Inequitable Inefficiency: A Case Study of Rail Transit Fare Policies

Abstract

Research on transit fare equity often measures equity based on a disparity in the fare per mile paid by different groups of riders. This cost-benefit measurement overlooks the cost sharing nature of transit; as more riders consume a service, the average cost per rider declines. Using an average cost per rider metric to assign trip costs, and origin-destination fare data to estimate trip-level cost recovery through fares, I estimate the spatial and temporal variability of cost recovery across two rail systems, BART and MARTA. I find that cost recovery patterns are spatially monocentric, and that the weekday peak period recovers more of its costs through fares than other time periods. I offer ideas on why these findings appear divergent to past research.

Introduction

Transportation finance policy in the United States treats travel as exogenous and goal-oriented around funding facilities and services to accommodate this travel. Hence, travelers generally do not directly pay for the costs of their travel and, to the extent that they do, what they pay is not proportional to the marginal costs of their travel. While this is broadly documented in research, surprisingly little research has evaluated how transport subsidies — the difference between the cost of providing transport and what travelers pay for the transport they consume — are distributed across space and time. Are travelers in different locations or at different times subsidized more than others? If so, this suggests that certain travel and development patterns may be facilitated through largely opaque and poorly documented transport subsidies.

Research on the incidence of transport subsidies is especially limited when it comes to public transit. This may be due to transit in the United States functioning both as a transportation lifeline for those who cannot or choose not to drive, and an alternative to driving that can reduce the negative externalities of driving (e.g., Meyer and Gomez-Ibanez, 1981 and Fielding, 1995, as referenced by Giuliano, 2005). Hence, with respect to the second objective, rather than upend broader transportation finance policy objectives to internalize travel costs, policy interventions prioritize marginally reducing the negative externalities of automobile travel by encouraging travelers to change their behavior. Luring travelers from driving to transit is one such example of these “second best” solutions. Furthermore, an oft-unstated constraint is that the travel behaviors people have become accustomed to must not be taken away or taxed — the so-called “Do No Harm” ethos in transportation policy (Altshuler and Luberoff, 2004; Altshuler, 2010). Because of these goals of transportation finance, transit is an especially subsidized travel mode in the United States. In 2017, nationwide highway costs were financed 53% by users (through tolls, fuel taxes, and license fees); transit, 21% (through fares) (United States Census Bureau, 2017).

On face value, those who travel more receive more transport subsidy. However, this overlooks the distributional spread of costs and user payments across a transport network. If there is any variability in these — for example, two bus routes with equal costs, but different levels of ridership and fare revenue, so different farebox recovery ratios — this implies that there is *inequitable inefficiency* in the network. Not only do both bus routes fail to recover their respective costs, making them inefficient, but the disparity in cost recovery level makes this network inequitably inefficient.

In a preceding study (Mallett, 2022), I showed that total and average per trip annualized operating costs of rail transit, inclusive of semi-fixed assets, is spatially and temporally variable.

I used the San Francisco Bay Area Rapid Transit District (BART) and Metropolitan Atlanta Rapid Transit Authority (MARTA) as case studies due to their heavy rail technology and service patterns being similar, which allows for a consistent and comparative analysis. In addition, their collection of highly disaggregate trip data (the origin, destination, and fare of every trip in the systems are recorded) and different fare structures allows for an assessment of trip-level subsidies, including whether this varies by fare structure. **Figures 1** and **2** show the BART and MARTA system maps, respectively, in operation during most of fiscal year 2019 (FY19) — July 1, 2018, to June 30, 2019 — which is the time period I used for both the cost allocation research and research of this paper. In the cost allocation research, I found that BART's total annual costs are highest, and average per-trip costs lowest, in the urban core of the system and during weekday peak periods; whereas MARTA's cost variability has no clear spatial pattern but similar temporal variability as BART.

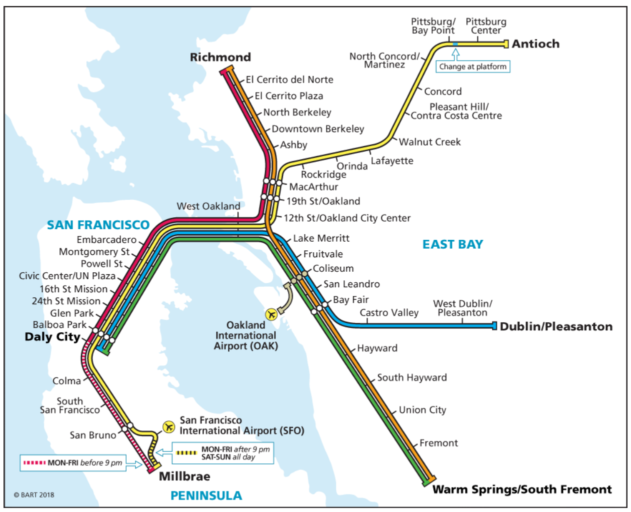

Figure 1: BART System Map effective July 1, 2018 to February 10, 2019

Figure 2: MARTA System Map effective in Fiscal Year 2019

In this paper, I combine findings from my cost allocation research with data on ridership and fare revenues to evaluate whether subsidies have a spatial or temporal pattern in the BART and MARTA networks. I hypothesize that travelers in outlying areas, travelers who travel longer distances, travelers and who travel outside of the weekday peak period each pay a lower share of their costs; but that the middle hypothesis is attenuated with fare policies that are distance-based. In other words, I hypothesize that non-peak period travel, long-distance travel, and travel to/from suburban and exurban areas are disproportionately subsidized relative to their counters, but that distance-based fare policies can reduce this effect for long-distance travel.

My hypothesis that peak period travelers are less subsidized than others is contrary to what past research suggests (e.g., Cervero, 1981; Parody et al., 1990), but aligns with my cost allocation findings that the peak period has the lowest average per trip cost of any operating time period in the two systems. Finally, unlike past research that measures the spatial equity of transit pricing based on whether there is parity in the fare paid per mile of travel consumed, I innovate by controlling for the cost sharing nature of transit. This allows me to account for the variability of costs across unique miles and stops of the network and measure equity based on whether there is parity or disparity in different riders covering the cost of their trip through the fare they pay.

Literature Review

Transit subsidies are generally categorized as either a supply-side subsidy or demand-side subsidy. Supply-side subsidies are intended to cover the costs of service provision that exceed income generated from fares, while demand-side subsidies are offered to riders to reduce the fares they pay. As reasoned, the incidence of both subsidies "trickle down" to riders; by partially covering operating costs, supply-side subsidies have the potential to reduce the fares needed to cover costs, while demand-side directly reduce fares paid by riders. However, supply-side subsidies have been shown to disproportionately benefit unionized transit workers and enable the industry to suffer from workforce inefficiencies and Baumol's Cost Disease (Jones, 1985; Wachs, 1989; Pickrell, 1985; Morales Sarriera and Salvucci, 2016). At various times, as transit subsidies have increased with the goal of increasing service levels, service levels were either maintained or reduced while unionized worker salaries were increased (Wachs, 1989). Morales Sarriera and Salvucci (2016) suggest this is partly driven by the industry being stagnant in technological and efficiency advancements relative to other industries; according to Jones (1985), because there are few promotional opportunities for operators and mechanics, across-the-board salary increases are the primary way frontline workers can increase their incomes.

As for whether and how subsidies trickle down to users, Serebrisky et al. (2009) conduct a critical literature review and surmise that most supply-side subsidy programs are socioeconomically neutral or regressive, while few demand-side subsidies are effective. More recently, federal transit subsidies have focused on capital projects rather than operations, often motivated by the short-term expenditure effects they generate rather than their long-term transportation effects (Taylor and Samples, 2002; Taylor, 2017). Respectively, these refer to the economic activity generated during the construction of a project (i.e., jobs, local spending from workers, etc.) and the project's effectiveness in addressing whatever transportation challenge it is purported to solve. In the long-run, major transit investments can create an operating "commitment trap" wherein an operator is de facto committed to providing the service even if usage does not warrant it (Schweitzer, 2017). The spatial and temporal variability of these long-term operating costs are what I seek to evaluate in this research.

Apart from these top-level considerations, user equity of different transit fare structures has also been studied. This was a popular research topic in the 1980s as transit agencies migrated from using time- and distance-variant fare structures to flat-rate fare structures. Researchers and policymakers sought to understand the equity and efficiency implications of the transition. Such studies typically entail using cost allocation models in which input-output relationships are used to allocate categories of agency costs to service outputs (i.e., occasions of costs), such as allocating power supply and labor costs to vehicle-miles and vehicle-hours, respectively. Once all costs are allocated to different dimensions of a network — time periods, service routes, or otherwise — the result can explain disparities in expenditure patterns. To test for socioeconomic impacts, some studies take an additional step by correlating the cost allocation findings with the makeup of ridership. To-date, research has at-best relied on weakly disaggregate agency data, such as allocating costs and fare revenues to just two time periods (e.g., Reilly, 1977; Parody et al., 1990) or to bus yard "cost centers" (e.g., Cervero, 1981) to evaluate temporal and spatial variability of costs and subsidies, respectively. In addition, the literature overwhelmingly focuses on bus transit and temporal variability.

The research broadly finds that it is more costly to serve peak period travel than off-peak period travel in both gross and net terms (Cervero, 1981; Parody et al., 1990; Taylor, 2000); and that flat rate fare policies relative to time- and distance-variant policies are less effective in recovering total costs, lead to peak period and suburban riders being more subsidized than off-peak and urban riders, and that this pattern is socioeconomically regressive since a higher share of peak period and suburban travel is consumed by higher-income, non-minority persons relative to off-peak and urban travel (Cervero, 1981; Meyer and Gomez-Ibanez, 1981). A small number of studies show that, because the fare revenue generated during the peak period well exceeds the additional costs of serving it, the peak period is less costly on a per rider basis (i.e., it is more cost effective) than the off-peak (Reilly, 1977).

A deciding difference in these findings is if and how both fixed and semi-fixed asset costs are included in the cost allocation step of the research. Cervero (1981) and Parody et al. (1990) allocate asset costs 85% to the peak period and 15% to the off-peak period in evaluating cost recovery variability; Taylor et al. (2000) use a marginal allocation method, derived from the Bradford Bus Study (Savage, 1989), based on the share of transit vehicles in-use during different time periods; and Reilly (1977) does not allocate asset costs. In these studies, only Cervero (1981), Parody et al. (1990), and Reilly (1977) made any measurement of the temporal variability of subsidy patterns; other studies focused on cost variability only. Of these, Cervero (1981) and Parody et al. (1990) clearly use the most robust cost allocation method in that they include capital costs. Even so, the peak-to-base "rule-of-thumb" ratio they use for allocating asset costs is primitive and discounts other time periods to which transit operators scale their resources. For example, in my cost allocation research (Mallett, 2022), I showed that BART and MARTA have eight and five unique operating time periods they scale to, respectively; not two! And while Taylor et al. (2000) used a more robust temporal cost allocation method that includes all capital and operating costs and accounts for there being more than two time periods, they do not evaluate temporal *cost recovery* (i.e., subsidy) patterns. Thus, no study that has measured temporal cost recovery patterns accounts for there being more than two operating time periods.

