Should Metra 10-ride ticket users not have any options for CTA travel using a fare integration add-on?

- Or should 10-ride ticket users be able to select the CTA fare product of their choice as an add-on at a discounted price?

How the agencies choose to approach these questions are important. Given the permutations of potential fare product combinations across all three agencies, prioritizing fare products for integration can make the burden smaller. Determining in which order to consider fare integrations helps create a path for continued integration and pinpoints the fare products that hold the most potential to riders and transit agencies. A priority ranking can guide not only integration decisions, but also where further research would be conducive. Pilot programs or customer surveys may help resolve the dilemma provided by the different fare product designs at the three transport agencies; a ranking of the best opportunities for integration can indicate transition points between high-potential, unknown potential, and low potential integration options.

7.3 Methodology

The methodological approach used for this case study is set out in this section, which is organized into three subsections roughly corresponding to each of the research questions previously listed. The methodology employed by this analysis relies on AFC travel records at the three Chicago-area transit providers. Identification fields are first established at each agency; these identifiers are used to observe and analyze how customers responded to previous improvements in fare integration. The insights gleaned are then extended to all riders traveling jointly on multiple transit systems, who are identified using AFC records available from the Ventra platform.

7.3.1 Evaluating Travel Changes on Existing Integrated Products

The transaction records available in Ventra provide detailed data on individual trip-making and fare payment. Unique identifiers for each account that travels on the CTA are available. Unique within-agency identifiers are also available at all three agencies at the level of payments for fare product purchases on the Ventra smartphone app. All identification fields are anonymized; their primary benefit to this analysis is to track changes over time.

As the effects of improved multi-agency fare integration constitute the primary analytical question of this case study, exploring previous improvements restricts the cohort of relevant travelers to those able to take advantage of those improvements. Therefore, this analysis narrowly defines the riders to evaluate as those that purchased the Regional Connect Pass at least once in the five months following its introduction. Riders who purchase the RCP are by definition multi-system riders, as only Metra monthly pass purchasers are eligible for the RCP add-on. After identifying the accounts that purchased the RCP add-on in at least one month, other fare payment decisions can be observed across time. The fare products from which RCP users have shifted can be used disentangle and quantify the benefits of an integrated offering.

7.3.2 Calculating Post-COVID Fare Elasticities for Integrated Products

Elasticity is an economic concept used to measure the change in demand for a product as the price of that product changes. When applied to transit, the product demanded is transit journeys and the demand level is quantified as trip volumes. The cost for that product is the transit fare. Application of elasticities to transit fares is well-established with a detailed literature developed over many years. The relevance of elasticity measurements to transit fare changes is very intuitive; the higher the cost to travel, the fewer journeys people will make. Discretionary trips may be foregone as the costs overtake any benefits. Transit substitutes become more competitive. Conversely, lower transit fares act to encourage additional transit use by reducing the cost burden and increasing overall competitiveness.

Modeling fare elasticities requires consistent units of measurement for both travel costs and trip volumes. These units of measurement can be defined in various way. The most common and straightforward technique evaluates elasticities through pay-per-use products. Periodic fare increases have a long history in transit, and they provide for direct measurement. Following a clear, measurable, distinct change in pay-per-use fares, assuming the rest of the transit environment remains constant, all corresponding shifts in travel volumes are measured. These shifts are attributable to the new higher cost of travel in aggregate, and the magnitude of travel volume changes are analyzed. Using changes in single-ride fares and pay-per-use travel volumes as the primary inputs, the rider fare-elasticity is estimated.

Given that fare elasticity measurements require a change in fare prices, fare elasticity is most commonly estimated in the time period surrounding a fare restructuring. In the long history of US mass transportation, such fare changes have normally only occurred in one direction – upward. The COVID-19 pandemic scrambled the logic on many traditional perspectives around mass transit, only some of which have been detailed within this thesis. One of the notions that has been reconsidered is a need for continual increases in fares in light of the steep declines in transit use. As a matter of fact, in Chicago there have been actual, observable decreases in the cost of some fare products. Chapter 5 discussed CTA pass price reductions in detail. Metra's deeply discounted monthly pass price, and the significant changes to fare integration represented by the RCP, are two additional instances of transit agencies actually reducing the cost to ride.

This chapter estimates post-COVID transit fare elasticities for commuter rail only riders as well as elasticity for multi-agency access as calculated using joint CTA/Metra riders. The detailed AFC data available is used to measure rider responses following the July 2022 pricing changes to two specific fare products: the Metra monthly pass and the Regional Connect Pass. The price change for monthly passes experienced by Metra users is not uniform. The monthly pass was set at a single price regardless of the number of zones for which it is used. Previously each of Metra's ten defined zones had a different monthly pass price. All customers that purchased a Metra monthly pass following the price reductions in July 2022 is included in the analysis, regardless of their prior fare product choice. These riders experienced a change in their fare per trip, and their corresponding travel changes are germane to ensure accurate elasticity estimation. Table 7.1 shows the change in prices for representative Metra fare products in July 2022. To measure the elasticity for post-COVID commuter rail, only riders that have experienced the change in fares are included, which are monthly pass users. For single-ride and 10-ride tickets pricing, differentiation for each of the ten zones remains. Any travel pattern shifts observed for the other riders are not being driven by fare changes. The other riders are excluded from elasticity estimation to avoid spurious impacts on the results.

	All Zones		**Up to 3 Zones**	
Metra Product	June 2022	July 2022	June 2022	July 2022
Monthly Pass	$275.50	$100.00	$159.50	$100.00
Ten-Ride	$90.25	$90.25	$52.25	$52.25
Day Pass	$10.00	$10.00	$6.00	$6.00
One-Way	$9.50	$9.50	$5.50	$5.50
CTA Add-On	$55.00	$30.00	$55.00	$30.00

Table 7.1: July 2022 Metra Fare Product Price Changes

For similar reasons, fare elasticity for multi-agency travelers is estimated using AFC data for all riders that purchased the RCP in the months following its introduction regardless of whether or not they had previously purchased the Link-Up pass.

All riders switching to the new fare product are experiencing the fare change it represents. Although the RCP replaced a number of different fare products for different riders, the price per ride can be measured regardless of fare product.

AFC data records for both RCP users traveling on the CTA and for Metra travelers that purchased a monthly pass are individually identifiable. The fare products selected, fares paid, travel patterns, and additional potential control variables were assembled into a panel dataset with both rider and trip information.

The regression model base formulation used to estimate elasticities is as follows:

$$log(monthlytrips)_{it} = \alpha_i + x_{it}\beta + \gamma log(avg.monthlyfare)_t + \epsilon_{it} \tag{7.1}$$

wherein

- $i = 1, ..., N$ representing all Ventra accounts
- $t = 1, ..., T$ representing all time periods, in this case months
- α_i = Ventra account-specific constant terms
- x_it = Control variables
- β_it = Coefficient estimates for all control variables
- γ = Estimated elasticity
- ϵ_{it} = Error term

In a bit more detail, the dependent variable is the logarithm of the number of monthly transit trips on the relevant network. RCP elasticity estimation analyzes CTA travel and Metra trips are used for the Metra monthly pass estimates. α_i represents a dummy variable that is specific to each Ventra account. It is used to control for any unexplained heterogeneity across riders and to prevent bias in the calculation of population-wide elasticities. The average fare per trip is used as the "price" variable; it is calculated on a monthly basis for each rider and the logarithm of this value is calculated. The coefficient estimate (γ) represents the overall estimated fare elasticity for the entire population of riders. Various control variables were also considered; the control variables explored will be discussed more completely in the Findings section of this chapter.

Taking the logarithm of both the demand variable (monthly trips) and the cost variable (fare per journey) allows for a direct interpretation of the estimated coefficients as the percentage change in demand given a percentage change in cost. This is the exact definition of elasticity. Transformation of the data is therefore a straightforward and beneficial tactic for elasticity estimation.

Previous fare elasticity findings have often been in the range of -0.3 to -0.6. This range is useful to provide some context on the results of model estimation in the post-COVID environment. Table 7.2 shows a variety of estimated elasticities cited in a handful of previous studies, including studies that reported results separately by transit mode. While this table is not an exhaustive list, it does serve as a useful reference point. Two additional benefits of the model formulated here is that they can be used to compare the elasticity of multi-agency riders against those of commuter rail only riders.

7.3.3 Clarifying Potential Market Size for Multi-System Riders

Building from the observed travel changes among Regional Connect Pass users alongside estimated post-COVID elasticities, the next step in advancing fare integration across the Chicago region is identifying what market exists for multi-system products outside of monthly pass purchasers.

Using the data at hand, the potential market size for multi-system riders will be approximated using rider segments defined by their multi-system transit use and

Transit Mode	Short Run v. Long Run	Elasticity Estimate	Source
All		-0.33	Cummings et. al. (1989) [67]
All		-0.30	Webster & Bly (1981) [68]
Bus	Short	-0.38 to -0.42	Balcombe et al. (2004) [24]
Bus	Short	-0.40	Dargay & Hanly (2002) [69]
Bus	Medium	-0.56	Balcombe et al. (2004) [24]
Bus	Long	-1.00	Balcombe et al. (2004) [24]
Bus	Long	-0.90	Dargay & Hanly (2002) [69]
Commuter Rail	Short	-0.37 to -0.6	Balcombe et al. (2004) [24]
Metro / Light Rail	Short	-0.29 to -0.30	Balcombe et al. (2004) [24]
Metro / Light Rail	Long	-0.6 to -0.65	Balcombe et al. (2004) [24]
Subway / Rail	Short	-0.33 to -0.46	Balcombe et al. (2004) [24]
Subway / Rail	Long	-0.65	Balcombe et al. (2004) [24]

Table 7.2: Previous Transit Elasticity Estimates

their fare product choice. There are other ways to define the potential market size, which could include geographic trip-making patterns or areas with overlapping transit services. However, this thesis focuses on recorded travel frequencies and observed fare product choices.

Having settled on a definition of potential market size, the newly established link within the Ventra database can be utilized to fully quantify the overall potential market for integrated products. The market size can be broken down at a more granular level, with subsections of the potential market grouped by existing fare product choices. The major advance offered by the recently introduced data link is the opportunity to quantify how many residents of the Chicago area are actually utilizing multiple public transportation providers, and how often they are doing so. The detailed nature of Ventra usage records provides for exploration based on all combinations of fare products offered across the three agencies. To give an example, the number of Metra monthly pass users can be subdivided based on how frequently they travel on the CTA, but they can also be subdivided based on which CTA fare products they choose for that travel. Such distinctions provide for a ranking of the fare products with the riders most likely to take advantage of multi-system, region-wide integration. This insight is valuable as the agencies collaborate to improve joint offerings, and can guide the sequence and extent of those integrations.

7.4 Findings

7.4.1 Travel Patterns of Multi-System Riders

Multi-system riders are the basis for pursuing regional fare integration. Integrated fares can be a tool to better serve existing customers, but they can also be used to make transit a more competitive form of travel and therefore capture additional ridership. To investigate travel patterns of multi-system riders, the riders can be partitioned into three segments. First, there are the riders that travel on multiple transit networks and already use some joint fare product conferring travel by all providers. The second group comprises those riders that travel on multiple systems, but purchase fare products for each system independently of each other. Expanding the extent of integrated fare products benefits both of these rider segments, but the second segment of riders offers a greater potential to increase travel volumes. Third, there are the riders that could be drawn into traveling on an additional system if offered an integrated product. These individuals present a completely untapped market from which to generate travel.

The Regional Connect Pass constitutes a specific fare product that provides improved fare integration between the CTA and Metra. Additionally, as an add-on product it is only available to multi-system travelers. The extensive AFC records available within the Ventra database provide clear identification of the riders that chose to purchase the RCP, and these riders are definitionally multi-system riders. Analysis of this rider cohort can show not only how CTA usage changed among these riders, but also how their fare product choice changed. The customers that have chosen to purchase the RCP are not comprised only of those that previously purchased the Link-Up pass; in fact, it is only a minority of RCP users that previously purchased the Link-Up pass.

Only riders that have purchased a Metra monthly commuter rail pass through the Ventra smartphone app are eligible to purchase the RCP. Concurrent with the introduction of the RCP, the prices of all Metra monthly passes were reduced to a uniform price of $100. Before diving into travel findings for RCP users, it is revealing to observe how all commuter rail riders responded to the price changes. Figure 7-3 shows how all Ventra accounts moved between primary Metra fare products on a monthly basis. The figure clearly shows that in the months leading up to the price reduction, monthly pass users were a small but steady fraction of all Metra riders. Single-ride tickets and 1-day passes served the majority of Metra accounts, with riders traveling on a 10-ride ticket comprising the remainder. Following the deep discounts to Metra's monthly pass, there are small but clearly observable shifts among riders to the monthly pass. Shifting to the monthly pass is exactly what one would predict to happen, although the magnitude of the resulting shifts may be smaller than expected. Nevertheless, riders continue to move into the monthly pass from the time of its price drop in July through November. Figure 7-3 further shows a large chunk of users switching into the inactive category and back each month. These riders predominately choose a 1-day pass or a single-ride ticket when they are active on the system, and their in- and out-flows are consistent over the entire observation period. This does show the inconsistent nature of much commuter rail use. It is possible some of these riders shown as moving into and out of the "unknown" category are sometimes purchasing tickets through the Ventra app, which is captured in the AFC data, and sometimes purchasing tickets on-board or through other channels. However, the vast majority of all Metra ticket sales occur through the smartphone app, and these riders clearly utilize the app, suggesting any switching between purchase channels is likely to only affect a small fraction of riders.

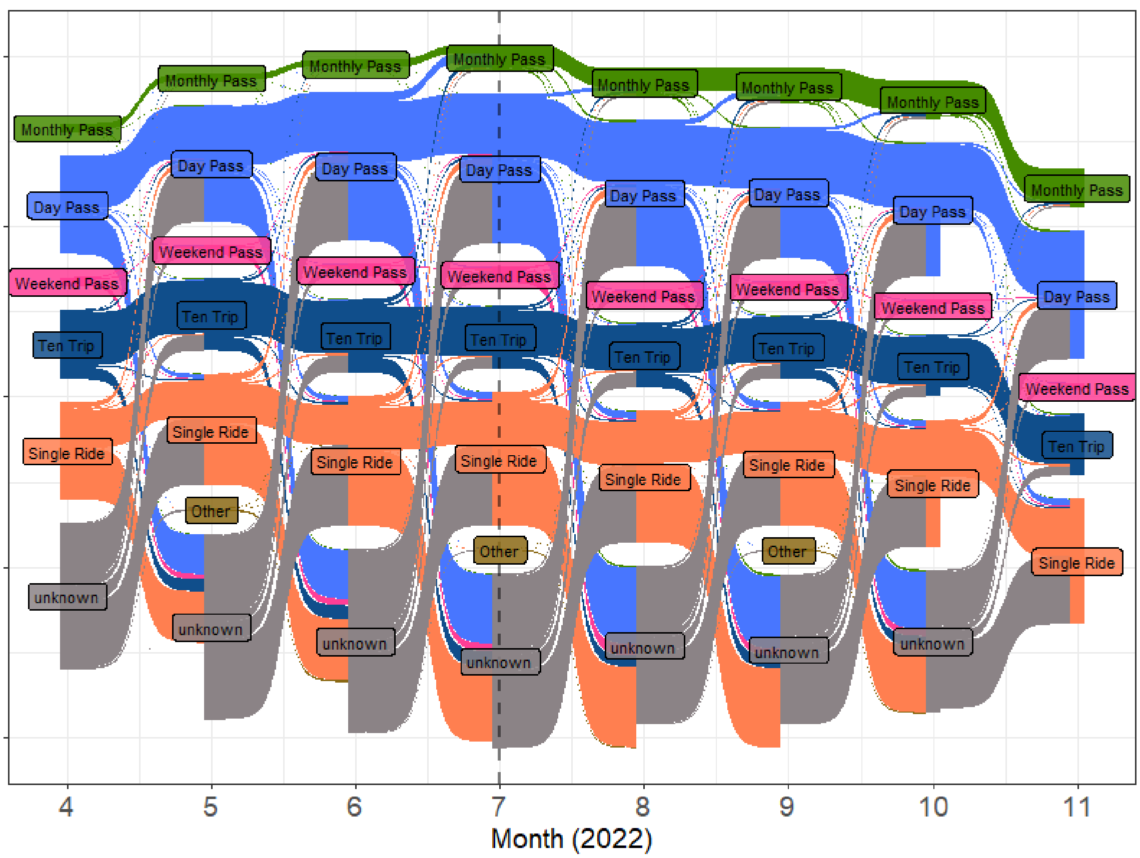

Figure 7-3: Metra Primary Fare Product per Month by Account

The flow of customers between fare products over time illustrates which products are gaining riders from which other products and to which products they are losing customers. The primary fare product choice of each Ventra account can also be used to categorize Metra travel volumes. Figure 7-4 shows how many Metra trips were taken each month, and the fare product used for each of those trips. This same data is presented alongside the number of active Ventra accounts broken down by primary fare product choice in each month. The second chart on the righthand side of the figure displays a stacked bar chart showing only active riders for each month. The juxtaposition presented in this figure helps clarify the scale of ridership offered by each of the fare products. Metra monthly pass users are a small portion of all Metra riders. This is true before and after the price decrease, although there has been a steady increase in the number of monthly pass purchasers since July 2022. Nevertheless, these riders are a large share of all trips taken on Metra. These monthly pass users are high-frequency riders, which is intuitive since they have decided to purchase a monthly pass. Conversely, single-ride tickets are a large number of all Ventra accounts, but a small portion of trips taken each month. Coupled with the shifts between single-ride users and inactive accounts seen in Figure 7-3, single-ride users are much less consistent, as well as less frequent, Metra travelers. These charts do reveal one surprising finding, which is 10-ride usage remained relatively steady in terms of both riders and trips even after the monthly pass price was heavily discounted. The monthly pass and the day pass both offer better options for round-trip travel following the July 2022 price decreases; currently the 10-ride pass is only optimal for regular one-way travel. Nevertheless, many users have remained committed to it over time. The 1-day passes offered by Metra comprise a large number of both trips and active accounts. This is true both before and after the monthly pass saw its price cut to $100.

Figures 7-3 and 7-4 help clarify the potential number of riders that are eligible to purchase the Regional Connect Pass, and how that number has changed over time. The next step is to explore the number of riders that chose to purchase a Regional Connect Pass, and how that impacted their CTA travel over time. Many commuter rail users that are commuting to the Loop or another business district can walk from the Metra station to their destination. And Metra monthly pass holders could be

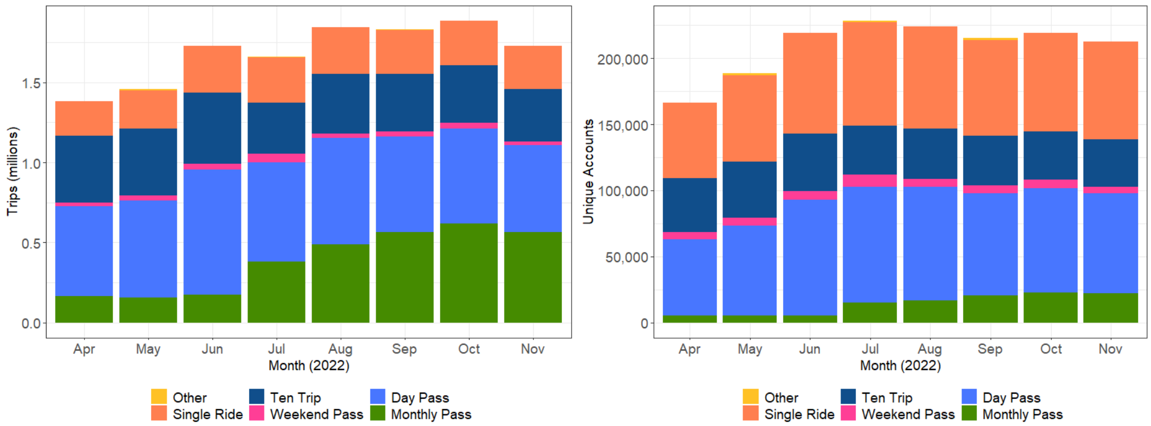

Figure 7-4: Metra Trips and Accounts by Fare Product

considered generally more likely to have a destination that has close proximity to the commuter rail system just by fact that they are high frequency users. The CTA network also offers many connections to Metra stations, providing an alternative. RCP holders have the ability to use the CTA, potentially leading to a large increase in CTA travel among monthly commuter rail riders. These riders still have the option to walk, but on any day when there is inclement weather or the rider has an errand to run, the CTA is right there.

As with the Metra fare product choice changes over time, fare product shifting for CTA travel can be charted for Regional Connect Pass users. The RCP was introduced in July 2022 and replaced the Link-Up pass, providing fewer restrictions at a lower price. Figure 7-5 shows that many riders have shifted to the RCP in the months following its introduction. It also indicates that very few of the riders that have adopted the RCP were previously traveling on the Link-Up pass. The Link-Up pass was always known to serve a small audience, but Figure 7-5 illustrates just how few of its potential riders it was reaching. There has been continual shifting into the RCP among these multi-agency riders, and only a small portion shifted away from the RCP. Two compelling takeaways from are that the RCP actually drew in individuals that previously had no identifiable CTA travel and that single-ride pay-per-use riders have seen the largest number of riders migrate to the RCP. Both of these are positive outcomes for riders and transit agencies. Among the riders that did not have any previously recorded CTA travel, the fact that they are now traveling on the CTA shows clearly that improved integration is expanding the transit market. Perhaps some portion of these riders previously traveled on the CTA system but did not use a Ventra account to do so; even if this were the case, moving riders onto Ventra is another agency benefit. It allows different journeys to be associated with each other, improving data analysis opportunities to improve service and understand the riding public. The continual flows into the RCP over time show that more and more riders are aware of its benefits, which speaks to the strength of the product.

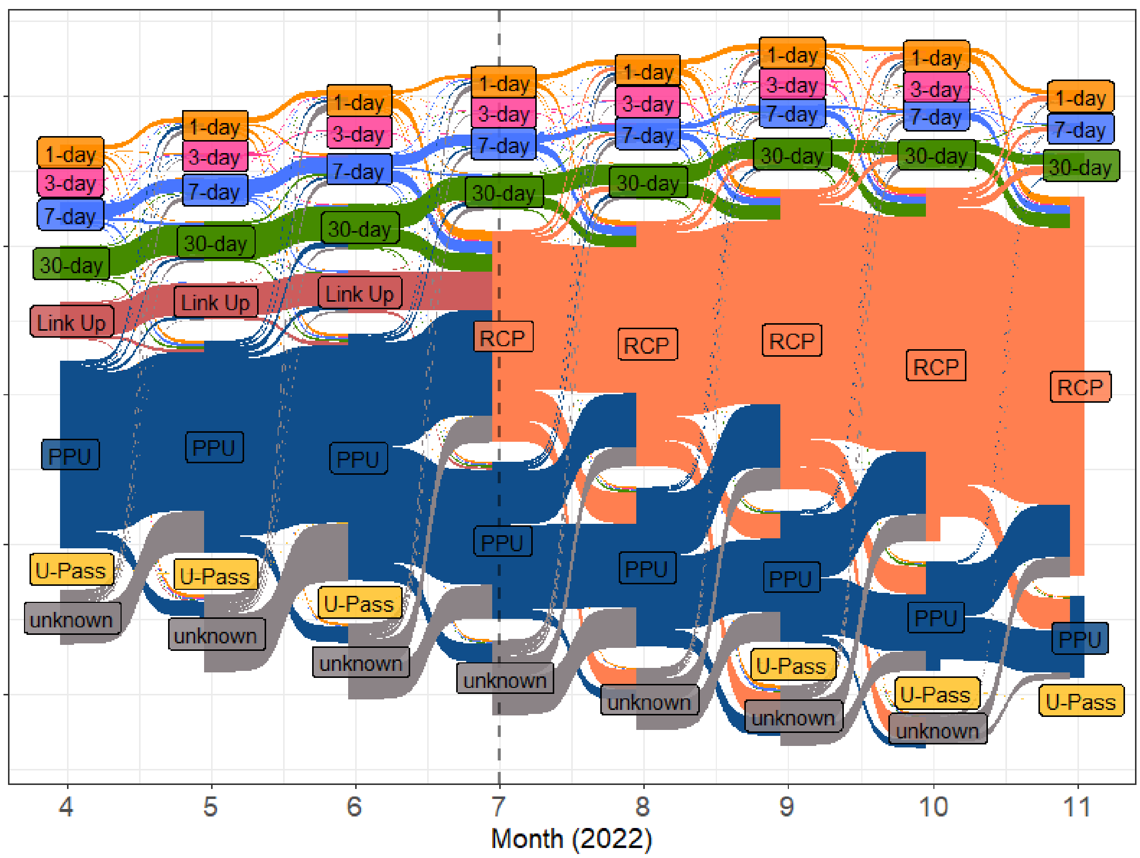

Figure 7-5: CTA Primary Fare Product per Month by Account

The RCP has clearly attracted users away from other products over the months since its introduction. This could be due to its lower price, the elimination of any time period restrictions on usage, or a combination of both. Further exploration of the impacts of the RCP on travel patterns in this section will focus on the cohort of riders displayed in Figure 7-5. This group of Ventra accounts comprise riders jointly traveling on the CTA and on Metra, by definition of their RCP purchase. Observable changes in the travel behaviors of this subset of riders after the introduction of the RCP, and their adoption of it, are indicative of the ridership benefits of fare integration. Additionally, shifts among RCP riders can inform expectations for additional integrated products.

Analyzing the change in trips across three distinct day and time periods can help clarify to what extent removal of usage restrictions drew riders in. The Link-Up pass only provided CTA travel on weekdays during the peak periods. It is worth investigating whether there were any observed changes in the days and times people chose to ride on the CTA following the Link-Up pass's replacement by the RCP. Figure 7-6 indicates that there was little change to the time of travel of RCP users. The proportion of trips taken on weekends remained steady before and after the RCP was introduced. This is also the case for weekday trips taken in off-peak periods and weekday peak trips. The number of CTA journeys taken in each of these three periods increased consistently in the months following the introduction of the RCP before leveling off in the late fall, but the proportional travel in each period was little changed. This suggests that loosened restrictions on CTA use were not the primary driver of RCP uptake, although the loosened restrictions help to serve travelers and draw them to the RCP.

Total travel on the CTA by the subset of multi-system riders that purchased the RCP increased in the months following its introduction, including across all days and times (Figure 7-6). Many riders that had used various CTA fare products, or had no identifiable travel at all, switched to the RCP (Figure 7-5). Figure 7-7 emphasizes another impact of the RCP; the interplay between travel and fare revenues. This

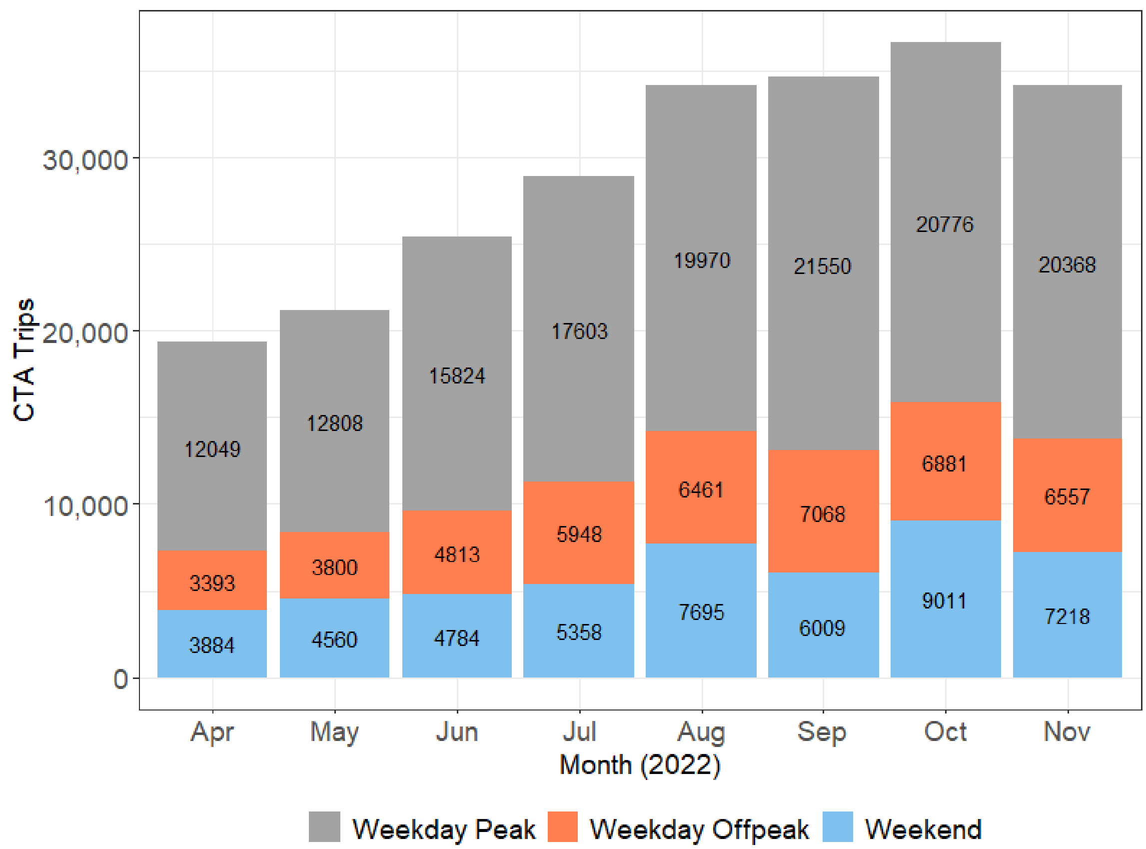

Figure 7-6: Regional Connect Pass Trip Volumes by Time of Day

chart shows total CTA travel by these multi-agency riders from April until November 2022. CTA travel volumes for this consistent set of customers increased noticeably in July 2022 and the months following. The colors in the stacked bar chart represent the CTA fare product on which those trips were taken. The Link-Up was a consistent but small portion of all trips pre-RCP, with single-ride pay-per-use constituting the most travel. There was a major jump in pass trips and a corresponding drop in in pay-per-use travel among these riders concurrent with the RCP's availability in July 2022. By October, the number of trips taken on the RCP outnumbered all trips taken in June, the month immediately preceding the RCP's debut.

Figure 7-7 also displays the CTA fare revenues generated from these trips on a monthly basis. The story told by the fare revenue numbers is somewhat different; fare revenues were highest in June immediately prior to the Regional Connect Pass becoming available. This is true even though June had the third fewest trips in the analysis period. Following the introduction of the RCP in July 2022, the fare revenues collected by the CTA from these multi-agency riders remained relatively steady as travel volumes climbed. The CTA received slightly lower fare revenues than in June 2022, but the collected fare revenues in each of the following months were still above April and May totals. Ultimately, the CTA witnessed clear increases in travel that were proportionally larger than corresponding drops in fare revenue.

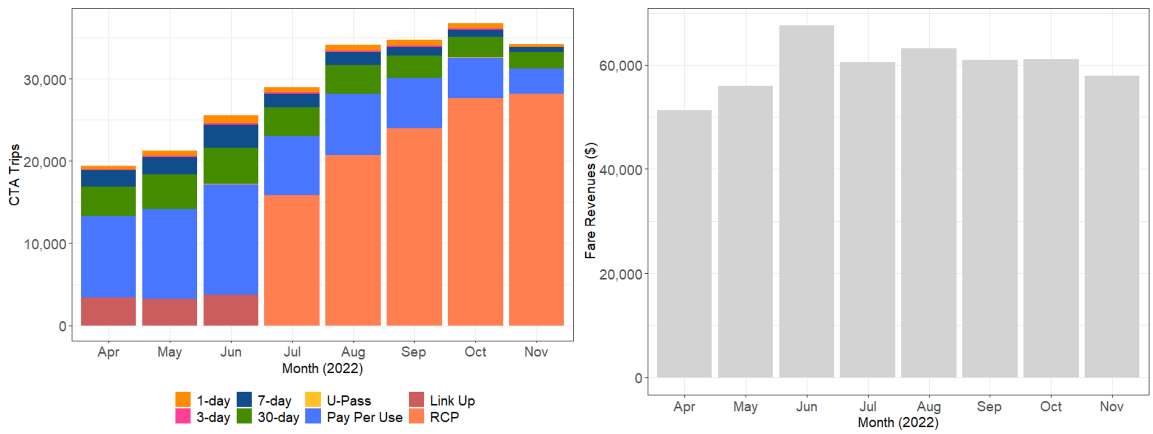

Figure 7-7: CTA RCP User Travel Volumes vs. Fare Revenues

Travel on the CTA network among these RCP users increased at an aggregate level. However, inspection of travel at the aggregate level does not reveal if travel increases occurred uniformly among all riders. Users that chose to purchase the RCP can be segmented in various ways. One particularly useful way to categorize these riders is by the fare product they chose to use for CTA travel before the RCP was available. Segmenting riders in this way is intuitive; these riders chose their fare product based on the frequency and times of day they needed access to CTA service, as well as the price point that was best. Given their choice from among the available fare products, riders can be sorted into three separate groups:

1 **Riders that *always* purchased the Link-Up pass** in every month preceding the introduction of the RCP. As shown in Figure 7-5 these riders all shifted into

the RCP. They are frequent CTA users that have a commuter travel pattern, using the CTA every month and always determining that the Link-Up with its time of day and day of week usage restrictions was an optimal product for them.