Studies on capital transit investment disparities have similar findings; the investment patterns tend to benefit higher income, non-minority persons more than lower-income, minority persons who are more transit dependent and use less capital-intensive transit services (Pucher, 1982; Garrett and Taylor, 1999; Taylor and Morris, 2015). However, the extent to which these costs are paid by the beneficiaries, such as through fares, is not considered in these analyses. So, while these studies demonstrate that certain areas and populations receive more transit capital investments in gross terms, if those same areas and populations contribute more through what they pay into the system, the equity implications in terms of net costs may be different.

Since the aforementioned studies, data granularity has greatly improved, which can allow for much more informative findings about how cost recovery and fare equity varies by time of travel. Yet, instead of using this newly available, highly disaggregate data, recent research has relied on past studies' findings to inform their analyses about the equity of current fare policies. For example, Brown (2018) uses the 45% peak-to-base net cost ratio from Parody et al. (1990) based on 1983 national aggregate data and 2012 California Household Travel Survey data to conclude that the current Los Angeles Metropolitan Transportation Agency's transit fare structure is regressive and would be more equitable if it adopted time- and distance-variant fares. Other recent studies are international and suggest that, contrary to findings in the United States, flat rate fares benefit lower income persons in Stockholm (Rubensson, 2020) and China (Zhou and Zhang, 2019) — likely due to lower income persons settling further from the urban core, unlike in the United States.

Finally, research on the spatial equity of transit fare policies is practically non-existent. Geopolitical considerations inform transport investments, as politicians and voters typically desire a visible return on their tax investments (Taylor, 2017; Taylor and Norton, 2009). However, the extent to which different geopolitical areas are subsidized at the user level is not a focus of much transit fare policy research. Among few exceptions are Hodge (1988), Iseki (2016), and the previously- referenced study of Cervero (1981). Cervero (1981) used bus facilities as "cost centers" and labeled them as suburban or urban to draw his conclusion that suburban riders pay less per mile of travel consumed, on average. In this way, he used highly aggregate geographies, vehicle-miles, and passenger-miles for his spatial analysis. Hodge studied the Seattle, Washington transit network, labeled every mile of service as urban or suburban using geopolitical boundaries, and found that when only fares are considered, core (urban) riders subsidize suburban riders; but that the net flow of subsidies goes from suburban areas to urban areas when the tax base is included in the analysis. However, Hodge's method of asset and operating cost allocations used only vehicle-miles and vehicle-hours as outputs, so was not robust. Iseki used a more robust cost allocation method and had similar findings in the Toledo, Ohio region; the urban area of Toledo is dependent on tax revenue from suburban areas, despite urban Toledo generating the highest ridership and fare revenue.

In sum, past literature suggests that, when asset costs are accounted for, the peak period is subsidized more than the off-peak period, but that the pattern is reversed when these costs are not accounted for; flat rate fare structures are socioeconomically regressive because they allow persons who consume longer and more suburban trips — whom tend to be white and wealthy — to pay less per mile; and urban areas pay a higher share of travel costs through fares than suburban areas, but a lower share on net when tax source revenues are accounted. Among missing elements in the literature are more thorough analyses of transit modes other than bus transit, as well as the use of available granular data to evaluate temporal and spatial variability of cost recovery patterns more precisely. For example, transit operators' time periods often are more nuanced than merely peak and base, and spatial variability is more varied than geopolitical boundaries. Most significantly, past research on spatial fare equity (e.g., Cervero, 1981) does not account for the cost sharing nature of transit — that is, how much the cost of each unique mile of travel is shared amongst the consumers of that unique mile of travel. Some game theory literature on how transit riders may change travel patterns in a cost sharing scheme (Rosenthal, 2017), as well as research on private sector "collaborative transportation" concepts (e.g., Frisk et al., 2010; Guajardo and Rönnqvist, 2016), exist; but no research on transit cost recovery and fare equity accounts for this.

Data and Methods

The objective of this work is to evaluate the temporal and spatial variability of farebox recovery of rail transit networks. Do travelers of different times and locations pay a different share of their travel costs through fares in any statistically significant way such that transit operations are not only inefficient (i.e., costs are not fully recovered through user payments), but inequitably so on spatial and temporal dimensions? To answer this, I rely on findings from my prior cost allocation research (Mallett, 2022) and trip-level ridership and fare data to estimate the costs and cost recoveries of unique origin-destination (OD) trips and time periods. I then weight trip-level findings to travel consumption patterns to evaluate what average riders of different links and stations of the network pay as a share of the cost of their trip. All original data came from the transit agencies, BART and MARTA. **Table 1** defines variables that are used in my analysis.

Variable	**Description**
$cost_{od}$	The cost of serving a particular OD trip, *od*.
$cost_t$	The total cost expended on serving a particular time period, *t*.
$cost_{tot}$	The total cost expended during the study period.
costppm	The cost per passenger-mile of a particular link; that is, the total annualized cost of a particular link divided by the sum-product of the length (in track-miles) and number of annual trips that traverse the particular link.
costppx $costppx_l$ $costppx_o$ $costppx_d$	The cost per passenger of a particular station or link; that is, the total annualized cost of a particular station or link divided by the number of trips into or out of the particular station or that traverse the particular link. Subscripts "l," "o," and "d" correspond to a unique link, origin station, and destination station, respectively.
$fare_{od}$	The average fare paid for consuming a particular OD trip, *od*.
$fares_t$	The total fare revenue generated during a particular time period, *t*.
$fares_{tot}$	The total fare revenue generated during the study period.
$percentpaid_{effective}$	The average fare paid for a particular OD trip divided by the cost of serving the OD trip.
$percentpaid_l$ $percentpaid_s$	The average cost recovery ($percentpaid_{effective}$) for all OD trips associated with a particular link, *l*, or station, *s*.
$percentpaid_t$	The total costs recovered through fares generated in a particular time period, *t*.
pxmiles	The total passenger-miles generated — that is, the sum-product of the count of trips associated with each OD trip and the trip length of each OD trip.
$subsidy_{effective}$	The cost of a particular OD trip not paid by the average fare of the OD trip; 100% minus the effective percent paid.
triplength	The length of a particular OD trip, in track-miles. Limited to mainline track only.
$triplength_{average}$	The average length, in track-miles, of all annual OD trips made that used a particular station or link of the network. Limited to mainline track only.
distancecore $distancecore_l$ $distancecore_s$ $distancecore_o$ $distancecore_d$	The distance, in track-miles, that a particular station or link associated with an OD trip is from a defined core station of the network — West Oakland Station for BART, Five Points Station for MARTA. Subscripts "l," "s,", "o," and "d" correspond to a unique link, station, origin station, or destination station, respectively. For links, the distance that the station farthest from the core station is defines this value. Limited to mainline track only.

$trips_{od}$ $trips_l$ $trips_s$ $trips_t$ $trips_{tot}$	The number of trips generated across the study period. Subscripts "od," "l," "s," "t," and "tot" correspond to unique ODs, links, stations, time periods, and the total for the entire study period, respectively.

Table 1: Table of Variables

The cost of serving a particular OD trip is equal to the sum of the average cost per rider of the origin station, destination station, and each link used to complete the trip, as derived from the cost allocation study (Mallett, 2022). This is modeled in **Equation 1**. **Equation 2** defines the effective percent paid for a particular OD trip, which is equal to the weighted average of what every rider who consumed a particular OD trip paid divided by the cost of serving the trip. This accounts for the different fares that different riders paid based on the discount programs, etc. that applied to their fare payment. One minus the effective percent paid is the effective subsidy given to consuming a particular OD trip.

$$cost_{od} = costppx_o + costppx_d + \sum_{l=1}^{n} costppx_l \tag{1}$$

$$percentpaid_{effective} = \frac{fare_{od}}{cost_{od}} \tag{2}$$

To test how much trip subsidies are explained by trip length and distance from the urban core of the network, I run ordinary least squares regressions. I regress the effective percent paid of an OD trip onto trip length, the distance the origin station is from the defined core station of the network, and the distance that the destination station is from the defined core station of the network. This is defined in **Equation 3**.

$$percentpaid_{effective} = \beta_0 + \beta_1 triplength + \beta_2 distancecore_o + \beta_3 distancecore_d \tag{3}$$

While Equation 3 will estimate how much an average consumer of a particular OD trip pays relative to the cost of serving them, it does not explain the geographic incidence of subsidies — that is, whether subsidies flow to particular areas of the network more than others, given the weighted pattern of OD trips throughout the network. To evaluate the geographic concentration of transit subsidies, I calculate the average cost recovery of all trips associated with every station and link of the network, as shown in **Equation 4** and **5**, respectively. A trip is associated with a station if the station is the origin or destination of the trip; with a link, if it traverses the link. To be clear, fare revenues are not allocated to links or stations, as there is no practical way to divvy up a fare into time, mileage, and station components for link and station fare revenue allocations. So, these are not link and station cost recoveries, per se; they are the weighted average cost recovery of all OD trips that used the link or station during the study period (i.e., FY19). Because there are just 38 stations and 37 links in the MARTA network, and 48 stations and 47 links (studied) in the BART network, there is an insufficient sample size to use multivariable regression analysis in this portion of the study. As an alternative, I provide descriptive findings and offer interpretation.

$$percentpaid_s = \frac{\sum_{s=o,d}(trips_{od} * percentpaid_{effective})}{trips_s} \tag{4}$$

$$percentpaid_l = \frac{\sum(trips_{od} * percentpaid_{effective}) \in l}{trips_l} \tag{5}$$

Finally, to estimate the incidence of temporal subsides, I divide the total fare revenue generated during a particular time period by the cost of serving that time period — the latter derived from the cost allocation study (Mallett, 2022). This is modeled in **Equation 6**.

$$percentpaid_t = \frac{fares_t}{cost_t} \tag{6}$$

As with the link- and station-level spatial analysis, there is an insufficient number of time periods — just five and eight for MARTA and BART, respectively — for any multivariant regression analysis. In lieu of this, I calculate the correlations between cost recovery and travel pattern variables — including the count of trips, total passenger-miles generated, average trip length, and average distance of origin and destination stations from the defined core station — for trips taken during each time period. This will not show the variation of net travel patterns, but the variation of travel patterns *across time periods*.