2 **Riders that *occasionally* purchased the Link-Up pass** in the months before July 2022. This group of riders predicted each month whether there travel needs would best be served by the Link-Up, and they did not always come to the same conclusion. These discerning customers use the CTA frequently, but not always in ways best served by the Link-Up.

3 **Riders that *never* purchased the Link-Up pass.** Most of these riders used the CTA before the RCP was offered, but they chose other products for their travel. This could be because of the price of the Link-Up or misalignment between its restrictions and their travel behaviors.

To help crystalize which riders drove the travel increases, the travel patterns of each of these three groups is explored after the RCP was introduced. Figure 7-8 charts the average number of CTA trips taken by each of these three groups in the months immediately preceding and following the replacement of the Link-Up pass with the RCP. Riders who always purchased the Link-Up pass have the highest average CTA travel each month. This makes intuitive sense; they should travel the most if the higher-priced, restrictive Link-Up pass made sense for them to purchase in the first place. However, this group did not see any uptake in CTA travel frequency. It is a similar story for the riders that selectively chose the Link-Up pass; their CTA use included some sporadic increases and decreases over time. Among those newly selecting the add-on option for CTA travel, there is a clear and consistent increase in CTA use following the RCP's introduction. And this latter group constitutes by far the largest number of individual accounts and passes purchased, and these accounts clearly show the benefits of integrated fares.

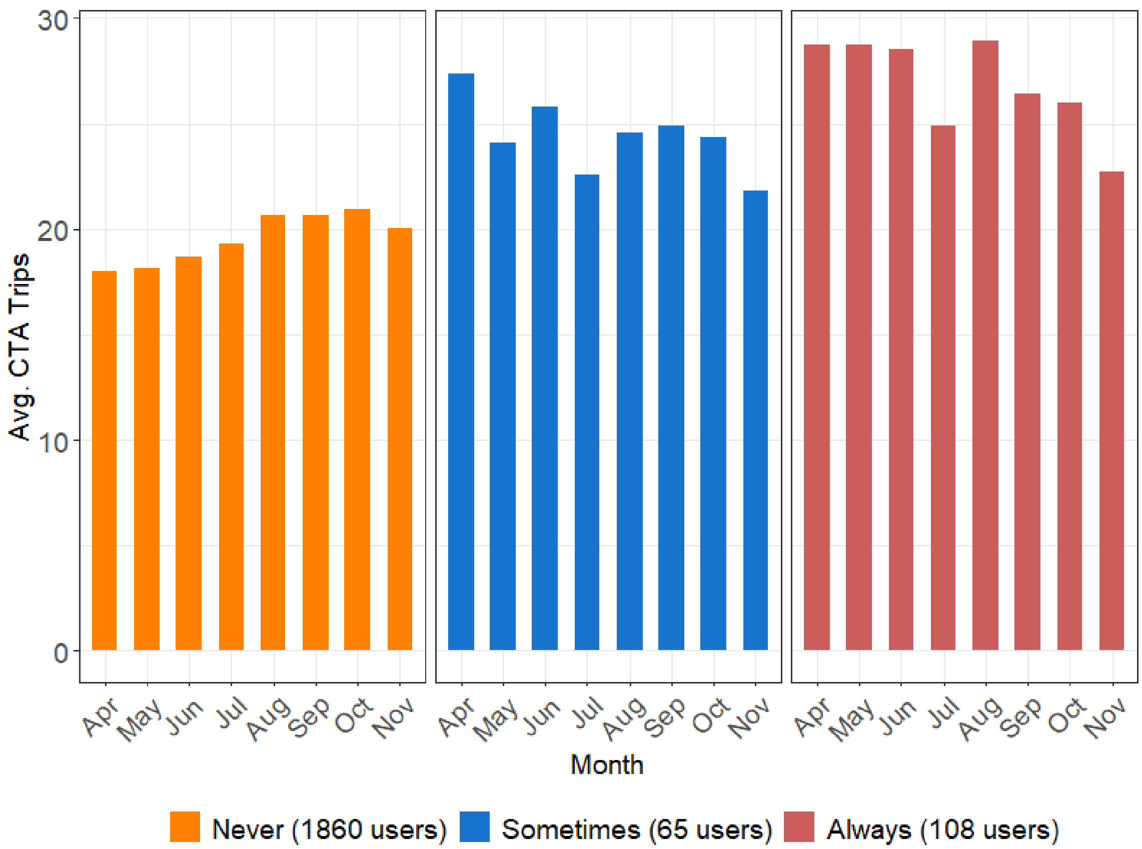

Figure 7-8: Average CTA Trips per Ventra Account

Figure 7-8 shows the mean number of trips across all riders based on their segment. However, it does not disclose the underlying distributions of CTA travel among these groups. Figure 7-9 does show the travel distributions for two months, one before the RCP's debut and one for a month following. Individuals are again partitioned according to their Link-Up use, but in this case the occasional and constant Link-Up purchasers are grouped together. Figure 7-9 corroborates the results from Figure 7-8: the RCP has encouraged Metra riders to travel more frequently on the CTA. There is a clear rightward shift in the distribution of the non-Link-Up riders. Previously, the majority of these riders took very few trips on the CTA each month. By November, the center of the distribution was just above 20 rides per account per month. These riders that never purchased a Link-Up were the target audience for the new, improved integration offered by the Regional Connect Pass. These riders responded by switching to the Link-Up from their various prior fare products, and then proceeded to use the CTA more. They are traveling more frequently, and doing so because of the advantages now available through the RCP. In contrast, the much smaller cohort of riders who had purchased the Link-Up pass show almost no change in the distribution of their CTA travel. The RCP has benefitted this group as well without the corresponding boosts to CTA trip volumes.

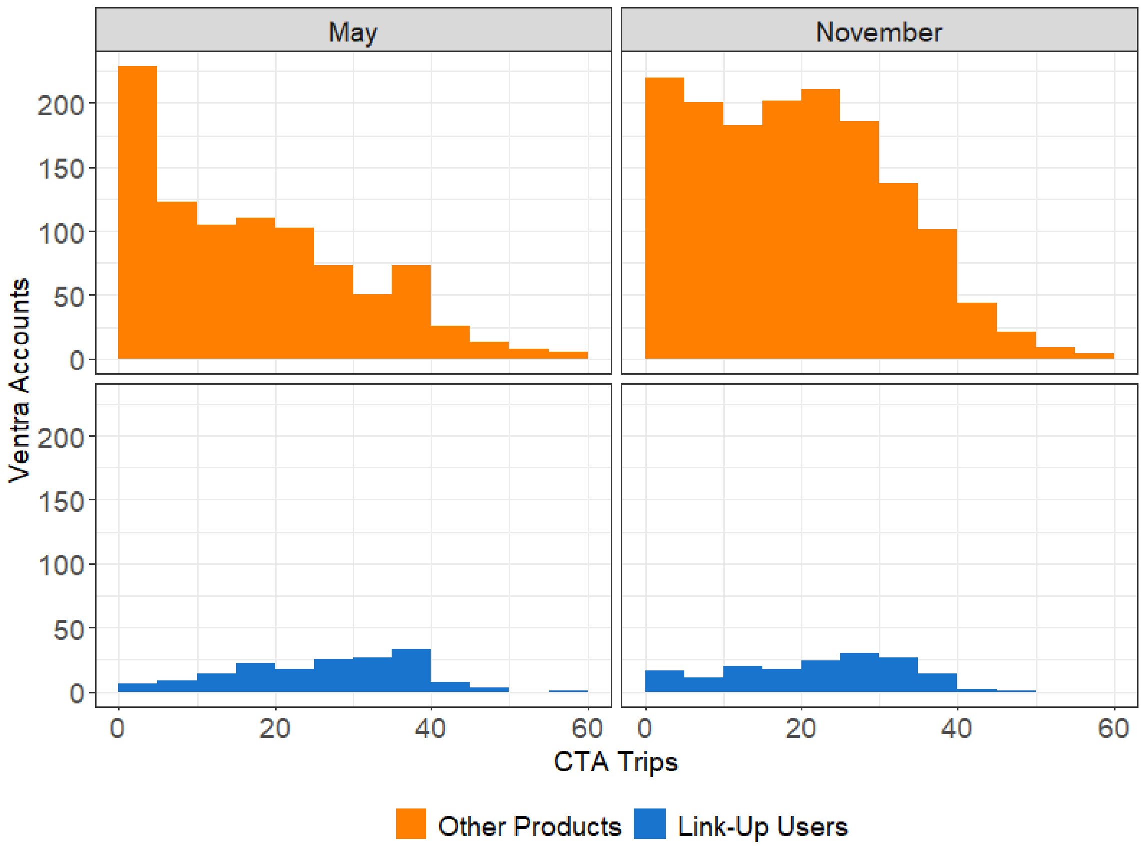

Figure 7-9: Distribution of CTA Travel by Ventra Account

The debut of the Regional Connect Pass in July 2022, coupled with deep discounts to the pricing for Metra monthly passes, have proven to be a boon to riders. The Metra monthly pass has drawn riders away from other Metra products, and this is especially true of high-frequency riders. The monthly pass users may remain a small portion of all Metra riders, but they make up a much more significant chunk of all Metra travel in the months following the pricing discounts. The fact that the Regional Connect

Pass is only available to monthly pass users may add an additional pull towards the monthly pass for high-frequency multi-system riders, potentially accounting for some of product shifting. Regarding fare integration, all of the observed changes on CTA travel by multi-system users indicate that the introduction of the Regional Connect Pass, and the improved fare integration it provided, has been a success. It has induced additional ridership, which is beneficial to the CTA. It has done this at a relatively low cost in terms of reduced fare revenues, which is again a win for the CTA. And it has given riders more travel options for seamless multi-agency transit use, improving their travel experience. Given these findings, extending integration to additional fare products beyond the current base of Metra monthly pass users would likely provide simultaneous benefits to both riders and transit agencies.

7.4.2 Post-COVID Fare Elasticities

The detailed data on fare product choice and travel behavior changes following July 2022, which concurrently saw changes to the Metra monthly pass and the RCP made available to purchasers of the commuter rail monthly pass, provides the basis for estimation of post-COVID fare elasticities. These two interrelated subsets of commuter rail users experienced the pricing changes in July 2022: first, the larger set of commuter rail riders who purchased a Metra monthly pass. Second, the subset of these riders who additionally purchased the Regional Connect Pass add-on. Utilizing the new pricing structure for Metra monthly passes, the elasticity of long-distance commuter rail travel in the post-COVID environment can be estimated. Travel patterns were changed by COVID, so it is possible that fare elasticities for commuter rail systems have similarly changed. Here we have a clear opportunity to explore this in the Chicago area. Similarly, commuter rail riders that utilize multiple transit systems experienced a price change in their access to a secondary system, in this case the CTA. This pricing change, and any corresponding changes in travel volumes, can be used to estimate fare elasticity for mass transit among multi-system riders.

The Ventra system provides the underlying data necessary for estimates of rider fare elasticity. The AFC data available through Ventra does require cleaning to ensure that calculations are made only within the pool of travelers that actually experienced the fare changes. As shown in Figure 7-4, Metra monthly pass users are a small number of all Metra travelers, but they account for a proportionally large number of commuter rail trips. This group of Metra monthly pass users are the key segment to analyze when estimating post-COVID commuter rail elasticities. Riders who did not ever purchase a monthly pass did not experience any change in the cost of travel during this time period.

The Metra riders who purchased a Metra monthly pass need to be subset further to adjust for changes in the number of zones covered by their monthly pass. When estimating elasticities, only the riders who consistently purchased products covering the same number of Metra travel zones were included. As Metra monthly pass prices were set to a uniform level ($100) in July 2022, the change in travel prices across Metra zones is not consistent. An apples-to-apples comparison requires riders travel the same distance, in terms of Metra zones, before and after the pricing changes. This is possible to identify, as even though pricing was standardized, Metra continued differentiating among fare products for each of the 10 zones in the system. Further, although the optimal decision for each rider would have been to start buying the pass covering all zones, most Metra riders did not do this. Instead, they opted to continue buying the same fare product covering the same number of zones. Among riders that purchased a Metra monthly pass following its price drop, 71% purchased only fare products covering the same number of zones between April 2022 and November 2022.

One final constraint was applied to these riders; any riders that switched from purchasing full-price products to reduced-fare products, or vice versa, were excluded. These riders similarly experienced an additional pricing change unrelated to the fare restructuring implemented by Metra. Fortunately, this is a very small percentage of riders. After applying each filter described, the resulting dataset comprises 11,324 Ventra accounts.

This dataset contains both observations across both time and across individuals, necessitating the panel formulation as described in the methodology section and shown in Formula (7.1).

Given the panel nature of the data, with each user being accounted for individually, control variables for the number of zones traveled and previous use of the Link-Up pass are unnecessary, and in fact are perfectly colinear with the individual-specific constants represented by α_i. This model is also referred to as a "fixed effects model", with α_i representing the fixed effect ascribed to each individual Ventra account. Nevertheless, to provide a fuller picture, regressions were also run using a model without individual-specific fixed effects. This model estimates rider elasticities at the aggregate level over time by comparing riders to all other riders, as opposed to tracking the changes of specific riders across months. Regression estimates for three different models are shown below in Table 7.3. First, for elasticity estimates at the aggregate level with no fixed effects. Second, with fixed effects for each Ventra account. And third, with fixed effects for both Ventra users and calendar months, resulting in panel data.

Model	**Description**	**Coefficient Estimate**	**Standard Error**	**T-Value**	**P-Value**
1	No Controls	-0.912	0.0048	-190.4	$<.0001$
2	Account-specific fixed effects	-0.779	0.9964	-121.0	$<.0001$
3	Account- and month-specific fixed effects	-0.783	0.0066	-118.1	$<.0001$

Table 7.3: Metra Elasticity Regression Results

Across the three models shown in Table 7.3, the coefficient estimate on the cost per trip variable is very consistent. This acts as the estimated elasticity in each model, and in all three of the displayed models it is highly significant. This coefficient estimate is found to be significant well beyond the 0.1% level. These findings indicate a very elastic demand for commuter rail travel in historical context. A short-run elasticity of just under 0.8 indicates a higher sensitivity price than is often observed for transit fares. This 0.8 estimate from the fixed-effects model is the lowest of the estimated elasticities, and it is controlled at the individual level for both accounts and travel months.

A regression model is also used to estimate the short-run fare elasticity of multi-system riders who are represented by Regional Connect Pass users. This is a subset of all Metra monthly pass purchasers, and preparing the data to regress requires less cleaning given that the CTA system does not charge fares based on any measure of distance. The panel data used to evaluate price sensitivity among multi-system riders is restricted to individuals that purchased the RCP at least once between its debut in July and November 2022. These individuals are not only multi-system riders by definition, they also experienced the price change and reacted to it. That reaction could take the form of benefiting from a lower price, as with Link-Up users, or the reaction could be adoption of the add-on pass in place of a different fare product. The second requirement for inclusion is identifiable travel on the CTA prior to July 2022, regardless of which fare product they used to travel. Riders with no previously identifiable travel have unknown travel demand before the introduction of the RCP, and there is no average cost per trip if there are no trips taken. Applying these requirements to the data leaves 1,380 unique accounts to analyze.

Table 7.4 shows the results of three different fare elasticity estimations. As discussed with regards to the Metra monthly pass, the fact that this is panel data reduces the number of variables potentially needing to be controlled in the fixed effects formulation. Previous use of the Link-Up pass, rider use of transit benefits, and general ridership trends are all captured by the individual-specific constant estimates

α_i. Nevertheless, the results from two simpler models are also included in the table. First, a model with no fixed effects which compares total travel and overall cost at the aggregate level for all ridership. Second, a model with rider-specific fixed effects, where each rider is compared to themselves. And third, the panel formulation with fixed effects for both the rider and the month of travel.

Coefficient estimate representing fare elasticity is 0.56. This is lower than the estimate for commuter rail users, but still on the high-side of historical estimated short-run transit fare elasticities. The results further show that at the aggregate level, ignoring fixed effects, the impacts of improved regional integration are even higher. Riders are still showing a clear sensitivity to the costs of multi-system integration, even if that sensitivity is less than among monthly commuter rail riders as a whole. This is understandable as Commuter Rail riders are choosing whether to use transit or another form of travel, whereas the RCP riders are choosing whether to *keep* using transit.

Model	**Description**	**Coefficient Estimate**	**Standard Error**	**T-Value**	**P-Value**
1	No Controls	-0.781	0.0145	-53.7	$<.0001$
2	Account-specific fixed effects	-0.567	0.0152	-37.3	$<.0001$
3	Account- and month-specific fixed effects	-0.550	0.0160	-34.3	$<.0001$

Table 7.4: Regional Connect Pass Elasticity Regression Results

7.4.3 Add-on Products

AFC data records from the Chicago area show clearly observable changes in use of CTA services by Metra riders after improved integration of the available fare offerings. There was an increase in multi-system travel across all days and times of the week, and customers that were new to the integrated product showed the largest response in terms of increased trip-making. These observed changes were used to model post-COVID fare elasticities for long-distance commuter rail travel and multi-agency integrated products. The modeled estimates showed relatively high sensitivity to transit fares in a historical context.

The impacts felt and insights learned from the Regional Connect Pass are useful to guide integration among other CTA and Metra fare products. Benefits from multi-agency fare products are currently restricted to individuals that choose to purchase a monthly pass for commuter rail travel. As shown, monthly pass purchasers are a small portion of all commuter rail riders, and a proportionally larger volume of commuter rail trips. Even so, monthly pass travel constitutes less than $\frac{1}{3}$ of all recorded commuter rail trips. That leaves a significant portion of trips, and an even great number of individuals, that would potentially benefit from multi-agency integrated fare offerings. Coaxing some of these riders to an integrated product presents opportunities to further boost ridership at each agency.

A case can be made that agencies should provide a variety of integrated add-on products that are available to mix-and-match between service providers. Riders would be able to choose the options that best meet their regional transit needs from an a la carte menu of offerings. For instance, if Metra monthly pass users could add on a discounted CTA day-pass when needed throughout the month that could capture a portion of users that have not chosen to purchase the Regional Connect Pass. Bringing these riders onto the CTA for some travel could create the habit of regular multi-agency use, perhaps convincing them of the RCP's utility such that they purchase it in a future period.

Assuming the agencies do not want to offer mix-and-match fare integration, which they have not as of yet shown a propensity to do, identifying similar products between agencies that could be leveraged as a jointly integrated offering. When evaluating additional integrated offerings, the fare products at each agency can be prioritized based on the propensity of riders towards multi-agency travel. This could include a fully integrated 1-day pass, offered at a slight discount when compared to purchasing each individually, or a single-ride ticket. How each of the potential like-product combinations should be ranked in order of potential market size is a question that must be answered to guide this process. Which Metra fare products serve riders that have high rates of multi-agency travel, and which do not, should be considered.

Figure 7-10 displays the number of CTA and Metra trips taken by each Ventra account. Travelers that use both agencies relatively equally would appear equidistant from both axes, while those riders located near an axis do not travel equally on each agency. High-frequency users of both the CTA and Metra are circled in green. These charts can help lay the groundwork for further exploration. It appears that promotional products, such as weekend and holiday passes, serve riders that travel on both agencies relatively evenly. In particularly, it appears a large number of Metra monthly pass users do travel on the CTA, even if they do not purchase the RCP. They travel using weekly passes, promotional passes, some pay-per-use travel, and (shockingly) travel on the CTA monthly pass. Any riders using both a Metra monthly pass and a CTA monthly pass should absolutely be using the RCP; the reason that they are not deserves some exploration. Similarly, Metra's 10-ride ticket and day pass offerings see riders that travel on the CTA at levels relatively close to their Metra travel. The results also suggest that CTA travelers are sometimes drawn to Metra for occasional travel when Metra offers promotional products, as demonstrated by the CTA 30-day pass users that travel using uncommon and/or promotional Metra fare products.

Preliminarily, 1-day pass offerings and promotional products both seem like promising areas to explore integrated products. In particular, promotional and weekend fare products can be seen as a first step in closer integration, or as a useful pilot through which to explore integration. Riders that take advantage of these are often infrequent transit users, so offering them a new and improved product and then tracking whether

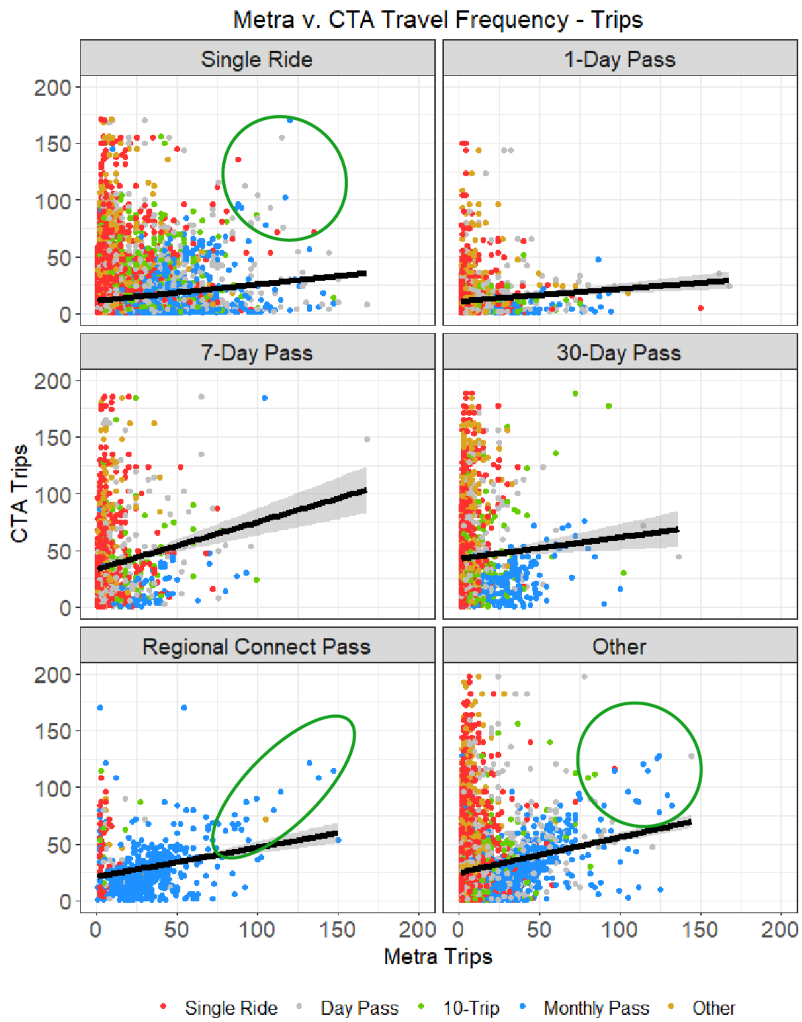

Figure 7-10: Metra and CTA Travel Frequencies

that product encourages them to use transit throughout the region more often could provide strong support for continuing widespread integration.

Next steps should focus on differentiating between multi-system riders that are CTA-first riders and Metra-first riders. The primary transit provider for each Ventra account can add context to any integrated fare offerings, as well as guiding agency strategy for rolling out and marketing products. Pilot programs for integrated fare offerings would be a boon, particularly among lower-frequency commuter rail travelers. If a similar number of 1-day pass and single-ride travelers chose an integrated offering, and responded by increasing travel as obviously as witnessed following the

RCP's introduction, a large market of potential ridership can be tapped, benefitting agencies. From the rider perspective fare integration improves the travel experience, reduces costs, encourages transit use, and expands travel possibilities. Regional fare integration makes sense to all involved parties and encourages overall transit use in the region benefitting operators as a whole.

Chapter 8

Conclusion

This thesis presents four case studies used to reevaluate major aspects of transit fare policy. This is done with particular attention to how fare-policy decision-making and public transport agency priorities have changed in the aftermath of the COVID-19 pandemic. Public perceptions of mass transit have continued to develop with greater focus placed on transit-dependent riders and front-line employees. Alongside societal awareness and the spotlighting of continuing inequities, questions of fare equity have only been magnified. Massive COVID-induced shifts in the travel patterns of the transit-riding public, and (at least temporary) reductions in the size of that group, have altered the utility of traditional pass products, providing a powerful impetus for agencies to provide additional offerings. Changes in ridership levels and travel behaviors further provide an opportunity for transit providers to collaborate more closely on integrated fares, incentivizing transit use across entire metropolitan regions and providing individual agencies avenues for capturing new customers.

In this chapter, the results of the four case studies presented in this thesis are summarized. Key findings from each of the analyses are restated and used to identify high-level commonalities. Drawing from these commonalities, generalized insights and recommendations are made on future directions for fare policies, fare products, and fare research across scales and geographies.

8.1 Key Findings

The case studies presented in this thesis each individually provide insight on specific aspects of effective fare policy in the current post-COVID environment. Compared to traditional agency positions or pre-pandemic studies, the findings from each case study reinforce the magnitude of change with which transit agencies are grappling. But these case studies also help outline strategies to encourage transit use, streamline usability, and ultimately grow ridership while better serving riders. The environment for travel has changed drastically, and transit providers must respond to remain competitive. Fare policies and fare products are one tool at their disposal.

Two examples of agencies going against traditional precedent have been witnessed in Chicago. The CTA's reduction in pass pricing and Metra's deeply discounted monthly passes were choices made to encourage transit ridership even at the expense of fare revenues. These were bold choices; agencies face a difficult balancing act in that they exist to provide transportation options but providing that service requires expenditures funded, at least in part based on past public policy, from collecting passenger fare revenues. Post-pandemic travel numbers, with ridership levels still well below pre-pandemic levels, suggest a primary short-term need to increase ridership even at the expense of traditionally robust fare recovery targets. A focus on serving riders, and growing ridership, can pay dividends.

Particularly important takeaways from each of the case studies presented are as follows:

- Given the local context in which WMATA operates, **fare structures that retain peak pricing differentiation improve equity for low-income riders**. Further, a straightforward move to all-day fares in a revenue neutral way would actively disadvantage low-income riders who would face fare increases for many trips. Additionally, under the existing distance-based fare structure the penalties for indirect Metrorail travel are inequitable. Low-income riders and non-white riders both take more circuitous trips, requiring more time and higher fares. Alternate fare structures that maintain higher rates for circuitous trips likewise maintain the inequitable effects of the existing distance-based structure. A simple, understandable solution that improves equity exists: **WMATA should switch to crow-flies miles as a measure of distance for travel**.

- The CTA's decision to reduce transit pass prices during the COVID-19 recovery was an effective step to improve the relative utility of those passes given COVID-caused changes in transit travel patterns. Prior to the price reduction, pay-per-use fares comprised not only the majority of pandemic-era travel, they were also the optimal choice for many individuals. However, opportunities to further enhance pass products remain. **A preset number of one-way journeys on the CTA, with an upfront payment at a reduced per-trip cost, would offer the most benefit to riders whose current best fare product choice is to ride pay-per-use.** Such a pass could induce additional travel, while simultaneously creating an upfront commitment to transit use by the customer.

- Existing short-term transit passes offered by the CTA do not provide the same incentives for purchase to all riders. This is due to **factors that include the fare differential between bus and rail, a rider's fare class, and travel expectations**. Single-day fare capping, with no changes to other pass products, would have small impacts on CTA fare revenues, but would offer opportunities to spur additional travel. For infrequent or occasional riders, it would also simplify travel decisions. Analysis shows that benefits from single-day capping would accrue to residents in the north and west of the CTA's service area. But that shines a light on two important aspects. First, **there is no 1-day pass for reduced fare riders**, leading to low pass adoption. If a single-day cap is not possible, **reduced fare options for all passes are necessary, and consistent with existing CTA strategy**. Second, **Ventra penetration appears much lower on the South Side and West Side in disadvantaged neighborhoods. Increased usage of Ventra would benefit these residents even** absent fare capping, and strides should be made to increase Ventra usage.

- Regional fare integration across commuter rail and mass transit systems is an effective way to both increase ridership at individual agencies and boost the appeal of transit travel to the riding public. Access to multiple regional transit providers at a single price has increased travel on each; collaboration between these agencies effectively creates an even larger transit network with each provider acting as a single component. However, the opportunity for unified travel at only a small price premium is currently only available to a small fraction of all transit users. The market for additional travel is much larger and remains untapped. **Integrated products beyond the Metra monthly commuter rail pass should be introduced. Opportunities for CTA users to access Metra at a lower cost premium need to be explored.**

Another substantive finding that repeatedly arises across each of the case studies, and indeed is key to this entire thesis, is the existence of rider-focused policies that simultaneously support substantial opportunities for transit providers to both boost ridership and align fare policies with rider behaviors. For example, the fare equity analysis conducted on Metrorail fares in Washington, DC pinpointed discrete aspects of fare policy that can be changed to improve fare equity. Fare structure changes were explored ranging from minor tweaks to extensive revisions, indicating the ability agencies have to be creative with fare policy to achieve their priorities.

Estimates showed that many of the alternate fare structures generated additional transit travel among certain rider segments as a positive knock-on effect. However, in other alternate scenarios, any additional travel had inequitable side effects causing large negative financial impacts to historically underserved groups. Ultimately, the analysis was able to illustrate fare policy decisions that would increase ridership without harming equity, benefit equity without decreasing ridership, or benefit both simultaneously. These opportunities to improve fare equity were possible within a revenue neutral framework, and some of the beneficial changes require only slightly adjusting existing structures. This is exemplified by the measure of distance used to charge fares, which is currently the composite mile. Simply switching to the crow-flies distance between origin and destination would both improve equity and more closely align the benefits and costs of transit.

Similarly, the exploration of multiple hypothetical "flex-pass" products that could be offered by the CTA illustrates the extent to which pre-pandemic pass products lost their competitive edge to pay-per-use travel. The CTA fare restructuring that was made permanent in fall 2021 successfully realigned the utility of various fare products, better incentivizing pass product purchases. This fare restructuring was a straightforward step taken by the CTA that was clearly rider-oriented, and it was successful in encouraging uptake of pass products. The methodology developed in this thesis can be further extended to fare product designs not analyzed here.

The third case study considers single-day fare capping as a potential fare product in Chicago. The riders that would be affected are identified, fare revenue estimates are generated, and differing impacts to different rider groups are explored. From here, the benefits of a single-day cap are made clear in relation to the existing short-term pass offerings which are not uniformly appealing to bus riders and rail riders, nor to reduced-fare riders and full-fare riders. A cap could address these differences, making single-day travel discounts the same for all riders, incentivizing more transit use, and leaving current differentiated fares intact.

In case study four regional fare integration across transit providers in the Chicago region is explored. The results indicate that by collaborating to offer fare products that confer travel on multiple transportation networks, transit agencies can successfully tap into new markets for ridership. The to-date measurable impacts of the multi-agency Regional Connect Pass clearly support this. CTA travel volumes among multi-system riders that adopted the Regional Connect Pass increased across weekdays and weekends. Additional riders were drawn to the CTA who had no identifiable previous travel. As with findings from the other case studies, rider-focused fare policies are not only able to coexist with but can actually complement agency goals around maintaining and even increasing ridership.

8.2 Commonalities and Differences Across Case Studies

Comparing individual results across the case studies can identify where commonalities exist, as well as to help pinpoint any contradictory findings. The most important, consistent finding across these case studies is the overlap between rider-focused fare policies and agency ridership goals. This is true when evaluating equity, when designing novel pass and fare product offerings, and when exploring regional fare integration. Additional commonalities exist and are important to recognize as well. Not every solution that aids riders doubles also supports internal agency goals, but exploration and analysis can identify synergies to leverage. Win-win policies are important to identify as the impacts of COVID-19 on transit seem to have reached an equilibrium; transit agencies should plan accordingly, instead of continuing to look to pre-pandemic norms for guidance.

This thesis has shown that fare equity in Washington, DC is dependent upon the geographic distribution of residences and workplaces, but post-pandemic it has also been impacted by differing reliance various rider segments have on Metrorail. Another case study shows that in Chicago, the utility of existing fare products has changed considerably due to pandemic-induced travel changes. In both case studies, not all pre-pandemic trips were equally likely to remain necessary after COVID-19, leaving agencies with tough decisions about the most effective and most equitable changes that are possible. Working within constrained resources, agencies must consider whether to target continuing riders with enhanced fare offerings, or instead focus on appealing to inactive and lapsed riders. To provide one specific example, commuting travel remains an important component of transit ridership, albeit much smaller than in the past. New fare products more closely aligned with post-pandemic commute patterns could attract infrequent commuters away from pay-per-use single-ride fares and induce more frequent use through the provision of new discounts. Such a hypothetical commuting-oriented product could also invite riders onto the transit system and generate an upfront commitment to at least a minimum amount of transit use. This, in turn, can help reinforce the practice of utilizing transit, creating longer-term habits of transit use. However, commuting-oriented products would not necessarily improve transit access or streamline travel for transit-dependent individuals that never stopped riding in the first place. Transport providers have limited resources, and decisions on how to prioritize different rider segments are driven by agency values and policies, but decisions must inevitably be made.

The case studies presented in this thesis do find tradeoffs between certain policies oriented towards increased ridership at an aggregate level and approaches to promoting fare equity among rider segments. This is most obvious in the findings from the equity analysis conducted on WMATA's system. The removal of peak pricing lowers the average fare for the largest number of trips and would be expected to result in higher systemwide trip volumes. But these ridership gains would occur at the expense of low-income riders, who would face disproportionate increases in fare levels as they take off-peak trips at higher levels. Looking to Chicago, the current fare integration offerings are targeted at frequent commuter rail users that purchase a monthly commuter rail pass. Riders that cannot afford a monthly pass, or CTA users with alternative service options from Metra, are excluded. Expanding the range of available integrated products could improve equity by providing options for these riders, while also increasing ridership overall by capturing new trips and inducing new travel, both of which have been demonstrated following the introduction of the Regional Connect Pass.