Descriptive Statistics

Table 2 provides agency profile information of the BART and MARTA rail networks, effective during FY19. Noteworthy is that, although BART has a network about 2.3 times the size of MARTA's, it generated 3.5 times as many vehicle revenue-miles and 3.9 times as many trip-miles, but only 1.9 times as many trips.

	BART	**MARTA**
Miles (mainline)	109.4 miles	48 miles
Stations	48 stations	38 stations
Fare Structure	Distance-based	Flat rate
Net Fare Revenue	$448,688,735	$58,576,496
Gross Costs	$847,799,127	$237,992,718
Farebox Recovery Ratio	52.9%	24.6%
Vehicle Revenue-Miles*	78M vehicle-miles	23M vehicle-miles
Annual Trips*	125M trips	65M trips
Annual Passenger-Miles*	1.756B passenger-miles	450M passenger-miles

**As reported in 2019 Agency Profiles of National Transit Database*

Table 2: Agency Profiles

During FY19, the MARTA system had a total of 38 stations, resulting in 1,406 OD pairs. By comparison, BART had a total of 48 stations, resulting in 2,256 OD pairs. However, BART's network includes three rail service types — standard BART trackage (mainline), a diesel multiple unit segment of service (eBART), and the Oakland Airport Connector (OAC) that operates as a cable-powered people mover. As in the cost allocation research (Mallett, 2022), while I include all stations, I only include mainline portions of the BART track network to ensure consistency in analysis. By extension, any trips that solely use non-mainline portions of track are excluded from my analysis in this paper. This results in a total of 2,248 OD pairs.

Finally, for any trips that partially used a non-mainline portion of track, I assign trip costs using only the mainline links of the networks — that is, a trip's cost is equal to the summation of the costs per rider of the origin station of the trip, destination station of the trip, and each mainline link used to fulfill the trip. As an example, a trip from Oakland International Airport Station to Antioch Station will exclude the Oakland Airport Connector and eBART segments of trackway, so will be charged as a trip that utilizes links from Coliseum Station to Pittsburg/Bay Point Station, plus the cost per rider of Oakland International Airport Station and Antioch Station.

The two agencies' fare structures are also different. MARTA's fare structure is flat rate at $2.50 per trip, regardless of trip length, origin, or destination. However, there are various discounts, including through transfer agreements with other transit operators, cooperative arrangements with area employers, multi-day passes, and more. At BART, base fares are distance-based, but follow a stepwise function. Riders pay a minimum fare for the first six miles of travel. Beyond six miles, riders pay for the first six miles plus a rate per mile up to 14 miles; beyond 14 miles, riders pay for the first 14 miles plus a lower cost per mile greater than 14 miles. Accordingly, although riders pay more for every additional mile traveled beyond six, longer distance travel is discounted on a per mile basis. There are also various fees, including a fee for use of the transbay tube, for travel to or from San Mateo County (a county that does not pay into the BART District through taxes), and for travel to or from San Francisco or Oakland International Airports; as well as various discount programs, including for senior and disabled riders, youth, and high-value discount (HVD) tickets. The effective percent paid calculations of every OD trip are reflective of the weighted average of these many discounts used by riders of the OD pairs.

Spatial Analysis

Tables 3 and **4** show descriptive statistics of key variables in the OD cost recovery analysis, including the count of trips associated with every OD pair, the effective percent paid of every OD pair, the trip length of every OD pair, the distance the origin station is from the defined core station of the network, and the distance that the destination station is from the defined core station of the network — for BART and MARTA, respectively. In these tables, I show both the unweighted distribution, which weights all OD trips equally so is irrespective of cumulative trip consumption patterns; as well as the distribution weighted to trip count, which reflects cumulative trip patterns rather than station-to-station pairs. So, while a particular OD trip may have a cost recovery higher than another, if the former has more consumption (i.e., trip counts) than the latter, the mean weighted to trip count will be higher than the unweighted mean. Similarly, the average, median, and standard deviation of trip length for all actualized trips will be different than the same statistics for the set of OD pairs without weighting. I use unweighted values to calculate how OD trip cost recoveries are associated with trip length and station distances from the core but show both unweighted and weighted statistics here to demonstrate how OD pair patterns differ from aggregate trip consumption patterns — which is pertinent for the station and link analysis that evaluates the incidence of subsidies.

	Minimum	Mean		Median		Standard Deviation		Maximum
		Unweighted	Weighted	Unweighted	Weighted	Unweighted	Weighted	
percentpaid$_{effective}$	10.10%	44.24%	63.31%	39.34%	57.21%	19.24%	25.59%	158.29%
distancecore$_o$	0	13.53	10.16	11.98	7.18	8.33	7.09	29.61
distancecore$_d$			10.04		6.68		7.04	
triplength	0.35	21.82	14.98	20.97	12.63	12.5	10.05	56.44
trips	176	52,257	—	15,874	—	92,539	—	751,744

N (OD pairs) = 2,248

Table 3: Descriptive Statistics of OD Pairs — BART

	Minimum	Mean		Median		Standard Deviation		Maximum
		Unweighted	Weighted	Unweighted	Weighted	Unweighted	Weighted	
percentpaid$_{effective}$	10.2%	30.10%	36.05%	27%	31.44%	13.38%	17.50%	105.25%
distancecore$_o$	0	5.66	5.75	4.73	5	4.61	4.72	15.95
distancecore$_d$			5.69		5		4.71	
triplength	0.38	9.62	9.13	9.05	8.45	5.95	6.05	26.39
trips	355	29,341	—	14,854	—	38,910	—	404,000

N (OD pairs) = 1,406

Table 4: Descriptive Statistics of OD Pairs — MARTA

One thing that stands out in these tables is the difference between the two networks in terms of how the unweighted (i.e., OD pairs, irrespective of trip count) and weighted (i.e., total OD trips) statistics vary. BART's weighted data are notably different than the unweighted data — implying that trip consumption patterns vary significantly across the network; trip consumption is not normally distributed. Specifically, the average cost recovery of trips consumed (weighted) is 1.4 times *more*, average trip length 32% *less*, and the distance from the core of the system that trips begin or end about 26% *less*, than if trips were equally spread across all OD pairs (unweighted). By comparison, MARTA's trip consumption patterns are similar to the network's unweighted spread of OD pairs. As I show later, this variance between the two networks reflects the difference in spatial concentration of travel across stations and links of the networks.

Focusing on cumulative trip patterns (weighted data), the mean OD trip cost recovery for BART and MARTA is 63% and 36%, respectively — higher than the systemwide cost recovery of 53% and 25%, respectively, as reported in Table 2. The unweighted mean is also different from the systemwide total. One might think that the mean trip-level cost recovery should be equal the systemwide cost recovery; that is, the mean of all trips' fares divided by their costs should equal the sum of fares divided by the sum of costs (**Equation 7**).

$$\frac{\sum fare_{OD}}{\sum cost_{OD}} \stackrel{?}{=} \frac{fares_{tot}}{costs_{tot}} \tag{7}$$

However, the sum of OD trip costs is *not* equal to aggregate systemwide costs because OD trip costs are derived from a subset of costs *after* costs have been allocated to stations and links and divided by ridership (see Equation 1). Hence, OD trip costs move across different subsets of costs (i.e., parts of the network). Furthermore, OD trips that utilize low cost per rider/high ridership stations and links will skew the OD trip cost recovery pattern higher than the aggregate systemwide cost recovery. That is, while more cost-effective trips may be consumed, cost-ineffective trips and segments of a network are still served. As a result, the average cost recovery of consumed trips will tend to be greater than the systemwide cost recovery. This is consistent with an objective of this research: the aggregate pattern of a network does not necessarily reflect patterns across its parts.

Apart from cost recovery, trips taken on BART, on average, begin or end just over ten miles, or 34% of the maximum possible distance, from the defined core station of the network. For MARTA, the average trip begins or ends about 5.7 miles from the defined core station, or about 36% of the distance the furthest away station is from the defined core station. Similarly, the average length of consumed trips taken on BART is about 15 miles or 39% of the maximum possible trip length; on MARTA, 9.1 miles or 35% of the longest possible trip. Thus, as a percentage of the maximum possible, trips are about as "suburban" and long in both networks.

On the other hand, the spatial pattern of how trips are concentrated in the two systems is distinct. **Figures 3** and **4** show the distribution of ridership patronage of stations — that is, out of all trips taken, the share that begin or end (inclusive) at each station — in the BART and MARTA network, respectively. We can see from this that about two-thirds of all trips taken in the BART system begin or end at BART's four busiest stations, which account for 8% of the system's stations, and that there is a notable drop off after these stations. By comparison, 53% of all trips taken on MARTA begin or end at its four busiest stations, which represent 11% of its stations, and the drop off is less abrupt. In both systems, these four stations are the only stations with a double-digit share of passenger patronage.