One other important commonality across all case studies in both Chicago and Washington, DC occurs when comparing the analytical findings to prior expectations. Fare policy can be used to improve the rider experience, boost travel volumes, and the negative revenue impacts are estimated to either be lower than prior assumptions or can be easily mitigated through modest changes to the fare structure. Traditional worries about significant ridership or revenue losses are not borne out by the data in any of the four case studies examined here. Instead, developing an equitable fare structure supports ridership and equity benefits in tandem. Novel fare products can capture riders who are not served by any existing products and have no lasting commitment to transit beyond a payment each time they board. Single-day fare capping could encourage discretionary travel that might not otherwise be taken, adding trips without reducing revenues. Fare integration has clearly increased ridership significantly with little cost in terms of revenue. A reevaluation of fare policy shows that, in the current environment, any negative impacts of rider-centric fares are small and are often overcome or outweighed by the benefits they provide.

8.3 Important Considerations for Post-COVID Fare Policy

This thesis has developed and presented four insightful analytical approaches that agencies can draw from to help guide their work. Nevertheless, the analytical frameworks must be adjusted to account for the local environment in which an agency operates. Local context is key in accurately diagnosing where policies can be improved. The geographies of residents, employment centers, and travel patterns within an agency's service area must be considered. These aspects, alongside the physical infrastructure of the transit network, all impact the utility of fare products and the equity of fare policies.

The importance of local context has been found consistently across the previous fare equity literature, and it is reinforced by the new methodology this thesis presents to analyze the equity of Metrorail's current fare structure in Washington, DC. The WMATA case study finds that a revenue-neutral switch from distance-based fares to flat fares would disproportionately disadvantage low-income riders. This finding is an artifact of where low-income populations live and how they travel in DC, meaning this same finding is not guaranteed to hold in other metropolitan areas. It is also important to note that if revenue neutrality is removed as a constraint, each of the alternative fare scenarios could be structured to benefit all riders compared to the existing baseline, albeit to differing extents. Revenue neutrality is necessary to contextualize the relative impacts on equity of various aspects of fare structure, but in practice revenue neutrality is an agency decision that could be abandoned in pursuit of both increased equity and ridership.

Along these lines, the choice to move away from a staunch focus on maintaining and maximizing fare revenue has been observed in the Chicago area, resulting in positive ridership outcomes. The tradeoff between maximizing fare revenues and spurring transit ridership has always been a balancing act for public transport. In our current environment the calculus may have shifted, given society's increasing attention to sustainability, environmental justice, and equity. The COVID-19 pandemic helped to thrust these topics further into the spotlight, while simultaneously upending transit norms. Transit agencies will need to review and reconsider longstanding policies through this prism to be responsive to local and societal priorities.

It is important for transport providers to be responsive because the many benefits provided by mass transit extend beyond those experienced by the riders themselves. For instance, public transport has been shown to expand the economic opportunities available to all citizens of metropolitan areas and the communities within them. Viable, effective transportation options enhance the economic vitality of regions and boost their attractiveness to employers. The advantages of agglomeration to both business and employees have long been studied and are well established; an efficient, frequent, and reliable transit network helps ensure that full advantage of these agglomeration effects can be accessed and exploited. Doing so can help support the business and political cases for continued funding to offset ongoing COVID-induced revenue losses.

Travel options beyond the traditional automobile can help lower carbon emissions, reduce local pollution, and support longer-term sustainability, all of which are paramount to combating climate change and advancing the continued livability of our urban areas. The equity, economic, and environmental impacts of transit use are external to any individual public transport agency, but they are nevertheless inherently and exceedingly valuable impacts. What is good for supporting transit use is good for society. Transit agencies must focus on rebuilding transit use, and fare policy is one tool that should be used.

As transit ridership continues to lag below pre-pandemic levels, societal benefits should be emphasized to justify and reset expectations around farebox recovery ratios. Some public transport agencies have mandated target fare recovery ratios set by governmental bodies. The CTA is one such agency, which is subject to a state law requiring that 50% of operating revenues be covered by fare revenue collection. Such requirements should be reconsidered, and additional external funding sources secured. While fares are a valuable funding source for transit, supporting effective transportation options, economic opportunity, and region-wide livability should be primary purposes of public transit. The current state of work travel, with work from home and hybrid schedules continuing to demonstrate staying power post-pandemic, suggests that a focus on measures of success beyond fare recovery is necessary.

Ridership has traditionally been used as a proxy measurement for the success of a transit agency, and within this thesis it has continued to be used as such. Nevertheless, it is only a proxy. The real goal should be to increase the percentage of all travel that is taken on mass transit systems. Given changes in travel and commuting patterns due to COVID-19, measuring travel volumes in absolute terms and comparing to a pre-pandemic baseline is no longer the most effective or insightful lens. Fare policies and products that encourage increased transit are a more applicable perspective for evaluating transit success.

Capturing travel, and growing ridership, should be primary goals of post-COVID transit fare policies. Agencies will always need to navigate the delicate tension between revenue generation and ridership, even if that balance may have shifted in the wake of COVID-19. New and reimagined fare products can contribute to the twinned goals of increasing ridership and generating additional fare revenue, but as in all realms of fare policy any changes must be introduced with attention paid to equity. Gaining new customers is important, but capturing additional trips made by existing transit riders is also essential. As with any proposed improvement, it is important to determine which riders are being served, and which do not experience any benefits. Not all individuals are equally in need of additional support. Discounted travel for its own sake is less compelling when the beneficiaries are already relatively well-off. Accordingly, setting any policy requires a high confidence that the desired beneficiaries will be effectively targeted. Fortunately, the case studies presented in this thesis demonstrate that fare equity improvements, new fare products, and enhanced regional fare integration provide substantial opportunities to benefit both rider and agency. Additionally, along each of these dimensions the magnitude of change required to improve equity is not singular, nor are the impacts. Instead, potential actions to improve equity span a range from small policy tweaks to major structural reconfigurations. There are a range of actions agencies can take, individually and in cooperation with each other, to improve fare equity that overlap with the goal of boosting ridership to some degree. These gains in total transit ridership across all agencies, which can be achieved through collaborative decision-making, can build support for transit funding at an aggregate level.

Finally, one still unknown but important consideration is whether transit ridership will continue to stabilize around its current levels, or whether there will be a move away from remote work and back towards pre-pandemic travel frequencies at some point soon. There are a number of potential transit ridership recovery scenarios that are yet to play out. This thesis, and the case studies it presents, focus on what I consider the most likely scenario, which also happens to be the current scenario – namely that travel behaviors have reached a stable equilibrium and this steady state is the environment in which transit providers will be operating for the foreseeable future. However, this outcome is not certain nor predefined. Mass transit has some influence over its own future; agencies are not completely powerless to outside forces. Operational and capital choices are two of the strongest levers transport providers can use to influence the future. Fare policy is another powerful tool at their disposal. Decision making around fares must consider multiple potential ridership scenarios. Nevertheless, the findings presented in this thesis would not be negated by a future that has higher than anticipated transit use. The methodologies developed and employed in each of the four case studies are purposefully responsive; they can be revised and reevaluated as more information becomes available. Policymaking and approval take time; the analyses in this thesis can continue to support and update findings throughout an iterative process. The case studies assume that existing pass products would remain available; should they again become useful to a large segment of riders those riders could simply switch back to them. Equity improvements are not based on ridership levels, they are informed by trip making across rider groups. The benefits to individual riders of regionally integrated fares will remain just as appealing regardless of how many riders are eligible to purchase integrated fare products. While the exact future of transit ridership is an unknown, the methodologies and findings detailed within this thesis can nevertheless guide fare policy decisions.

8.4 General Recommendations and Research Extensions

This thesis lays out approaches to evaluating key aspects of transit fares and does so within our post-COVID environment. Nevertheless, there remain many additional aspects of fare policy worthy of detailed investigation.

Throughout the case studies presented, the benefits fare products could provide were used as a measure of utility to ridership, but questions of how riders choose a fare product were set aside. This second question is integral to ensuring that the benefits of transit fare products reach the intended audience. Understanding why riders make the choices they do can inform better design of new fare products, and potentially allow those products to be more closely tailored to specific rider segments.

One particular aspect of transit user fare product choice that is perennially debated is the value individuals place on fare simplicity. It remains undetermined how important a simple fare structure is for generating transit use. Many people assume that there is some inherent value associated with fare simplicity, but the measurement of that value remains unknown. Further, there may be a subtle but important difference between fare simplicity and fare legibility. The number of inputs used to calculate the cost of a trip and the ways they interact can be complex, but perhaps that complexity can be masked or negated by having the cost clearly communicated. As technology continues to become more ubiquitous, fare information that is currently presented through fare tables might be communicated in more legible formats. In the context of WMATA, the distance-based fares charged for Metrorail use provide opportunities to improve fare equity. The value, if any, of fare simplicity is assumed to be constant across rider segments. However, this may not be the case. It is possible that the importance of fare simplicity or fare legibility varies substantially across rider segments. Any such differences are not accounted for when estimating the impacts of alternate fare structures on ridership. An estimated value for simplicity could recontextualize the findings. Perhaps a move to a simple fare structure such as a flat fare would recoup some of the estimated ridership losses by attracting riders drawn in by a simple, understandable cost. Similarly, a high value for simplicity may suggest that single-day fare capping will spur even more additional trips than might otherwise be expected.

Beyond exploring how riders choose between multiple fare products and seeking to quantify the value riders place on simplicity, there are non-monetary costs associated with transit use that are not captured by the analyses presented here. The time required to travel from origin to destination is its own form of cost and perhaps the primary determinant of mode choice. Since these costs are in addition to the monetary fare due, they are not incorporated into the results presented here. Under a distance-based fare structure, there is the probability that certain riders pay higher monetary fares and bear higher costs in terms of travel time, resulting in a doublepenalty. Synthesizing the interactions between fares and non-monetary costs could help come up with a more holistic measure of transit costs, including a breakdown of the portion of those costs that are governed by agency fare policy.

Regarding agency fare policies, long-standing expectations and requirements around farebox recovery ratios need must be reevaluated. While ridership levels remain low, the idea of a 50% recovery ratio seems insurmountable, even though the services provided by transit agencies remain integral. Even a lower recovery ratio requirement may be unsustainable. The costs of public transit have a history of rising more quickly than the rate of inflation, and fare revenues have not kept pace. Given this trend, and the need to ensure transit is accessible to all riders, questions about the underlying concept of the recovery ratio arise. Other funding streams to support the provision, operations, and maintenance of transit services must be identified. This creates a quandary, as political support for transit likely relies on effective systems and high ridership, and system operations and ridership are impacted by fares, creating a difficult feedback loop. Aggressively courting ridership may be worth considering in the short-term to build ridership and corresponding political support for transit, but research would be necessary to identify effective strategies for doing so.

In addition to transit fare elasticity, estimates of cross-elasticities would help inform society-wide transportation decisions. New regulations affecting future sales of automobiles, in particular mandated shifts away from internal combustion vehicles, will impact the consumer market in ways unknown at present. Widespread electric vehicle adoption, which is anticipated, will require households to purchase vehicles with price tags that are (at present) clearly higher. It also requires costs associated with installation of all new charging infrastructure, or access to installed infrastructure somewhere outside the home. Given these cost increases, and questions about charging infrastructure, perhaps effective public transit systems can act in place of second (or third, or fourth) cars, reducing the number of average cars per household.

Cross-elasticities for transit and newer technologies could guide policy, especially if long-term elasticities show an ability to impact long-term spatial decisions.

Questions also remain about cash-paying transit users. Each of the four case studies in this thesis make extensive use of AFC records to analyze travel patterns, fare levels, and transit usage in both Chicago and Washington, DC. The available data records are comprehensive and provide powerful insights on most transit users. They do not capture all transit riders, however, and the riders that are absent are more likely to come from historically disadvantaged groups. In DC, this manifests clearly in the different ratios of fare payment media used between rail and bus. Cash payments are not accepted for rail, but they make up a very high proportion of all bus fares. Bus riders, especially those without SmarTrip AFC cards, are disproportionately low-income. In Chicago, the geographic distribution of Ventra AFC cards is not uniform across the CTA service area. Minority and low-income areas to the south and west side of downtown show much higher rates of travel on paper tickets and other fare products that are not associated with any Ventra accounts. Account-based analysis of travel patterns misses a higher percentage of trips taken in these areas, potentially underestimating any impacts on these riders. Increasing the number of riders using AFC cards in both systems can help build up datasets for more complete analyses in the future, but the additional benefits extend beyond fare research. The riders currently paying for rides with cash or riding on paper tickets deserve to access the same benefits and transit modes other riders already experience. There are few, if any, benefits for regular riders continuing to travel via cash or limited-use tickets. Further, the entire cost-burden of transit for these riders could be calculated to arrive at a more holistic understanding of transit use. The non-monetary costs associated with transit travel often compound inequities caused by fare policy decisions. Slower travel is often the cheapest in monetary terms, sometimes excluding riders from trips that are more efficient for travelers to take and for agencies to operate.

Strategies to significantly increase the use of employer-based transit benefits are worth considering, both to support riders and as a funding source for agencies to tap. The economic benefits from agglomeration in large urban areas and the negative costs imposed by significant congestion suggest the need for this expansion from an employer perspective. In addition, the pre-tax nature of transit benefits can be leveraged to make existing pass products a competitive choice, even if office commutes remain less frequent than prior to COVID.

In this same vein, providing transit benefits programs through non-employer channels or introducing low-income fare products deserve continued consideration. Existing employer-based transit benefits allow pre-tax money to be used to purchase travel, lowering the relative cost of transit fares for eligible riders. Riders from disadvantaged groups are less likely to have access to employer-based transit benefits programs, excluding riders with high need from access to lower transit prices. Reduced fare offerings for such riders would work to level the playing field. Additionally, fare elasticity estimates for lower-income riders and riders from other disadvantaged groups may be subject to income effects as transit fares correspond to a more significant portion of their income. Such impacts have been explored, and differentiated elasticity estimates have been used in this thesis where available, but continued refinement can further demonstrate opportunities for rider equity and ridership gains. The opportunity to save on transit fares can spur ridership across all rider segments, likely at even higher rates across disadvantaged rider groups. Advantages should not be restricted to only traditional commuters that currently have access to transit benefits programs. Expanding access to transit benefits and/or providing reduced-fare offerings can improve equity outcomes, and it can be considered a move towards fairness for all riders as opposed to a give-away targeted at only certain groups.

Lastly, dissimilar ridership patterns and rider segments across transit modes deserve further attention to ensure effective and equitable transportation is supported by an agency's fare policy and fare products. As mentioned, across WMATA there are clear differences between the makeup of bus users and rail riders. In Chicago, the quest for regionally integrated fares has focused on opening CTA travel to Metra commuter rail travelers. CTA users could technically access the Regional Connect Pass, but the required cost premium is substantially greater for CTA monthly users than it is for Metra monthly users. This makes the barrier to regionally integrated fares much higher for CTA users. Piloting additional integrated products, especially those with shorter validity periods, may provide more equitable access than is currently available. When considering both metropolitan areas and the transport agencies serving them, the modal differences between ridership is an area rich for further exploration. Such exploration can add further context to discussions of fare equity, and guide development of effective fare offerings.

Appendix A

Updated Post-COVID Fare Equity Results

Chapter 4 develops a methodology for analyzing transit fare equity which is then deployed in the Washington, DC region. WMATA's distance-based Metrorail fares were evaluated through an equity lens at the time of the most recent regional travel survey ('RTS'). Detailed more fully in Chapter 4, the RTS was conducted by the Metropolitan Washington Council of Governments between October 2017 and December 2018. For this reason, it is important to specify that this most recent travel survey, and the contemporaneous AFC records used to contextualize fare equity, preceded the COVID-19 pandemic. While the geographic distribution of jobs and residences are unlikely to have changed substantially between the analysis period and the present, travel frequencies and commute patterns have changed substantially. Further, such changes may have been inconsistent across rider segments. For example, the ratio of peak and off-peak trips is likely to have shifted with pandemic-induced increases in remote work. The transit riders most able to take advantage of remote and hybrid work opportunities are likely concentrated among certain ridership subsets. To understand Metrorail fare equity in a present-day context, it is important to reevaluate the aspects of the existing fare structure and understand which advance equity and which harm in the aftermath of COVID-19. Ideally, revisiting the analysis with post-pandemic travel will reconfirm that the evaluated fare structure features retain the same impacts in terms of direction and magnitude.

The same methodology is reapplied to post-COVID AFC records to understand the extent of any COVID-induced fare equity changes. However, this requires once again that rider demographics be confidently linked to specific journeys taken on the Metrorail network. The RTS is conducted approximately once per decade, meaning that no post-COVID survey has been administered and no survey responses are available. Nevertheless, the SmarTrip card numbers identified in the pre-COVID period provide one potential source of rider demographics. These SmarTrip card numbers all have associated rider demographic information. The rider demographics were confidently assigned to representative AFC trips in one of two ways. The first approach required matching travel-diary entries made by RTS respondents to travel records in WMATA's Trace database. The second, and more straightforward, approach simply uses demographic information provided by SmarTrip card holders when they register their card with WMATA. The second approach is easily repeatable across time periods, and can be reapplied to post-COVID data. The first approach, which resulted in SmarTrip card numbers matched to RTS responses, provided a large sample of representative trips during the RTS period. However, some RTS responses returned matches to multiple trips recorded in the Trace database. While any of those matched trips can be randomly selected as representative of the demographic group to which the RTS respondent belongs, the exact card belonging to the RTS participant remains unknown.

This issue with multiple matching records was most common for RTS travel diary records occurring between high-ridership stations during the busiest parts of the day. While a random matching trip can be used during the RTS period, multiple possible SmarTrip card number matches prevent a specific card from being identified and reused for analysis in other time periods. Fortunately, a number of RTS travel diary responses did return only a single SmarTrip card match from the Trace database. Among these singly identified cards, the demographics of the owner are known with high confidence. These singly matched cards allow inferred demographics to be assigned to an anonymized SmarTrip card number itself, not only to a representative trip. Figure A-1 illustrates this approach, with the intersection of Trace records and RTS responses identifying one set of SmarTrip cards, and the second group of cards with known demographics being identified once again directly from the labeled, registered SmarTrip cards in WMATA's dataset.

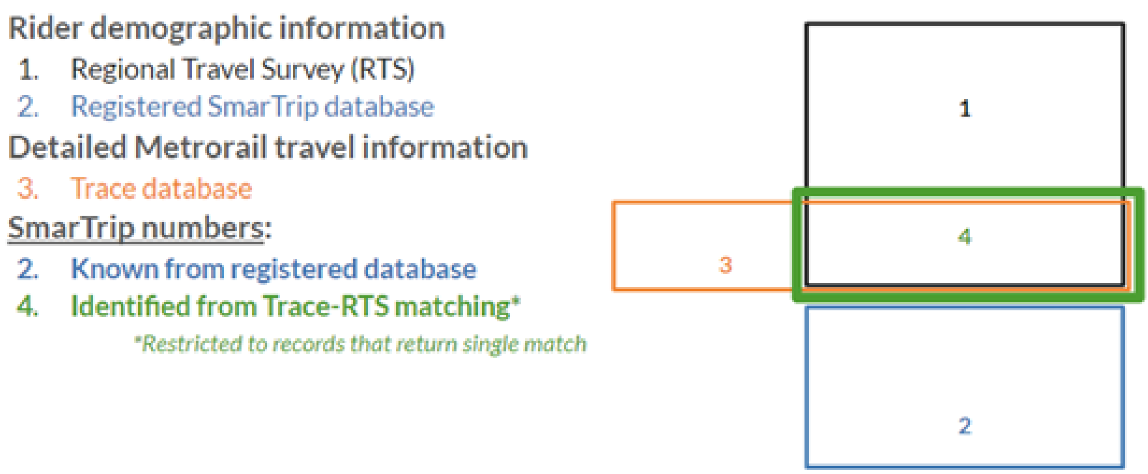

Figure A-1: Schematic of SmarTrip Card Demographic Identification

After linking anonymized SmarTrip card numbers with cardholder demographic information, the travel diary approach used to generate synthetic RTS data in Chapter 4 was replicated. The time period selected for analysis post-COVID runs from January 2021 through September 2022. The original Regional Travel Survey spanned only 14 months, but a longer period was chosen post-pandemic so that a larger sample of Metrorail travel could be generated considering lesser transit use post-COVID. Transit-dependent Metrorail riders who remained active throughout the pandemic would be captured by a smaller analysis window. However, the majority of riders who stopped using transit after the onset of COVID-19 returned slowly, and some still remain inactive. Figure A-2 helps illustrate the importance of a longer analysis window by showing the first month after January 2021 that each individual SmarTrip card with associated demographics traveled on Metrorail. Ultimately, by the end of the analysis window, slightly over 2,000 cards with known demographics had used Metrorail at least once. Of this cohort, roughly 1 out of 5 were already riding by January 2021. By May 2021, 50% of all cards that ended up returning had taken a trip. Nevertheless, riders continued to trickle in and slowly return to Metrorail over the entire period, with each calendar month witnessing at least some riders return.

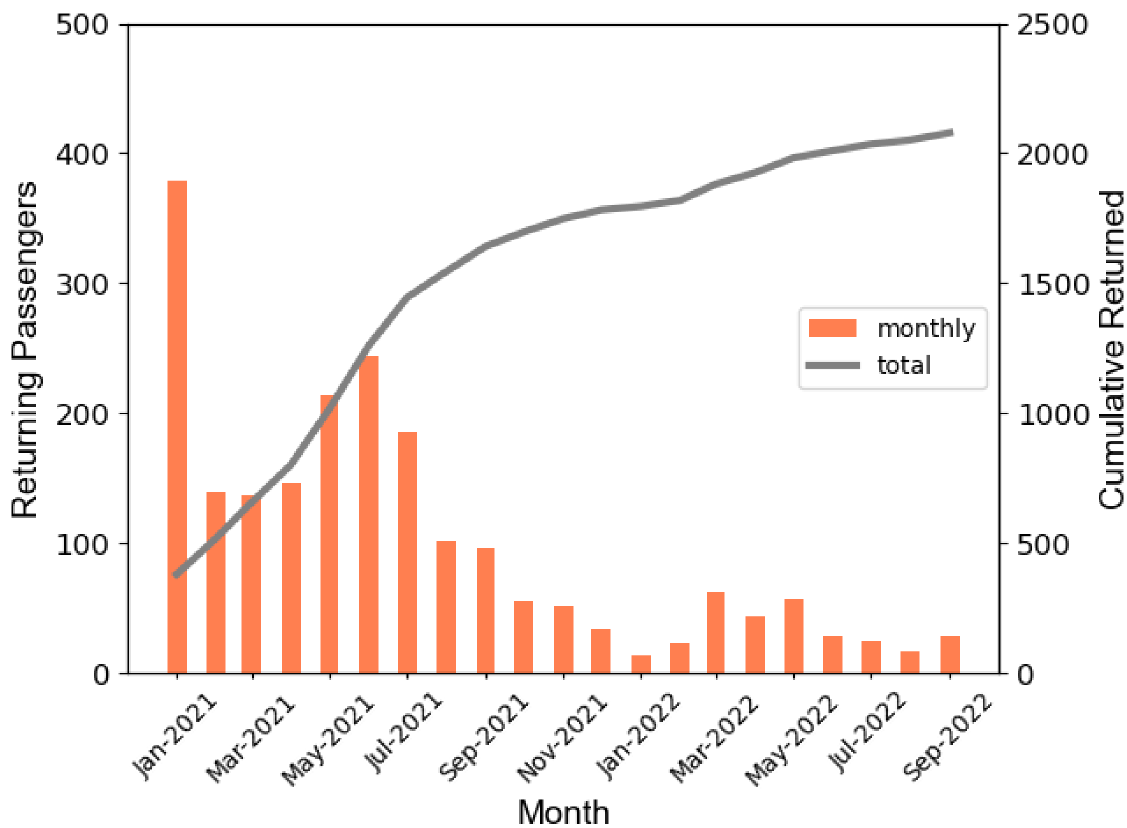

Figure A-2: First Active Month in Analysis Period for Card with Known Demographics

Using these roughly 2,000 SmarTrip account holders with known demographics that used Metrorail to travel between January 2021 and September 2022, synthetic travel diary entries were generated from WMATA's Trace database. The travel diary format mimicked the synthetic data generation described in Chapter 4 and used during the RTS period. In brief, each SmarTrip card holder with known demographics and less than 8 registered SmarTrip cards was assigned a random weekday on which they traveled during the analysis window. On that random weekday, all transit use records for the card holder were queried from the Trace database. The resulting dataset comprises individual Metrorail trips with associated rider demographic information, allowing the methodology used for the pre-COVID fare equity analysis to be reimplemented. The post-COVID compiled trips dataset is, however, much smaller than the one created during the RTS period. Table A.1 below shows a comparison between the two periods.

	RTS Period	Post-COVID
Unique Riders	6,174	2,007
Trips	12,129	3,710
Origin Stations	91	91
Destination Stations	91	91
Unique Station Pairs	1,973	1,486
OD Combinations (*Potential*)	48.2%	36.3%
OD Combinations (*Weighted by Travel Volume*)	86.5%	69.5%

Table A.1: Pre- and Post-COVID Data Sample Summary Statistics

As shown, the post-COVID sample size has roughly $\frac{1}{3}$ as many unique riders and 30% as many trips. All of the stations remain present as both an origin and a destination, but the number of represented station pairs is lower. Further, when the station pairs are weighted by actual travel volumes between each pair, the post-COVID data represents just under 70% of all Metrorail travel during the period. Post-COVID, the data generated by this approach to rider demographics seems less comprehensive, and likely somewhat less representative of overall ridership, than the comparable dataset compiled for the RTS period. Nevertheless, it is still a relatively large sample to explore.

However, one important indicator that this approach results in a less representative dataset post-COVID can be observed in the average travel distance. The average travel distance for a journey in this post-COVID dataset is just over 6 miles, but the average travel distance for all Metrorail trips taken during the same period is 5.5 miles. The trips with known demographic information are, on average, 0.5 miles longer than systemwide average trips. This seemingly small difference is important in the context of distance-based fare analysis, however. It amounts to as much as $0.16 per trip under the current system, meaning the average fare paid by these riders with known demographic information is measurably different than the overall average fare. When evaluating alternate fare structures, the magnitude of this discrepancy could grow or decrease.

The post-COVID synthetic dataset was then weighted along three dimensions using iterative proportional fitting, following the same procedure used in Chapter 4. Weights were calculated for each individual trip such that the results match the income and race values from the 2016 WMATA rail ridership survey. While the demographic information from that survey also comes from a period before the COVID-19 pandemic, it is the closest data to being contemporaneous. The ground truth on rider demographics may have shifted somewhat, but the best observations of that ground truth still come from prior to the pandemic. On the other hand, weighting also accounted for peak and off-peak trips. This dimension was weighted using the observed peak ratio of all trips taken on Metrorail during the analysis window between January 2021 and September 2022. After calculating weights, and applying those weights to each individual trip, the post-COVID travel patterns of rider segments as defined by race and income are observable. A summary of average travel taken by the rider segments is shown in Table A.2.

The observed travel differences across demographic groups are insightful, and largely remain consistent with findings from the pre-COVID RTS period. Low-income

Weighted Metrics	All riders	People of Color: Below Livable Wage	People of Color: Above Livable Wage	White: Below Livable Wage	White: Above Livable Wage
Trips	3,710	390	1,372	122	1,826
Pct. of total		10.5%	37.0%	3.3%	49.2%
Average distance *(Crow-flies miles)*	6.06	5.50	6.31	5.75	6.02
Travel circuity *(Track miles / crow-flies miles)*	1.28	1.37	1.31	1.27	1.27
Trips by period					
Peak	61.6%	56.1%	62.1%	56.2%	62.9%
Off-peak	38.4%	43.9%	37.9%	43.8%	37.2%
Average fare ($)	3.23	3.14	3.32	3.06	3.21

Table A.2: Weighted Post-COVID Travel by Rider Demographics

riders of all races continue to take more trips in off-peak periods than their higher-income peers. This is true even though overall the ratio of peak travel has decreased. Low-income riders of color continue to take the most circuitous trips on the Metrorail network, and will therefore be the most negatively impacted by the choice of composite miles as the distance basis used to measure and calculate fares. Riders of color earning above the livable wage continue to have the longest observed trips. While the pre-COVID and post-COVID findings are directionally consistent, there are nevertheless some slight shifts. In particular, low-income white riders are found to take trips that are equally circuitous with their higher-earning peers post-COVID. Prior to COVID, low-income white riders had trips that were marginally more circuitous than trips taken by white riders with incomes above the livable wage.

In addition to the sampling approach described above, a second approach to inferring rider demographic information for individual Metrorail trips was developed based on previous work done by Cambridge Systematics in cooperation with WMATA [70]. The Cambridge Systematics team developed demographic probabilities for individual WMATA journeys based on a hierarchical evaluation of trip characteristics and AFC information. Up to six aspects of a WMATA journey are considered to determine the probability that a traveler has certain demographic attributes. These six aspects of a trip used to estimate probabilities are, in order: transit mode, origin, destination, fare class, time of travel, and multimodal transfers. The aspects available from the data are ordered based upon their explanatory power regarding rider demographics. From these observable trip characteristics, demographic probability weights were generated that matched actual travel records available using WMATA's Trace database, one bus passenger survey from 2018, and two rail passenger surveys from 2016.

From the previously developed demographic probability model, a methodology to apply these probabilities to understand travel patterns of demographic groups was finalized and deployed by Nineveh O'Connell[1]. This methodology apportions fractional transit trips to predefined rider demographic group based on the calculated demographic probabilities. As applied to this fare equity analysis, each Metrorail journey is split into four unequal pieces. Those unequal pieces are then assigned to rider groups defined by income and race. Table A.3 illustrates how this would be done for a single transit journey. For the sample record shown, the probability of the rider being low-income and white is 20%; each of the remaining combinations of rider income and race also have probabilities attached. Together, the probabilities sum to 100%. In this way, the entire trip is accounted for even though it is split up and then assigned as a fractional probability to the rider segments.

This approach to inferring demographic information for Metrorail trips was used to create an alternative second post-COVID dataset to analyze. The approach was extended to all trips for which demographic probabilities could be estimated during a

[1]Current Masters of City Planning student at MIT

CARD_ID	TRIP_ID	*Other Variables*	LOW_INCOME	WHITE	WEIGHT
3733673	12028	...	Y	Y	0.20
3733673	12028	...	Y	N	0.10
3733673	12028	...	N	Y	0.45
3733673	12028	...	Y	N	0.25

Table A.3: Probability-based Trip Apportionment Example

one-week analysis window. Demographic probabilities were estimated at the level of individual SmarTrip cards, as opposed to being estimated for each individual transit journey. This technique was decided upon to ensure model consistency, as demographics belong to the individual making a trip; demographics do not belong to the trip itself. The most common originating station was used as the basis to estimate demographic probabilities for individual SmarTrip card numbers. Resulting estimated demographic probabilities were then applied to all trips taken by that SmarTrip card during the week beginning Monday, October 3, 2022 and ending Friday, October 7, 2022 [71].

This second approach to assembling a dataset of Metrorail trips with inferred rider demographics has the benefit of generating a much larger sample, ensuring that travel patterns within the sample are representative of all Metrorail trip making. On the other hand, the demographic information linked to trips within this second sample is less certain in comparison to the rider-reported demographics available in the original, matched travel dataset. Nevertheless, both datasets together can be used to support and reinforce each other. Comparing findings results from both approaches provides a convincing robustness check with each providing cross-verification for the other.

After applying both approaches, the summary travel statistics for Metrorail rider segments defined by race and income are generally consistent with the values shown in Table A.2. The differences between rider demographic groups are directionally uniform, although some magnitudes vary between the two approaches. Figure A-3 illustrates this using the average circularity for trips as calculated for four rider segments. As can be seen, the measured differences between all four groups are smaller for the larger, probability-based sample. Conversely, the smaller sample utilizing rider-reported demographics finds much more circuitous travel for low-income riders of color. Nevertheless, the ranking of trip-circuity among the groups is the same regardless of approach. Table 4.3 in chapter 4 shows a similar chart displaying the average trip circularity during the RTS period; comparing the results before and after COVID-19 similarly indicates that while trip circuity may have changed slightly, it is the same groups taking indirect trips in both periods.