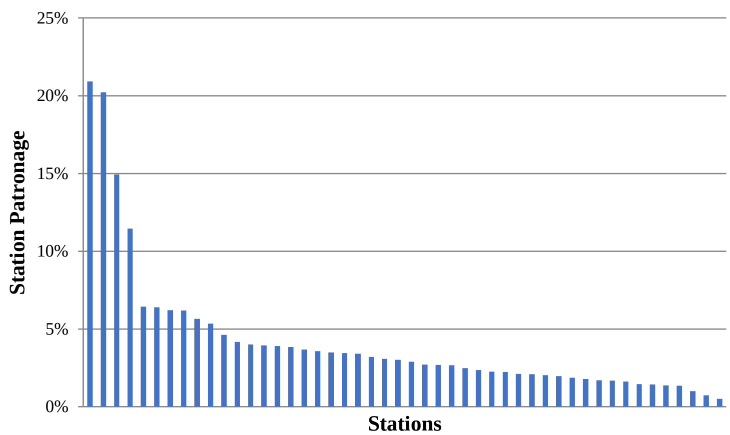

Figure 3: Ridership Patronage by Station — BART

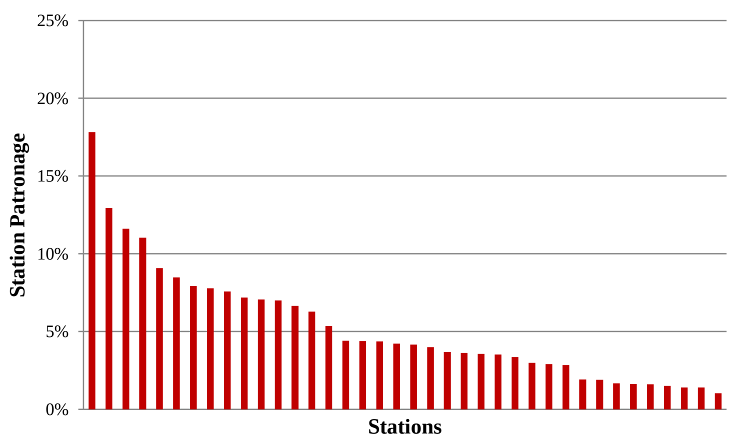

Figure 4: Ridership Patronage by Station — MARTA

BART's four busiest stations — Embarcadero, Montgomery Street, Powell Street, and Civic Center/UN Plaza Stations — are adjacent to one-another in Downtown San Francisco. By contrast, the four busiest MARTA stations — Five Points, Airport, Peachtree Center, and College Park Stations — are associated with two different epicenters of ridership, Downtown Atlanta and the airport. The ridership at the eight adjacent Downtown Atlanta stations (21% of stations) — North Avenue, Civic Center, Peachtree Center, Five Points, Garnett, Georgia State, Dome/CNN Center, and Vine City — individually range from first (or upper third percentile) to 35^{th} (or lower 11^{th} percentile) in passenger patronage; collectively, 51% of trips begin or end in Downtown Atlanta. Further, exclusive of the four busiest stations, every BART station's ridership has its highest destination relationship with one of the four busiest stations, ranging from 10% to 30%, with an average of 19%. When the busiest stations are analyzed as a group, every other BART station has between 29% and 73% (with a 50% average) of its trips destined to Downtown San Francisco. At MARTA, just 27, or 79%, of stations exclusive of the busiest four, have their strongest destination relationship with one of the four busiest stations. When the four busiest stations are grouped, their destination market-share of trips originating at other stations ranges from 15% to 42% with an average of 28%. Focusing on the eight Downtown Atlanta stations, 22, or 73%, of other stations have one of the Downtown Atlanta stations as their highest market-share destination station. As a group, Downtown Atlanta accounts for between 18% and 47% of the destination market-share of other stations with an average of 31%. It is clear from this that BART's OD ridership patterns revolve around a monocentric center, Downtown San Francisco, whereas MARTA's OD ridership patterns are polycentric if not broadly dispersed. This can influence how origin and destination station distance from the urban core of each network, as well as trip length, influences spatial cost recovery patterns.

It is important to note that, while the level of monocentricty in origin-destination travel patterns is markedly different between BART and MARTA, this does not necessarily extend to the links used to complete various trips. **Figures 5** and **6** geographically show how travel patterns on BART and MARTA affect station patronage and link usage (note that the link usage scales vary significantly). These figures suggest that, while station patronage is relatively monocentric for BART and dispersed for MARTA, the bi-directional flow of ridership along links is centrally concentrated in both networks. This makes sense, as even in a polycentric or dispersed travel environment, trip paths can be centrally concentrated along links of a network even if origins and destination are not. This is especially subject to occur if the network has a defined central node that many origin-destination trips must pass through, as is the case for MARTA. I investigate these observations more thoroughly in the Results section.

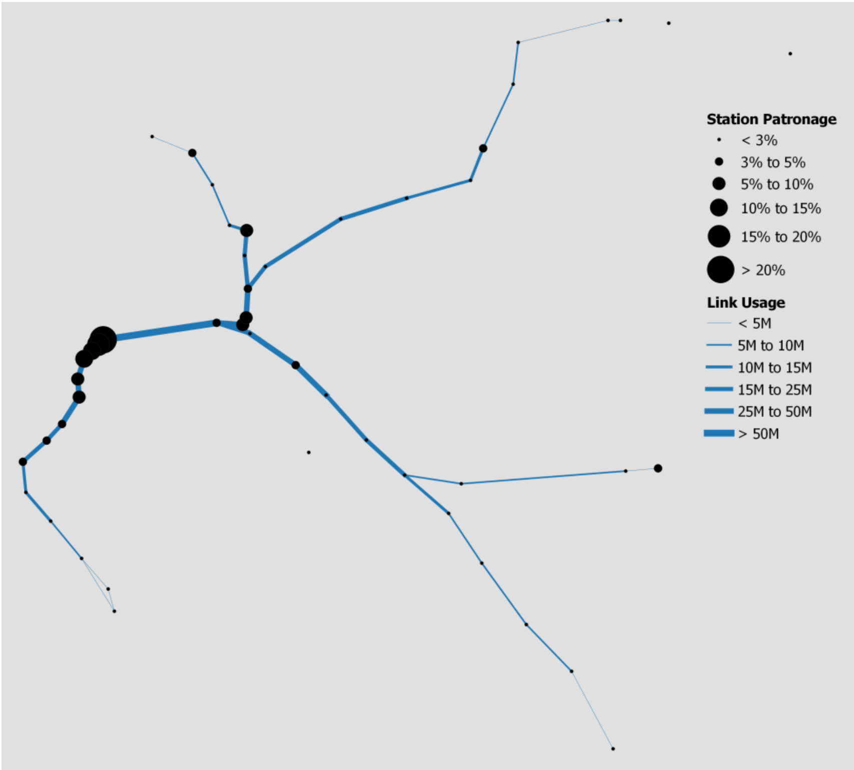

Figure 5: Spatial Pattern of Station Patronage and Link Usage — BART

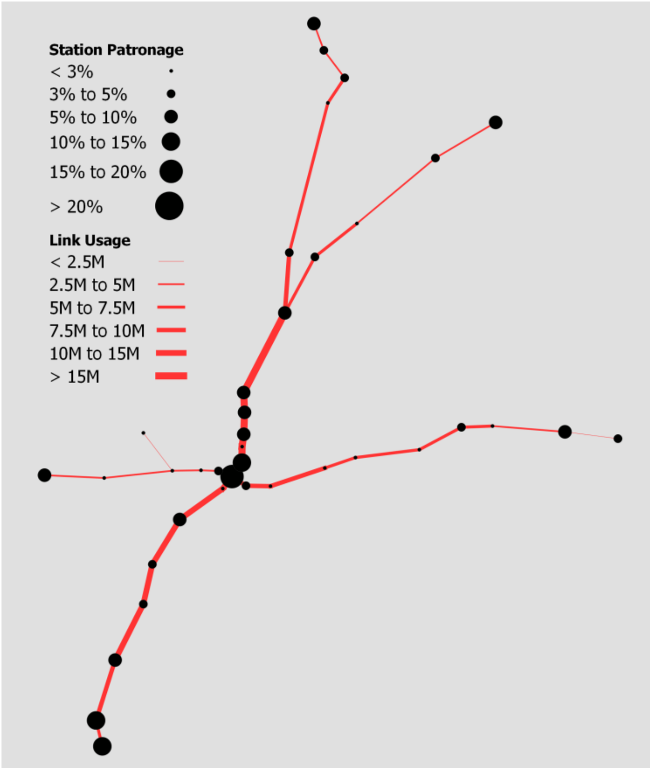

Figure 6: Spatial Pattern of Station Patronage and Link Usage — MARTA

Tables 5 and **6** show the resulting descriptive statistics for the station-level analysis — that is, each station's distance from the defined core station of the network, the weighted average percent paid for all trips that began or ended at the station, the weighted average trip length for all trips that began or ended at the station, and the number of trips that began or ended at the station — for BART and MARTA, respectively. These data are generated from the above-summarized OD data, conditional on a station being an origin or destination station. **Tables 7** and **8** show the same descriptive statistics for the link-level analysis — in this case, corresponding to all trips that traversed a particular link, which is easily derived through trip path assignment. The objective of this part of the analysis is to evaluate whether user subsidies have a spatial pattern. Since fares are not allocated to links and stations, the *percentpaid* variable in these tables does not (and cannot) reflect station and link cost recoveries; they are a cost recovery "profile" of the average rider of each unique station and link. Therefore, these numbers will not necessarily align with OD trip or systemwide aggregate cost recovery statistics.

Variable	Minimum	Mean	Median	Standard Deviation	Maximum
$percentpaid_s$	24.86%	54.47%	54.57%	14.36%	99.79%
$distancecore_s$	0	13.56	12.07	8.44	29.61
$triplength_{average}$	7.67	17.36	15.25	7.8	42.25
$trips_s$	617,004	4,897,543	3,473,553	5,053,904	24,571,444

N (stations) = 48

Table 5: Descriptive Statistics of Station Profiles — BART

Variable	Minimum	Mean	Median	Standard Deviation	Maximum
$percentpaid_s$	13.38%	29.37%	29.93%	6.27%	40.71%
$distancecore_s$	0	5.66	4.73	4.67	15.95
$triplength_{average}$	5.02	8.81	7.47	3.21	16.36
$trips_s$	430,987	2,171,240	1,728,397	1,529,423	7,356,599

N (stations) = 38

Table 6: Descriptive Statistics of Station Profiles — MARTA

Variable	Minimum	Mean	Median	Standard Deviation	Maximum
$percentpaid_l$	24.86%	47.61%	48%	9.29%	64.58%
$distancecore_l$	1.59	13.61	12.16	8.21	30.02
$triplength_{average}$	1.82	22.14	19.88	7.36	39.78
$trips_l$	154,130	18,199,306	13,028,305	15,014,428	64,656,418

N (links) = 47

Table 7: Descriptive Statistics of Link Profiles — BART

Variable	Minimum	Mean	Median	Standard Deviation	Maximum
$percentpaid_l$	13.38%	23.61%	23.83%	2.85%	27.35%
$distancecore_l$	0.38	5.82	4.8	4.64	15.95
$triplength_{average}$	6.37	12.28	12.42	2.45	16.57
$trips_l$	430,987	8,194,683	6,654,606	4,931,187	17,505,807

N (links) = 37

Table 8: Descriptive Statistics of Link Profiles — MARTA

Comparing these statistics with the weighted OD trip statistics in Tables 3 and 4, station and link mean cost recovery profiles are clearly less than the mean weighted OD trip cost recovery. With smaller standard deviations, the spread is also much less than in the case of weighted OD trip cost recoveries. This is principally explained by these being different measurements: the latter is a measurement of trip cost recoveries, while the former are measurements of trip cost recovery *averages* for trips that are associated with the link or station. As a result, although high cost recovery OD trips "crowd out" low cost recovery OD trips in the OD trip cost recovery statistics, those same OD trips share links and stations with low cost recovery OD trips. Hence, even the links or stations with the highest cost recovery profile will inherently have a value less than the highest cost recovery OD trip. Furthermore, there are far fewer links and stations than there are OD pairs, and high cost recovery trips may be concentrated across few links and stations, leaving the remainder of links and stations to be associated with lower cost recovery OD trips. This last concept is evaluated in the Results section.

Figures 7 and **8** geographically show the distribution of the *percentpaid* variable across stations and links in the BART and MARTA network, respectively.

Figure 7: Average Percent Paid Across Links and Stations — BART

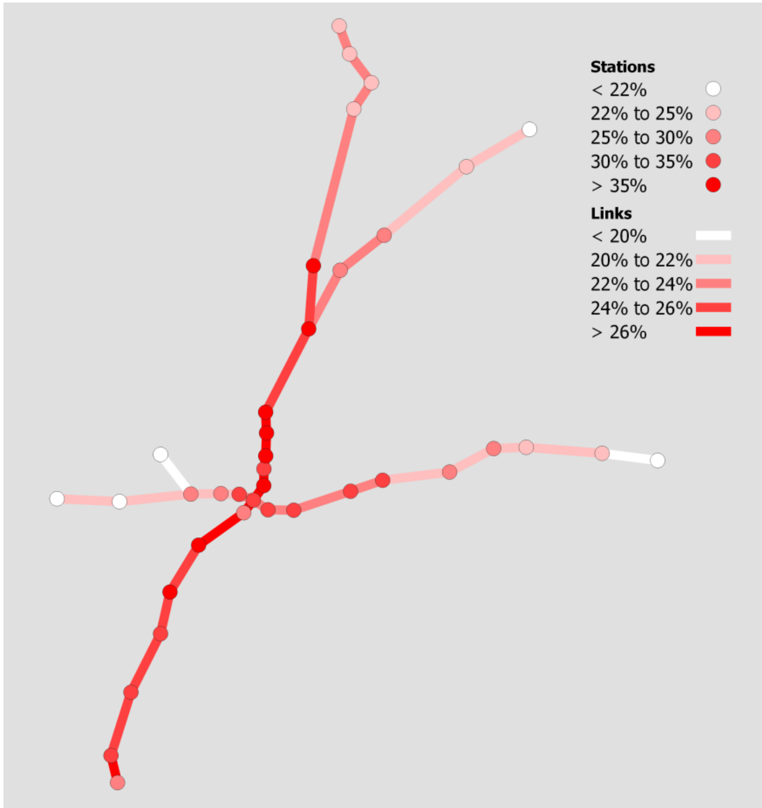

Figure 8: Average Percent Paid Across Links and Stations — MARTA

To illustrate interpretation of these maps, the "average rider" who travels to or from MARTA's Five Points Station pays 34.1% of the costs of their trip, while the "average rider" who traverses BART's link between West Oakland and Embarcadero Stations (the transbay tube) pays 63.9% of their total trip cost. With few exceptions, these figures surmise that the average cost recovery of riders of different links and stations generally *declines* with distance from each system's core — though, as suggested by the different scale on the maps, the magnitude of variance is significantly less for MARTA relative to BART. Also notable and consistent with my cost v. cost per rider analysis (Mallett, 2022), BART's most expensive stations and links to operate — and likely to build due to their being tunneled — also have the highest level of ridership.

Temporal Analysis

For temporal analysis, **Figures 9** and **10** show the total annualized allocated costs (left axis) and average cost per rider (right axis) of the different time periods for BART and MARTA, respectively. These data are derived from the cost allocation study (Mallett, 2022). **Tables 9** and **10** show full descriptive statistics for trips consumed during different time periods for each agency. Five of the variables — cost, fare revenue, cost recovery, trip count, and passenger-miles — are aggregated to the time period; they reflect the total costs, fare revenue, etc. for the time period. The other three variables — origin station distance form the core station, destination station distance from the core station, and trip length — are disaggregated to trips consumed during each time period. Accordingly, the variation in these values reflects how trip consumption patterns vary across time periods.

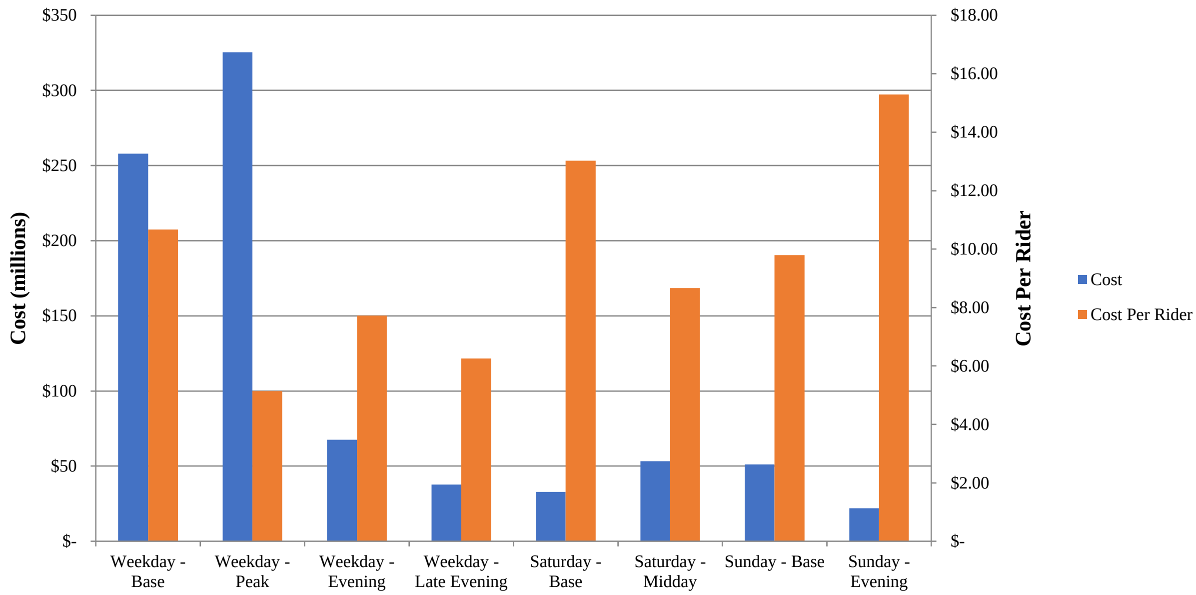

Figure 9: Cost v. Cost Per Rider by Time Period — BART

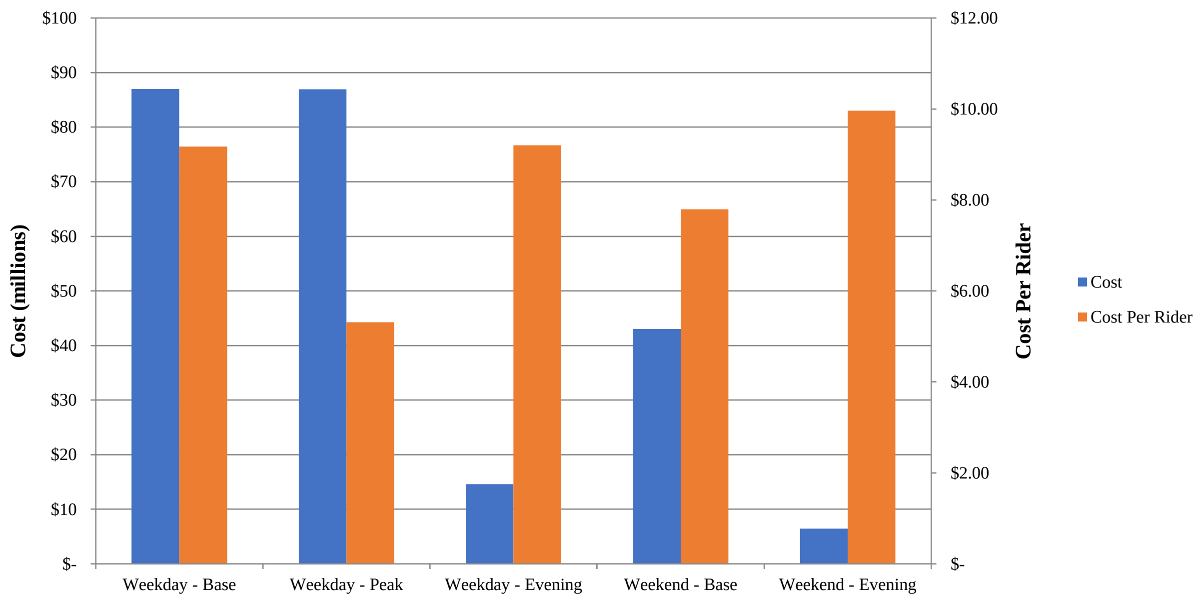

Figure 10: Cost v. Cost Per Rider by Time Period — MARTA

Time Period		Weekday				Saturday		Sunday/Holiday		Total
		Base	**Peak**	**Evening**	**Late Evening**	**Base**	**Midday**	**Base**	**Evening**	
$cost_t$		$258,037,518	$325,430,391	$67,444,710	$37,682,645	$32,792,752	$53,279,185	$51,104,798	$22,027,128	$847,799,127
$fares_t$		$89,697,710	$243,146,564	$35,114,074	$23,259,610	$9,622,910	$22,504,842	$19,464,467	$5,878,558	$448,688,735
$percentpaid_t$		34.76%	74.72%	52.06%	61.72%	29.34%	42.24%	38.09%	26.69%	52.92%
$trips_t$		24,177,808	63,242,768	8,732,436	6,017,829	2,514,618	6,141,670	5,208,938	1,438,202	117,474,269
pxmiles		343,857,182	970,670,140	138,290,468	89,440,215	36,825,580	84,332,354	72,345,257	22,255,671	1,758,016,867
triplength	**Minimum**	0.35				0.35		0.35		0.35
	Mean	14.22	15.35	15.84	14.86	14.64	13.73	13.89	15.47	14.97
	Standard Deviation	10.16	9.96	10.45	10.15	10.01	9.84	9.86	10.6	10.07
	Maximum	56.92				56.92		56.92		56.92
$distancecore_o$	**Minimum**	0				0		0		0
	Mean	11.04	10.27	8.14	8.4	9.4	10.34	10.44	9.82	10.16
	Standard Deviation	7.59	7.18	5.38	5.54	6.4	7.03	7.12	6.8	7.09
	Maximum	29.61				29.61		29.61		29.61
$distancecore_d$	**Minimum**	0				0		0		0
	Mean	8.89	9.99	12.41	11.79	10.89	9.54	9.8	11.46	10.04
	Standard Deviation	6.09	7.06	8.13	7.89	7.47	6.46	6.71	7.69	7.04
	Maximum	29.61				29.61		29.61		29.61

Table 9: Temporal Descriptive Statistics — BART

Time Period		Weekday			Weekend/Holiday		Total
		Base	**Peak**	**Evening**	**Base**	**Evening**	
$cost_t$		$86,991,391	$86,950,887	$14,582,218	$43,032,908.02	$6,435,315	$237,992,718
$fares_t$		$15,839,966	$29,559,350	$2,626,714	$9,450,864	$1,099,602	$58,576,496
$percentpaid_t$		18.21%	34%	18.01%	21.96%	17.09%	24.61%
$trips_t$		12,391,117	19,898,647	2,057,750	6,985,699	841,506	42,174,719
pxmiles		109,206,799	188,547,757	17,270,565	61,508,060	6,977,054	383,510,235
triplength	**Minimum**	0.38			0.38		0.38
	Mean	8.81	9.48	8.39	8.8	8.29	9.09
	Standard Deviation	6.22	5.94	5.86	6.04	5.79	6.04
	Maximum	26.39			26.39		26.39
$distancecore_o$	**Minimum**	0			0		0
	Mean	5.56	5.96	5.17	5.69	4.81	5.74
	Standard Deviation	4.57	4.9	4.31	5.54	4.27	4.71
	Maximum	15.95			15.95		15.95
$distancecore_d$	**Minimum**	0			0		0
	Mean	5.57	5.77	8.39	8.8	8.29	5.68
	Standard Deviation	4.61	4.87	5.86	6.04	5.79	4.7
	Maximum	15.95			15.95		15.95

Table 10: Temporal Descriptive Statistics — MARTA

As evident from these data, the peak period is the most expensive to operate by far for BART and almost as expensive to operate as the weekday base period for MARTA. But in both cases, the weekday peak period serves so many more riders and passenger-miles that the cost per rider and per passenger-mile is lowest during the weekday peak period. Similarly, by serving so many more riders and passenger-miles, much more fare revenue is generated during the weekday peak period — so much that it offsets any additional costs of serving the weekday peak period, even when semi-fixed costs are accounted for. As a result, the weekday peak period recovers the greatest share of costs of any time period. In the following section, I show correlates between the descriptive statistics and offer interpretation.