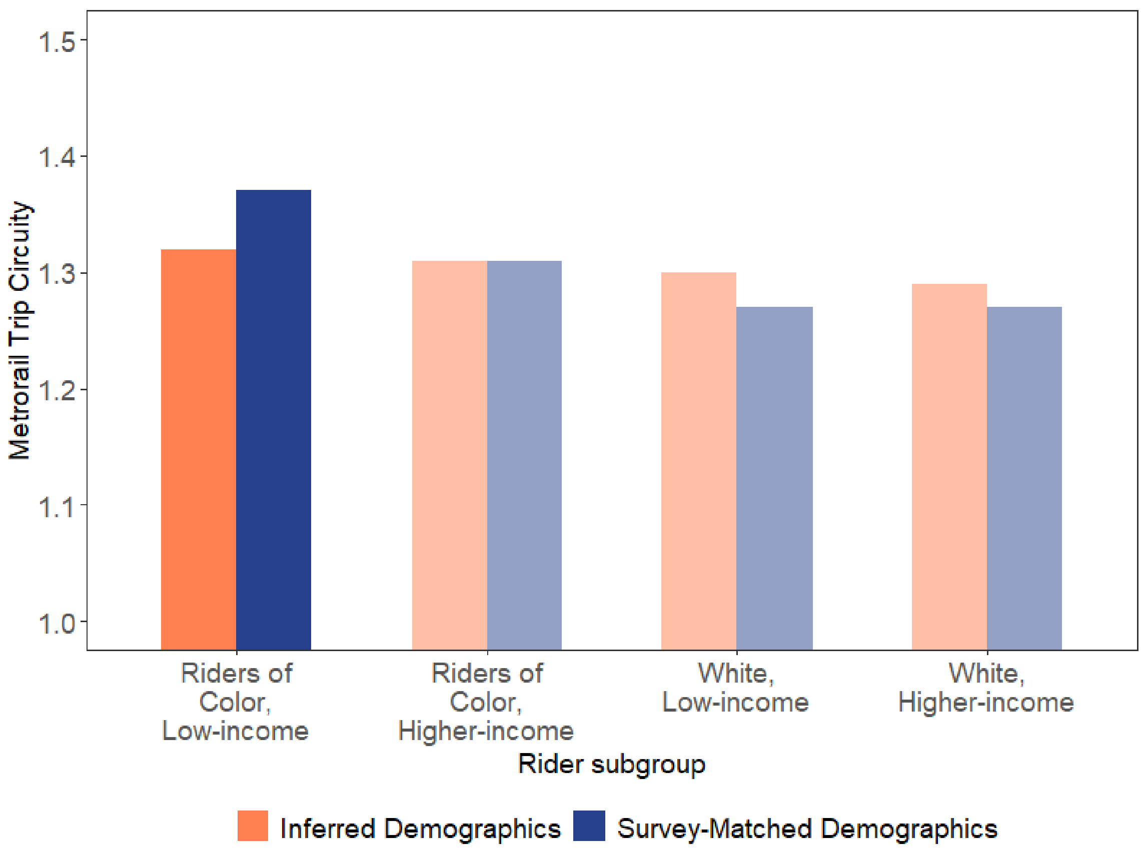

Figure A-3: Post-COVID Metrorail Travel Circuity

Regardless of approach to linking travel and rider demographics, after compiling the travel dataset revenue-neutral rates for each, alternate fare structure must be recalculated for the post-COVID period. The procedure used to calculate these rates during the pre-pandemic RTS period is described in Chapter 4, and the same methodology was reimplemented using systemwide post-COVID travel. Utilizing post-COVID travel to calculate alternate fare rates results in little observed change compared to the pre-COVID period.

Using the recalculated post-pandemic revenue neutral fares, fare changes under each alternate scenario were estimated for all journeys with associated demographic information. Elasticity estimates were then applied[2] using rider and trip characteristics to capture responses to any changes in the overall fare levels. With these estimated changes to average fare and travel volumes for each rider segment, the equity of each structure can be understood given current post-pandemic travel patterns. Further, the equity of the current distance-based fare structure can be contextualized.

Although revenue neutral fare rates for all of the alternate fare structures were recalculated post-COVID, ultimately four representative fare structures were chosen as the focus of this updated analysis. One fare structure each was chosen from scenarios 1-4. The representative fare structures were chosen to capture the major equity impacts of the scenario within which they were situated. The current fare structure, represented by Scenario 1a, was chosen from the composite-miles distance-based fares. Scenario 2a, where crow-flies miles is used as the distance basis, was selected as a counterpoint. Among the zonal implementations 3c was selected; this zonal structures utilizes a boundary-crossings basis with peak differentiation. As a reminder, this boundary-crossings implementation charges a higher rate for the first 5 zonal crossings and a lower rate for any further crossings. Finally, Scenario 4b was selected, which uses a simple flat fare with no peak-pricing differentiation. One additional fare structure was included, but only for post-COVID travel. Labeled Scenario 5, this new structure represents WMATA's proposed Metrorail fare structure for fiscal year 2023. This proposal would remove peak-pricing differentiation, retain composite-miles as a distance-basis, use the current off-peak minimum fare as the

[2]For any questions about methodology, please reference Chapter 4 where the process is described in detail

minimum at all times of day, and raise the maximum fare per trip to $6.50.

The four representative structures, as well as WMATA's proposed fare structure, are represented in Figures A-4 and A-5. As shown, both of the approaches used to compile a sample of Metrorail travel with linked rider demographic information find largely similar fare equity impacts on rider groups defined by race. The largest discrepancy between approaches occurs when looking at the revenue-neutral flat fare (4a) and WMATA proposal (5). The updated sample with self-reported demographics shows a proportionally smaller decrease in travel volumes given the new fare rates. This discrepancy can be ascribed to the dimension in which this smaller sample is not representative of overall Metrorail travel: journey length. As previously described, the trips in this sample are slightly longer than the systemwide average. This means that these riders are already paying slightly higher than average fares. A systemwide revenue-neutral flat fare would therefore be relatively less costly to the trips in this sample, meaning that any expected decreases in travel volumes would be less drastic given this sample of trips. Similarly, the removal of peak pricing and the increase in the maximum fare in Scenario 5 disproportionately impacts off-peak trips and long-distance trips. Regardless which of the two approaches is used to create a sample of demographic-linked travel, the findings are consistent that riders of color take the longest trips. For that reason, this segment of riders is expected to see a larger reduction in travel volumes under Scenario 5 when compared to white riders. The most obvious discrepancy between the techniques for identifying trip demographics occurs when comparing alternate Scenarios 4b and 5. The larger, inferred demographic approach shows a flat fare (4b) depresses travel more than WMATA's proposed structure (5). The smaller sample, using rider-reported demographics, finds that Scenario 5 would have a much larger chilling effect on travel than a simple flat fare (4b).

The results tell a similar fare equity story when considering riders defined by income as opposed to race. Scenario 2a, which uses crow-flies miles as a distance basis,

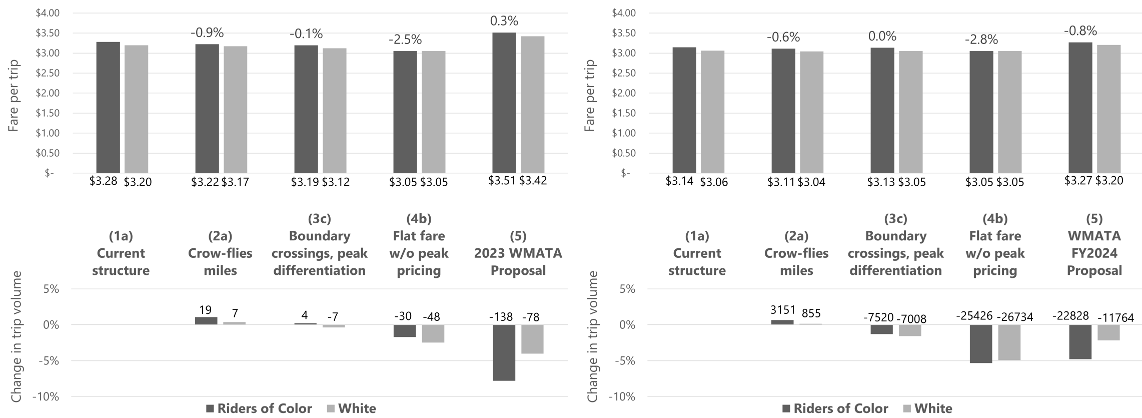

(a) Matched Demographics (b) Inferred Demographics

Figure A-4: Differences Across Identification Technique by Race

results in slight benefits to fare equity and travel volumes. The zonal implementation considered (3a) has negative impacts on ridership and on equity. The two scenarios that would depress travel volumes to the largest extent are the revenue-neutral flat fare (4b) and WMATA's proposed fare structure (5). Both would also have a negative impact on equity, with the larger impacts, in proportional terms, expected to be felt by low-income riders. As when considering race, the two approaches to creating a demographic-linked travel sample come to different conclusions about which of these two scenarios (4b, 5) is worse for ridership. WMATA's proposed fare structure (5) is expected to result in the largest travel volume impacts when considering the smaller sample containing rider-reported demographics. The larger sample, with inferred rider demographics, indicates that it is the flat fare (4b) that will cause the largest travel volume decreases.

Both approaches to generating a post-COVID sample of Metrorail travel with rider demographic information offer provide advantages and drawbacks. The benefit of using both together is cross-verified results. While the estimated equity metrics change depending on which sample of Metrorail travel is used, the story does not. The differential impacts between rider groups is similar, and the equity impacts of each representative fare structure remain consistent.

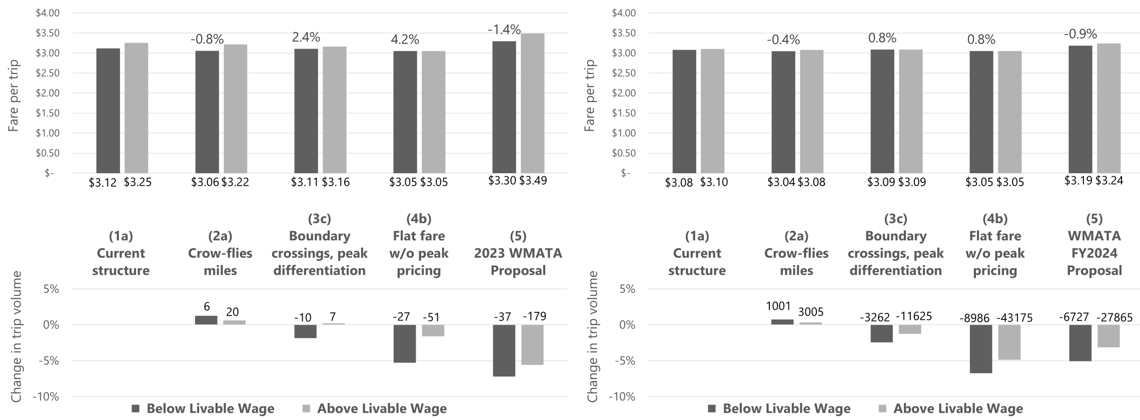

(a) Matched Demographics (b) Inferred Demographics

Figure A-5: Differences Across Identification Technique by Income

Another important question to consider with these post-COVID findings in hand is whether, and to what extent, fare equity has changed since the pre-pandemic period. To explore this, pre- and post-COVID fare equity calculated using the smaller sample of rider-reported demographic information is used. Results from either of the two approaches to link demographics and travel post-COVID could be used as the findings are consistent. The rider-reported sample is chosen because it matches the approach used prior to COVID. As a reminder, pre-COVID fare equity was measured using travel records during the time period covered by the most recent Regional Travel Survey. Comparing the pre-pandemic and post-pandemic equity metrics is a useful approach to capturing any fare equity changes. In particular, it is important to identify if the equity impacts found in any of the alternate fare scenarios have flipped, or whether they have continued to remain similar in both periods.

Figure A-6 shows the four representative fare structures (excluding WMATA's fiscal year 2023 proposal) and the corresponding equity metrics calculated for both pre- and post-COVID time periods. The first key takeaway, when considering riders as defined by race, is that the equity benefits of a crow-flies distance structure remain unchanged. The fare differential between rider groups decreases, a higher percentage of trips are generated among riders of color, and there are no decreases in trip volumes among white riders. On the other hand, the findings regarding the zonal implementation (3c) have changed somewhat. Prior to COVID this structure was expected to decrease travel among all groups, with a slightly larger proportional decline among white riders. Post-COVID, this zonal implementation shows at most a marginal impact on ridership and no impact on the fare differential between groups. The impacts of a revenue-neutral flat fare (4b) have remained constant across periods; decreased trip volumes and larger impacts on white riders.

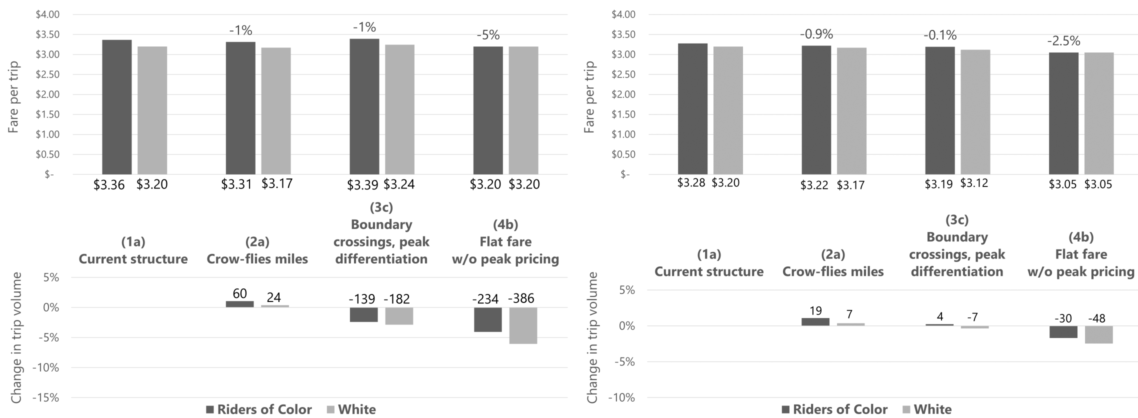

(a) Pre-COVID (b) Post-COVID

Figure A-6: Changes Across Time by Rider Race

Figure A-7 shows the same four representative fare structures, but with riders segmented based on their income level instead of their race. Here the pre-pandemic and post-pandemic equity findings are once again very similar. A crow-flies distance basis (2a) would benefit low-income riders who would experience a decrease in the fare differential and an increase in travel volumes. Higher-income riders would also potentially benefit from a marginal increase in trip volumes. A zonal system that charges for each zonal boundary crossed (3c) once again has disproportionately large negative impacts on low-income riders, who would reduce travel volumes at higher levels. Finally, revenue-neutral flat fares (4b) are expected to lower travel volumes, with low-income riders forgoing trips at much higher levels.

It is reassuring that all three approaches to associating rider demographics and

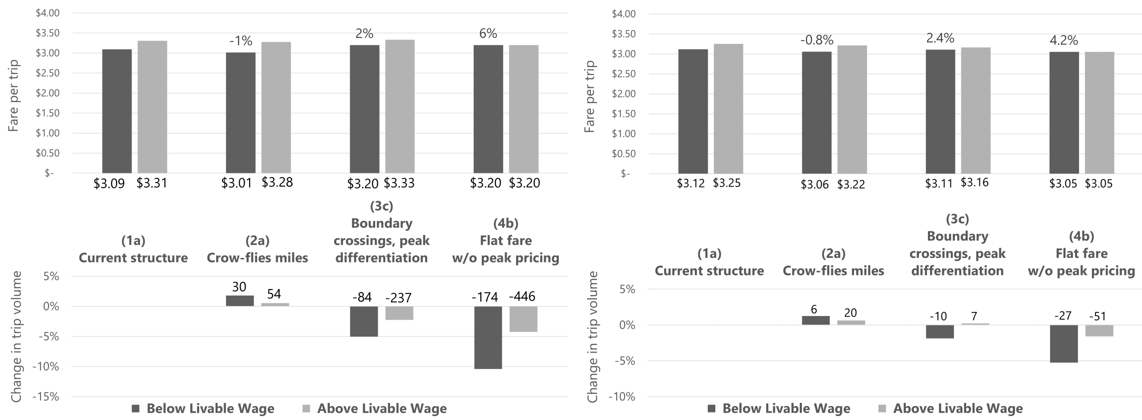

(a) Pre-COVID (b) Post-COVID

Figure A-7: Changes Across Time by Rider Income

Metrorail travel come to the same recommendations when evaluating the equity of WMATA's distance-based fare structure. The aspects of the fare structure that can be leveraged to improve fare equity remain stable regardless of time period or travel sample. Flat fares (4b) are bad for travel volumes, as the revenue-neutral rate is much higher than the current cost for most short-distance trips, which would result in many of these trips being foregone. The decrease in travel volumes would impact rider groups as defined by race and income disparately, with higher travel volume decreases for white riders and also for low-income riders. Crow-flies as a distance basis (2a) shows both equity and ridership benefits across all approaches and time periods, suggesting this a simple and straightforward way to improve fare equity. Zonal implementations based on the number of boundaries crossed (3c) provide no strong benefits to riders of color, but would negatively impact low-income riders.

Ultimately, updating the analysis to consider post-COVID travel affirms the findings of chapter 4. In the aggregate, low-income riders take less direct trips on Metrorail, so structures that penalize indirect travel have negative outcomes for fare equity. Such penalties for indirect travel include using composite-miles (or track-miles) as a distance basis or charging for each zonal boundary crossed. At a high-level the equity of a distance-based structure within the geography of Washington, DC is dependent on whether rider groups are defined by race or by income. This is true both before and after the onset of the COVID-19 pandemic. There are avenues to improve equity without harming any other disadvantaged groups by tweaking the existing system.