Results — Spatial Analysis

Tables 11 and **12** show the results of the OD cost recovery model defined in Equation 3 for BART and MARTA, respectively. In both systems, for a given OD trip, its trip length is negatively associated with cost recovery, though the magnitude is markedly greater for MARTA. Whereas every mile longer that a trip is, is associated with a 0.07 percentage point decrease in cost recovery for BART; it is associated with a 2.1 percentage point decrease in cost recovery for MARTA. On the other hand, the distance that origin and destination stations are from the core of the networks has an opposing effect. For BART, every additional mile an origin or destination station is from the core station is associated with a 0.9 percentage point reduction in farebox recovery for the trip. Furthermore, this has a greater statistical significance on cost recovery than trip length for BART. In stark contrast, for MARTA, every additional mile further that an origin or destination station is from the core is associated with a 0.7 percentage point *increase* in cost recovery and has the same degree of statistical significance as trip length. This implies that travelers who use outlying stations of the MARTA network pay a greater share of their trip costs, on average, after controlling for other factors.

Variable	Coefficient	Standard Error	95% Confidence Interval
triplength	-0.067*	0.04	(-0.14, 0.01)
$distancecore_o$	-0.94***	0.05	(-1.04, -0.85)
$distancecore_d$	-0.94***	0.05	(-1.04, -0.84)
constant	71.15***	0.83	(69.53, 72.77)
N	2,248		
Adjusted R-squared	0.3577		

Statistical Significance: *** p ≤ 0.001, ** p ≤ 0.01, * p ≤ 0.1

Table 11: Origin-Destination Cost Recovery Model Results — BART

Variable	Coefficient	Standard Error	95% Confidence Interval
triplength	-2.12***	0.07	(-2.25, -1.99)
$distancecore_o$	0.73***	0.07	(0.59, 0.87)
$distancecore_d$	0.73***	0.07	(0.59, 0.87)
constant	42.27***	0.52	(41.25, 43.28)
N	1,406		
Adjusted R-squared	0.5073		

Statistical Significance: *** p ≤ 0.001, ** p ≤ 0.01, * p ≤ 0.1

Table 12: Origin-Destination Cost Recovery Model Results — MARTA

However, what more reasonably explains this counterintuitive result for MARTA is a combination of the collinearity of the independent variables, travel patterns in the network, and

MARTA's flat rate fare structure. In both networks, the one-to-one relationship that OD trip cost recovery has with each independent variable is negative and statistically significant. For MARTA, the one-to-one correlation between cost recovery and trip length and between cost recovery and station distance from the core, is -0.68 and -0.26, respectively. For BART, these numbers are -0.45 and -0.42, respectively. At the same time, the station distance variables have a statistically significant positive relationship with the trip length variable — 0.53 for MARTA and 0.5 for BART. This makes intuitive sense, as suburban travel is associated with longer trips, in general (e.g., Sultana and Weber, 2007). Thus, despite both independent terms having a negative influence of cost recovery, there is a cancelling out effect between them in the regression, and trip length has a greater level of influence for MARTA partly due to its stronger one-to-one influence to begin with.

In addition, spatial patterns of travel in the MARTA network are markedly more dispersed than in the BART network. The more evenly dispersed travel in a network is, the less influence a node's distance from the core of the network will have on the cost recovery, all else being equal. Finally, in a flat rate fare structure environment, like MARTA, *all else being equal*, every additional mile traveled on a network will be associated with a lower cost recovery. That is, the consumer is consuming more wear, labor-hours, etc. and not paying anything more for it, unlike in the BART network. Taken together, trip length will reasonably have a greater influence on cost recovery outcomes in the MARTA network than in the BART network.

My analysis heretofore has measured how the cost recovery of single OD trips are explained by their lengths and the distance that their origin and destination stations are from the core of the network. This does not explain the geographic incidence of rail transit subsidies. If, for example, a station or link has a disproportionate share of riders who consume low cost recovery OD trips compared to other stations and links in the network, consumers of the subject station or link, on average, are disproportionately subsidized relative to consumers of other stations or links. Evaluating this requires weighting OD trip findings to net travel patterns (i.e., Equation 9), then estimating how cost recoveries (and subsidies) flow to different geographic areas of the network. To evaluate this, I calculate the average cost recovery of every trip that begins or ends at a station, which defines the cost recovery profile of the station (i.e., Equation 4). Similarly, I define the cost recovery profile of a link as the average cost recovery of every trip that traverses the link (i.e., Equation 5). **Figures 11** to **14** show the relationship that station and link cost recovery profiles have with the distance they are from the core station of the network (right axis) and the average trip length of all trips associated with them (left axis).

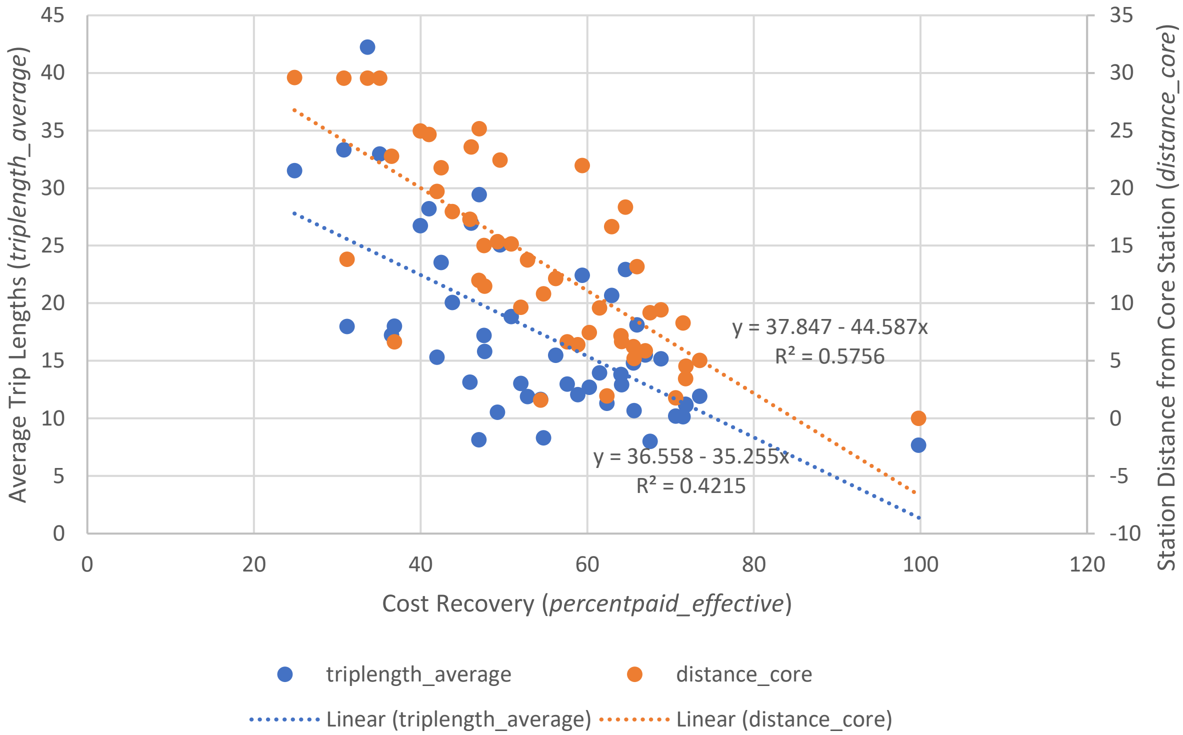

Figure 11: Station Cost Recovery Profile Relationship with Average Trip Lengths and Station Distance from Core of Network — BART

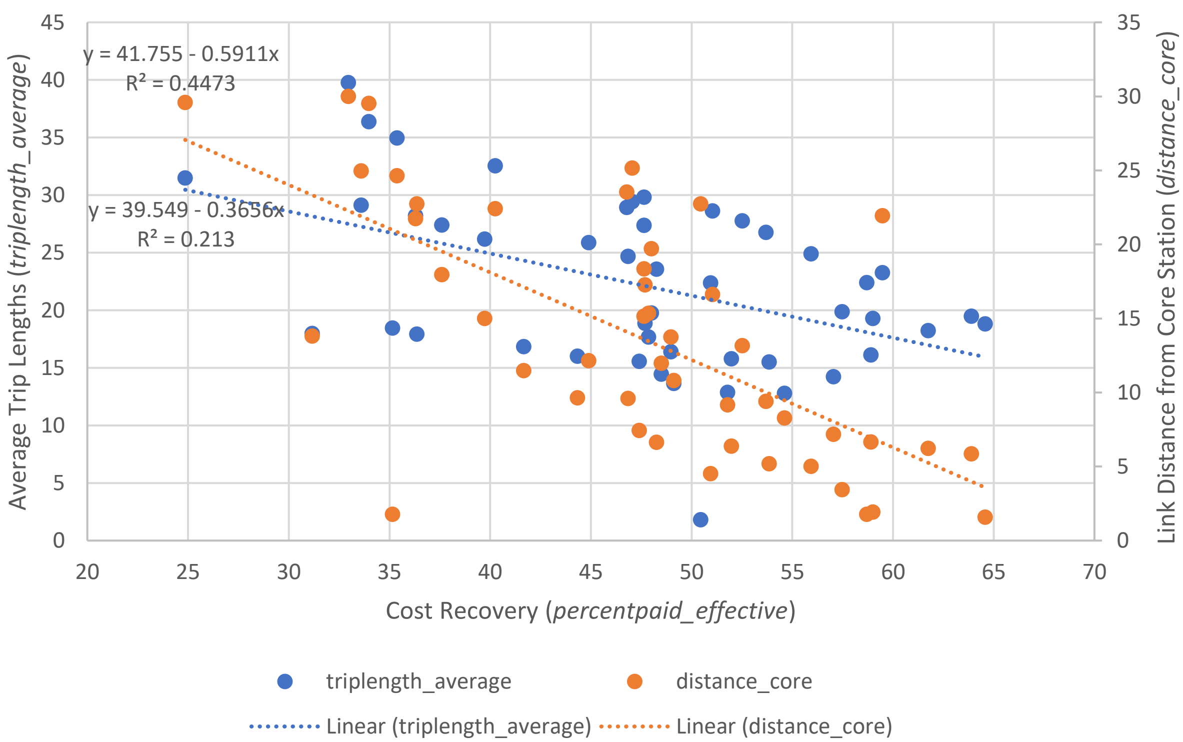

Figure 12: Link Cost Recovery Profile Relationship with Average Trip Lengths and Link Distance from Core of Network — BART

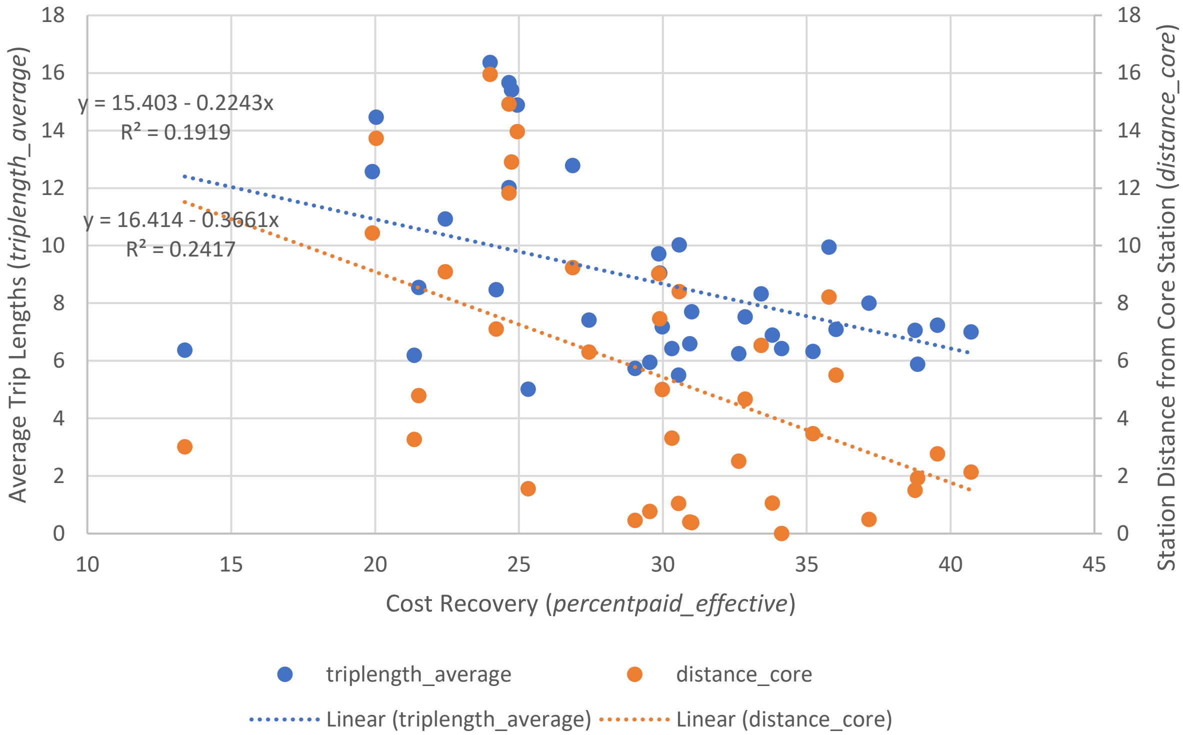

Figure 13: Station Cost Recovery Profile Relationship with Average Trip Lengths and Station Distance from Core of Network — MARTA

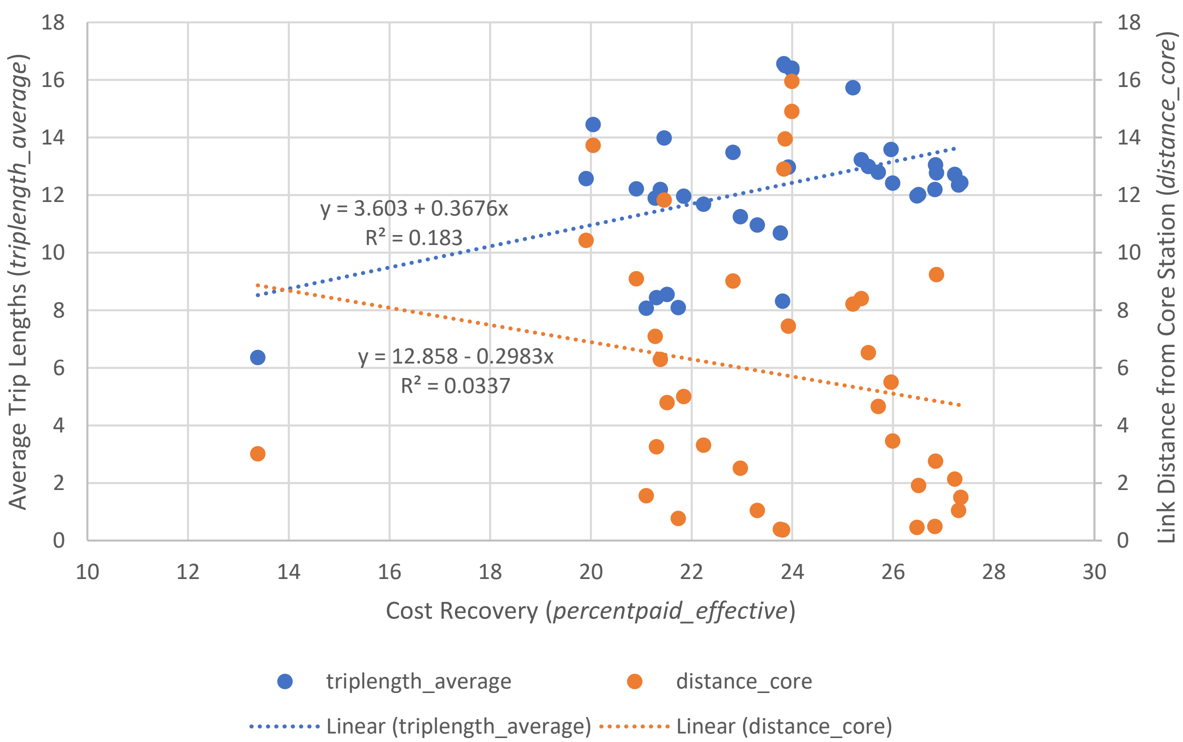

Figure 14: Link Cost Recovery Profile Relationship with Average Trip Lengths and Link Distance from Core of Network — MARTA

These plot charts show that, in almost all instances, station and link distances from the core of the network, and the average trip length for trips associated with them, have negative relationships with the cost recovery profile. That is, the further away links and stations are from the core station, and the longer the average trip associated with the link or station, the lower the cost recovery profile. The one exception is the positive influence average trip length has on link cost recovery profiles in the MARTA system. You will note that MARTA has one station and link outlier, which is Bankhead Station and the spur link that serves it. Removing this observation does not greatly change the resulting relationship.

On the other hand, the strength of the relationships is notably different between the two systems. In the BART system, both correlations have high R^2 values, especially for stations. The R^2 is consistently greater than 0.4, except for trip length's relationship with link cost recovery profiles, which is around 0.2. In stark contrast, the R^2 of the relationships for MARTA never reach 0.25 and is particularly low for the link cost recovery profile relationships. The dispersion of independent variable values across a narrow range of cost recovery profile scores explains this result for MARTA links.

Thus, with high confidence, station cost recovery profiles appear to decrease the further the stations are from the core of each system. In fact, the distance the station is from the core of the network has both stronger magnitude and statistical significance of influence on cost recovery profile than average trip length does. The same story general holds for links in the BART network. However, the relationship that distance from the system core and average trip length has on links for the MARTA system is probably too statistically insignificant to draw meaningful conclusions from.

Results — Temporal Analysis

Figures 15 and **16** show the total costs that are allocated (left axis), cost recovery (right axis), and resulting monetary subsidy (left axis) of different time periods in the BART and MARTA systems, respectively. While one time period's ridership may pay more of their costs on a percentage basis (cost recovery), an agency may still expend more money subsidizing travel during that time period (monetary subsidy).