Bibliography

[1] Thomas Doyle. Public transportation ridership rises to more than 70 percent of pre-pandemic levels. URL https://www.apta.com/news-publications/press-releases/releases/public-transportation-ridership-rises-to-more-than-70-percent-of-pre-pandemic-levels/.

[2] Chicago Transit Authority. Annual ridership report, 2019, . URL https://www.transitchicago.com/assets/1/6/2019_Annual_Ridership_Report.pdf.

[3] Chicago Transit Authority. Monthly ridership report, October 2022, . URL https://www.transitchicago.com/assets/1/6/Monthly_Ridership_2022-10(Final).pdf.

[4] CTA L system map. URL https://www.transitchicago.com/assets/1/6/ctamap_Lsystem.png.

[5] Fare recovery ratio: What it is and why it must be reformed, . URL https://www.rtachicago.org/blog/2023/01/05/fare-recovery-ratio-what-it-is-and-why-it-must-be-reformed.

[6] Nicole Nutter. Quarterly financial and performance report. URL https://www.rtachicago.org/uploads/files/general/Transit-Funding/Q3-2022-Quarterly-Financial-Report.pdf.

[7] WMATA ridership snapshot - August 2022, . URL https://www.wmata.com/initiatives/ridership-portal/upload/August-2022-Ridership-Snapshot.pdf.

[8] WMATA ridership snapshot - August 2019, . URL https://www.wmata.com/initiatives/ridership-portal/upload/August-2019-Ridership-Snapshot.pdf.

[9] Garrett Shrode. The mass transit fiscal cliff: Estimating the size and scope of the problem. URL https://www.enotrans.org/article/the-mass-transit-fiscal-cliff-estimating-the-size-and-scope-of-the-problem/.

[10] WMATA rail system map. URL https://www.wmata.com/schedules/maps/upload/2022-System-Map.pdf.

[11] Jarrett Walker. Can fares be fair? URL https://link.springer.com/book/10.5822/978-1-61091-174-0.

[12] Nihit Jain. Assessing the impact of recent fare policy changes on public transport demand in London. URL http://hdl.handle.net/1721.1/66866.

[13] Andrew Stuntz. Transit fare policy: Use of automated data to improve incremental decision making. URL https://dspace.mit.edu/handle/1721.1/119275.

[14] Stephanie Lotshaw and Kirk Hovenkotter. A fare framework: How transit agencies can set fare policy based on strategic goals. URL https://transitcenter.org/how-transit-agencies-are-getting-strategic-about-fare-policy/.

[15] Robert Cervero. Flat versus differentiated transit pricing: What's a fair fare? 10(3):211–232, . ISSN 0049-4488, 1572-9435. doi: 10.1007/BF00148459. URL http://link.springer.com/10.1007/BF00148459.

[16] Robert Cervero. Efficiency and equity implications of transit fare policies. 1 (1):48–48, . ISSN 0739-456X. doi: 10.1177/0739456X8100100120. URL https://doi.org/10.1177/0739456X8100100120. Publisher: SAGE Publications Inc.

[17] Candace Brakewood and George Kocur. Modeling transit rider preferences for contactless bank cards as fare media: Transport for London and the Chicago, Illinois, Transit Authority. 2216(1):100–107. ISSN 0361-1981. doi: 10.3141/2216-11. URL https://doi.org/10.3141/2216-11. Publisher: SAGE Publications Inc.

[18] Raymond Chan, Maulik Vaishnav, Scott Wainwright, Paul Murray, and Alex Cui. Data-driven opportunities from an account-based fare payment system. 2672(8):536–545. ISSN 0361-1981. doi: 10.1177/0361198118798242. URL https://doi.org/10.1177/0361198118798242. Publisher: SAGE Publications Inc.

[19] Nigel H.M. Wilson, Jinhua Zhao, and Adam Rahbee. The potential impact of automated data collection systems on urban public transport planning. In *Schedule-Based Modeling of Transportation Networks: Theory and applications*, Operations Research/Computer Science Interfaces Series, pages 1–25. Springer US. ISBN 978-0-387-84812-9. doi: 10.1007/978-0-387-84812-9_5. URL https://doi.org/10.1007/978-0-387-84812-9_5.

[20] Jason B. Gordon, Harilaos N. Koutsopoulos, Nigel H. M. Wilson, and John P. Attanucci. Automated inference of linked transit journeys in London using fare-transaction and vehicle location data. 2343(1):17–24, . ISSN 0361-1981. doi: 10.3141/2343-03. URL https://doi.org/10.3141/2343-03. Publisher: SAGE Publications Inc.

[21] Jason B. Gordon, Haris N. Koutsopoulos, and Nigel H. M. Wilson. Estimation of population origin–interchange–destination flows on multimodal transit networks.

90:350–365, . ISSN 0968-090X. doi: 10.1016/j.trc.2018.03.007. URL https://www.sciencedirect.com/science/article/pii/S0968090X18303255.

[22] Gabriel E. Sánchez-Martínez. Inference of public transportation trip destinations by using fare transaction and vehicle location data: Dynamic programming approach. 2652(1):1–7. ISSN 0361-1981. doi: 10.3141/2652-01. URL https://doi.org/10.3141/2652-01. Publisher: SAGE Publications Inc.

[23] Catherine Vanderwaart, John P. Attanucci, and Frederick P. Salvucci. Applications of inferred origins, destinations, and interchanges in bus service planning. 2652(1):70–77. ISSN 0361-1981. doi: 10.3141/2652-08. URL https://doi.org/10.3141/2652-08. Publisher: SAGE Publications Inc.

[24] R Balcombe, R Mackett, N Paulley, J Preston, J Shires, H Titheridge, M Wardman, and P White. The demand for public transport: a practical guide. URL https://discovery.ucl.ac.uk/id/eprint/1349/1/2004_42.pdf.

[25] Todd Litman. Transit price elasticities and cross - elasticities. 7(2):37–58. ISSN 1077-291X, 2375-0901. doi: 10.5038/2375-0901.7.2.3. URL http://scholarcommons.usf.edu/jpt/vol7/iss2/3/.

[26] Paul Schimek. Dynamic estimates of fare elasticity for U.S. public transit. 2538 (1):96–101. ISSN 0361-1981, 2169-4052. doi: 10.3141/2538-11. URL http://journals.sagepub.com/doi/10.3141/2538-11.

[27] Cynthia Chen, Don Varley, and Jason Chen. What affects transit ridership? a dynamic analysis involving multiple factors, lags and asymmetric behaviour. 48 (9):1893–1908. ISSN 0042-0980, 1360-063X. doi: 10.1177/0042098010379280. URL http://journals.sagepub.com/doi/10.1177/0042098010379280.

[28] Adam Rosenfeld. Driving change: How workplace benefits can nudge solo car commuters toward sustainable modes. URL https://dspace.mit.edu/handle/1721.1/117826.

[29] Yi Hong. Transition to smart card technology: how transit operators can encourage the take-up of smart card technology. URL http://hdl.handle.net/1721.1/38237.

[30] Yan Liu, Siqin Wang, and Bin Xie. Evaluating the effects of public transport fare policy change together with built and non-built environment features on ridership: The case in South East Queensland, Australia. 76:78–89, . ISSN 0967070X. doi: 10.1016/j.tranpol.2019.02.004. URL https://linkinghub.elsevier.com/retrieve/pii/S0967070X18307194.

[31] Richard L. Oram, United States, Charles River Associates, and R.L. Oram Associates. *Deep discount fares: building transit productivity with innovative pricing*. U.S. Department of Transportation, Urban Mass Transportation Administration, Office of Budget and Policy ; National Technical Information Service, distributor. URL https://catalog.hathitrust.org/Record/012354530.

[32] Texas A&M Transportation Institute. Commuter tax benefits policy brief. URL https://docs.google.com/viewerng/viewer?url=http://tti.tamu.edu/documents/TTI-2015-5.pdf&hl=en_US.

[33] Tony Dutzik, Elizabeth Berg, Alana Miller, and Rachel Cross. Who pays for parking? URL https://transitcenter.org/wp-content/uploads/2017/09/TC-Who-Pays-For-Parking-Publication-20170912-1.pdf.

[34] Transportation Research Board and National Academies of Sciences, Engineering, and Medicine. *Analyzing the Effectiveness of Commuter Benefits Programs*. Transportation Research Board. ISBN 978-0-309-43152-1. doi: 10.17226/21979. URL https://www.nap.edu/catalog/21979. Pages: 21979.

[35] Paola Carolina Bueno, Juan Gomez, Jonathan R. Peters, and Jose Manuel Vassallo. Understanding the effects of transit benefits on employees' travel behavior: Evidence from the New York-New Jersey region. 99:1–13. ISSN 0965-8564. doi: 10.1016/j.tra.2017.02.009. URL https://www.sciencedirect.com/science/article/pii/S0965856416303986.

[36] Asha Weinstein Agrawal, Evelyn A Blumenberg, Sarah Abel, Gregory Pierce, and Charles N Darrah. Getting around when you're just getting by: The travel behavior and transportation expenditures of low-income adults. page 126.

[37] Songhua Hu and Peng Chen. Who left riding transit? examining socioeconomic disparities in the impact of COVID-19 on ridership. 90:102654. ISSN 1361-9209. doi: 10.1016/j.trd.2020.102654. URL https://www.sciencedirect.com/science/article/pii/S1361920920308397.

[38] Transit. Who's left riding public transit? a COVID data deep-dive. URL https://site.transitapp.com/whos-left-riding-public-transit-hint-it-s-not-white-people-d43695b3974a/.

[39] Luyu Liu, Harvey J. Miller, and Jonathan Scheff. The impacts of COVID-19 pandemic on public transit demand in the United States. 15(11):e0242476, . ISSN 1932-6203. doi: 10.1371/journal.pone.0242476. URL https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0242476. Publisher: Public Library of Science.

[40] Karla Therese L Sy, Micaela E Martinez, Benjamin Rader, and Laura F White. Socioeconomic disparities in subway use and COVID-19 outcomes in New York City. 190(7):1234–1242. ISSN 0002-9262. doi: 10.1093/aje/kwaa277. URL https://doi.org/10.1093/aje/kwaa277.

[41] Rebecca Brough, Matthew Freedman, and David C. Phillips. Understanding socioeconomic disparities in travel behavior during the COVID-19 pandemic. 61(4):753–774. ISSN 1467-9787. doi: 10.1111/jors.12527. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/jors.12527. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/jors.12527.

[42] Federal Transit Administration. Title VI requirements and guidelines for federal transit administration recipients. Circular 4702.1B. URL https://www.transit.dot.gov/sites/fta.dot.gov/files/docs/FTA_Title_VI_FINAL.pdf.

[43] TransitCenter. Going above and beyond: Advice to USDOT on measuring transportation equity, . URL https://transitcenter.org/going-above-and-beyond-advice-to-usdot-on-measuring-transportation-equity/.

[44] TransitCenter. *Equity in Practice: A Guidebook for Transit Agencies*. TransitCenter, . URL https://transitcenter.org/publication/equity-in-practice-a-guidebook-for-transit-agencies/.

[45] Brian D. Taylor. *The Geography of Urban Transportation Finance*. ISBN 978-1-59385-055-5. URL https://trid.trb.org/view/756098.

[46] Cornelius Nuworsoo, Elizabeth Deakin, and Aaron Golub. Equity impacts of transit fare proposals: A case study of AC transit. 47(3):23–43.

[47] Anne E. Brown. Fair fares? how flat and variable fares affect transit equity in Los Angeles. 6(4):765–773. ISSN 2213-624X. doi: 10.1016/j.cstp.2018.09.011. URL https://www.sciencedirect.com/science/article/pii/S2213624X17302298.

[48] Steven Farber, Keith Bartholomew, Xiao Li, Antonio Páez, and Khandker M. Nurul Habib. Assessing social equity in distance based transit fares using a model of travel behavior. 67:291–303. ISSN 0965-8564. doi: 10.1016/j.tra.2014.07.013. URL https://www.sciencedirect.com/science/article/pii/S0965856414001785.

[49] Isak Rubensson, Yusak Susilo, and Oded Cats. Is flat fare fair? equity impact of fare scheme change. 91:48–58. ISSN 0967-070X. doi: 10.1016/j.tranpol.2020.03.013. URL https://www.sciencedirect.com/science/article/pii/S0967070X19306808.

[50] Ahmed El-Geneidy, David Levinson, Ehab Diab, Genevieve Boisjoly, David Verbich, and Charis Loong. The cost of equity: Assessing transit accessibility and social disparity using total travel cost. 91:302–316. ISSN 0965-8564. doi: 10.1016/j.tra.2016.07.003. URL https://www.sciencedirect.com/science/article/pii/S0965856416305924.

[51] Matt Johnson. Why a flat fare is a bad idea for metro. . URL https://ggwash.org/view/12882/why-a-flat-fare-is-a-bad-idea.

[52] Matt Johnson. Metro's union wants a flat fare. here's why that's a bad idea. . URL https://ggwash.org/view/62915/metros-union-wants-a-flat-fare-heres-why-thats-a-bad-idea.

[53] Matt Johnson. Should metro switch to zones? (no.). . URL https://ggwash.org/view/62979/should-metro-switch-to-zones-nope.

[54] Michael Eichler. Metrorail fares are complicated, for good. URL https://planitmetro.com/2013/12/06/metrorail-fares-are-complicated-for-good/.

[55] Luke J Juday. The changing shape of American cities. URL https://demographics.coopercenter.org/sites/demographics/files/ChangingShape-AmericanCities_UVA-CooperCenter_February2015.pdf.

[56] Alan Ehrenhalt. *The Great Inversion and the Future of the American City*. Vintage, illustrated edition. ISBN 978-0-307-47437-7.

[57] Aaron Renn. The new donut. URL https://www.aaronrenn.com/2014/09/14/the-new-donut/.

[58] Jeff Allen and Steven Farber. Suburbanization of transport poverty. pages 1–18. ISSN 2469-4452, 2469-4460. doi: 10.1080/24694452.2020.1859981. URL https://www.tandfonline.com/doi/full/10.1080/24694452.2020.1859981.

[59] 2017/2018 Regional Travel Survey. Technical report, Metropolitan Washington Council of Governments, January 2021.

[60] Amy K. Glasmeier. Living wage calculator for the District of Columbia. URL https://web.archive.org/web/20171003093917/https://livingwage.mit.edu/states/11.

[61] National Center for Smart Growth Research and Education. Origin-destination land use ridership model presentation.

[62] National Center for Smart Growth Research and Education. Origin-destination land use ridership model for fare policy analysis. URL https://planitmetro.com/wp-content/uploads/2015/12/151215_LandUseRidershipModel_FinalReport.pdf.

[63] Hugh Clark. Who rides public transportation. URL https://www.apta.com/wp-content/uploads/Resources/resources/reportsandpublications/Documents/APTA-Who-Rides-Public-Transportation-2017.pdf.

[64] CTA welcomes back customers with "when you're ready, we're ready!" campaign, . URL https://www.transitchicago.com/cta-welcomes-back-customers-with-when-youre-ready-were-ready-campaign/.

[65] New agreement will streamline and enhance CTA and Pace unlimited-ride passes in 2023, . URL https://www.transitchicago.com/new-agreement-will-streamline-and-enhance-cta-and-pace-unlimited-ride-passes-in-2023/.

[66] Jacqueline Serrato, Pat Sier, and Charmaine Runes. Mapping Chicago's racial segregation. URL https://interactive.wttw.com/firsthand/segregation/mapping-chicago-racial-segregation.

[67] C. Philip Cummings, Malcolm Fairhurst, Sarah LaBelle, and Darwin Stuart. Market segmentation of transit fare elasticities. 43(3):407–420. URL https://babel.hathitrust.org/cgi/pt?id=uc1.c104685756&view=1up&seq=432.

[68] F. V. Webster and P. H. Bly. The demand for public transport: Part i. the changing environment in which public transport operates. 1(4):323–351. ISSN 0144-1647, 1464-5327. doi: 10.1080/01441648108716470. URL http://www.tandfonline.com/doi/abs/10.1080/01441648108716470.

[69] Joyce M. Dargay and Mark Hanly. The demand for local bus services in England. 36(1):73–91. ISSN 0022-5258. URL https://www.jstor.org/stable/20053893. Publisher: [London School of Economics, University of Bath, London School of Economics and University of Bath, London School of Economics and Political Science].

[70] Xiao Yun Chang, Scott Boone, Shrikant Fulari, Maximilian Alfaro, Michael Marks, and Michael Eichler. Inference of transit rider demographic profiles using farecard data and rider surveys through hierarchical weighting.

[71] Nineveh O'Connell. Fare equity from inferred post-COVID rider demographics.