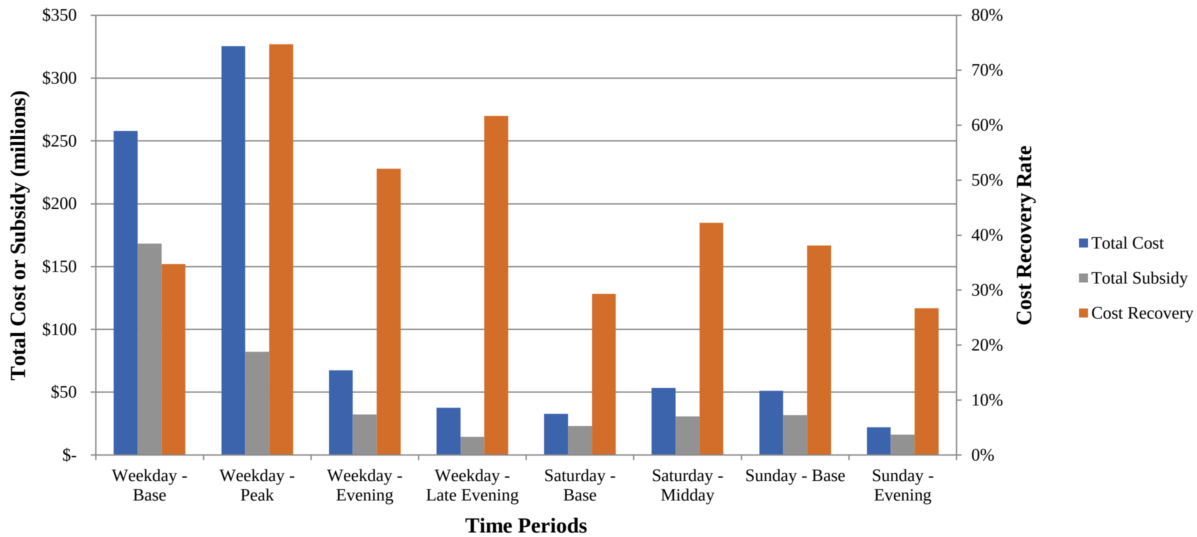

Figure 15: BART Time Period Costs, Subsidies, and Cost Recoveries

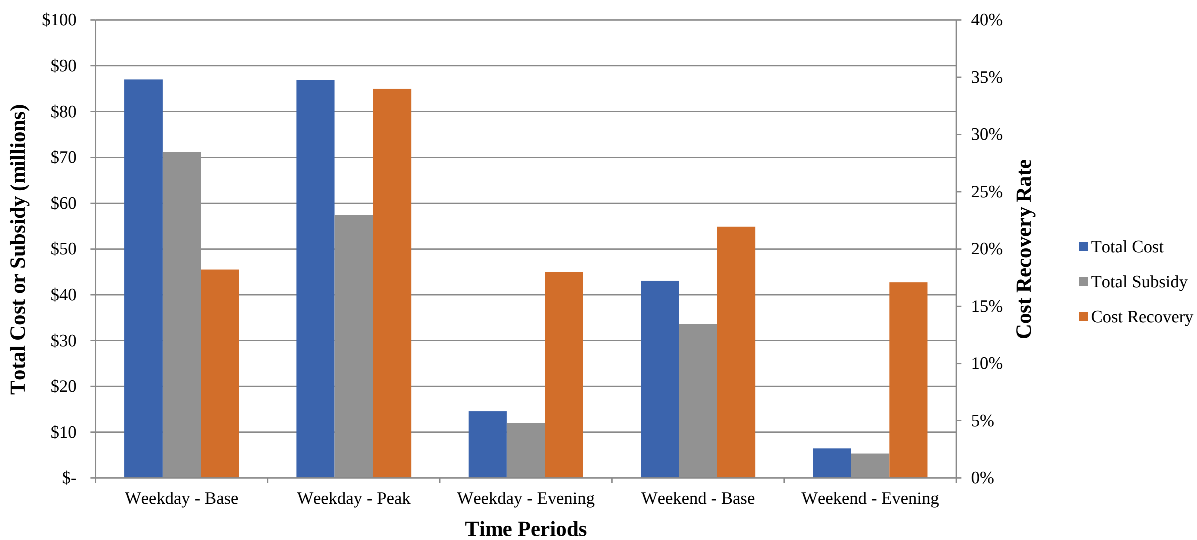

Figure 16: MARTA Time Period Costs, Subsidies, and Cost Recoveries

These findings show that, although the weekday peak period is the most expensive to serve for BART and about equally as costly as the weekday base period for MARTA, it recovers the highest amount of its costs through fare payment in both systems — 75% at BART and 34% at MARTA. In fact, the weekday peak period recovers twice as much of its costs compared to the weekday base period in both systems — so much that even the monetary subsidy expended on weekday peak service is less than weekday base service. Other time periods' total and marginal costs — that is, the additional costs of serving a time period relative to the next highest cost time period — are much smaller compared to weekday base and peak periods. Even so, it is noteworthy that the weekend base period at MARTA and Sunday base period at BART also have higher cost recoveries than the weekday base period.

Figure 17 shows the correlations between the cost recovery of different time periods (*percentpaid$_t$*) and time-variant travel patterns. In this figure, BART is represented in blue; MARTA, in red.

	percentpaid_t	trips_t	pxmiles	triplength	distancecore_o	distancecore_d
percentpaid_t						
trips_t	0.6893*					
	0.8367*					
pxmiles	0.6998*	0.9995***				
	0.8576*	0.999***				
triplength	0.3321	0.2238	0.2449			
	0.924*	0.9704**	0.9751**			
distancecore_o	-0.2828	0.3403	0.3192	-0.6138		
	0.7772	0.8831*	0.8763*	0.9273*		
distancecore_d	0.1423	-0.3625	-0.3386	0.7554*	-0.9407***	
	-0.3275	-0.5693	-0.5436	-0.5436	-0.7978	

Figure 17: Correlation Matrix of Temporal Variables

This matrix shows the relationship between variables *across time periods*. Looking at the bottom-right, at BART, time periods in which the origin or destination station of an average trip is further from the core of the network will have the other station in the OD pair closer to the core of the network and in a statistically significant way. By contrast, this OD relationship across time periods has no statistical significance for MARTA.

The only variables that relate with cost recovery of time periods in any statistically significant way for BART is the number of trips and passenger-miles generated during a time period. Specifically, time periods with more trips and passenger-miles generate higher cost recoveries, on average. For MARTA, in addition to these two influencing factors, time periods having longer trip lengths is positively associated with cost recovery.   However, with both networks, the travel pattern variables that correlate with cost recovery in a statistically significant way also correlate with each other in a statistically significant way, suggesting potential interaction between or across these terms.

Importantly, the distance that origin or destination stations of average trips are from the core does not correlate with temporal variability of cost recovery in any statistically significant way. This suggests that, to the extent cost recovery is influenced by how "suburban" a trip is, the influence does not vary across time.

A Synthesis of Cost Recovery Measurements

In the preceding sections, I introduced multiple ways of measuring cost recovery, each with a different purpose and finding. I then used a select subset of these in my analyses. The measurements of cost recovery included an aggregate or systemwide measurement, a measurement for each OD trip (unweighted and weighted to trip count), a measurement based on station cost recovery profiles, a measurement based on link cost recovery profiles, and a measurement across time periods. **Table 13** offers a brief breakdown of definitions each of these.

Cost Recovery Measurement	Description
Systemwide cost recovery	The share of total expended costs covered through fares
OD trip cost recovery	The cost of a unique OD trip that is paid for through the fare charged for consuming it.
Station cost recovery profile	The weighted average OD trip cost recovery of all trips consumed that began or ended at a particular station.
Link cost recovery profile	The weighted average OD trip cost recovery of all trips consumed that traversed a particular link.
Time period cost recovery	The share of total expended costs in a particular time period that were covered through fares generated during that particular time period.

Table 13: Cost Recovery Measurements

Systemwide cost recovery measures how much of an agency's total costs are covered by its total fare revenue. Based on the costs I included in my cost allocation study (Mallett, 2022), I find that BART and MARTA recover 52.6% and 24.6% of their costs through fares, respectively. However, as is a motivation of this research, systemwide cost recovery overlooks dimensional variability, such as spatial and temporal variability. It is possible that certain times of travel or locations of travel recover more of their costs such that these riders subsidize others. Systemwide cost recoveries combine all fare revenues and costs together, so obscure such variability.

The unweighted OD trip cost recovery measures how much of the cost of serving a particular OD trip is recovered through the fare paid for that trip. The calculation depends on dividing a system into its parts to estimate costs of each "part" (i.e., stations and links), then estimating the cost per rider of each station and link as the ratio of its cost to the number of trips that use it. I calculate a cost for every possible OD trip as the sum of the cost per rider of every link and station used to complete the trip (i.e., Equation 1). Importantly, even if no real traveler ever consumes a particular OD trip, the trip still has a cost, fare, and associated cost recovery. The unweighted OD trip cost recovery represents the cost recovery of a one-off OD trip taken in a system as it is built and priced, treating station and link usage as fixed. When weighted to actual consumption patterns, the resulting statistics will reflect net consumption patterns rather than a baseline distribution. This is reflected in Tables 3 and 4, where, for example, the unweighted mean OD trip cost recovery for MARTA is 30.1%, while the weighted mean OD trip cost recovery is 36.1%.

Both unweighted and weighted OD trip cost recovery statistics will deviate from the systemwide cost recovery principally because the source of costs is different. The systemwide cost recovery bundles all parts and their costs together, whereas OD trip cost recoveries only account for costs of the parts of the system that the OD trip utilizes. Furthermore, by using a cost per rider metric, OD trip costs account for how all OD trips share in the costs of parts of the system. When weighted to actual ridership, OD trips along high usage/low cost per rider parts of the network will influence the OD trip cost recovery statistics. In most instances, this will lead to a mean cost recovery greater than the systemwide cost recovery because higher rates of trip consumption will typically occur along low cost per rider parts of the network, whereas the systemwide cost recovery must also account for low ridership/high cost per rider segments of the network.

Station and link cost recovery profiles are the average cost recovery of all trips taken that are associated with the station or link. The purpose of this measurement is to evaluate the spatial incidence of subsidies. In the absence of being able to allocate fare revenues to links and stations like costs are so that a formal cost recovery can be calculated, taking the average cost recovery of all users of a link or station is an alternative. Users of stations and links with higher cost recovery profile scores are less subsidized, on average, because they pay a higher percentage of the cost of their trip through their fare, on average. Similarly, users of stations and links with lower cost recovery profiles are more subsidized, on average, because they pay a lower share of the cost of their trip through their fare, on average. My findings show that station distance from the core of both networks has the strongest influence on the spatial incidence of subsidies. In particular, the further away that an origin or destination station is from the core of the network, the more subsidized that station's riders are, on average.

Finally, time period cost recoveries focus on a different dimension of the network — time periods. Unlike with the spatial cost recovery metrics, there are not multiple combinations of time "parts" to define time period cost recoveries. Instead, it is merely the sum of fares collected during a time period divided by the costs allocated to that time period. My findings show that the weekday peak period has the highest cost recovery of any time period in both networks.

Discussion and Conclusion

Using a long-run partially allocated cost model, I have shown that cost recovery patterns on BART and MARTA have spatial and temporal dimensions. My research is cross-sectional and OD trip costs are defined by how travelers use the system. Thus, findings are inherently based on — and likely sensitive to — travel patterns in the networks being fixed. And by not allocating fares to links and stations, I do not estimate the cost recovery of links and stations, so focus on travelers that use each link and station instead. Nevertheless, during the study period, travelers with trips to or from outlying areas of the networks and who travel during off-peak travel periods pay a lower share of their costs relative to travelers in core areas and during peak travel periods — meaning the former are subsidized more.

In the case of MARTA, although station distance from the core has a positive relationship with cost recovery at the OD level, this appears to be driven by the positive correlation between trip length and the distance that origin and destination stations of a trip are from the core station. Given that trip length has such a large and negative influence on OD trip cost recovery, the smaller magnitude positive influence that origin and destination station distance from the core have on cost recovery effectively offsets this while still leaving their shared effect negative. The same story holds in the analysis of cost recovery profiles of stations and links, though in these instances trip length carries the positive coefficient, while distance from the core has the negative coefficient. Given this, when the distribution of consumed trips is accounted for at the link and station level, the monocentric spatial pattern of cost recovery is evident; although origins and destinations are "scattered" in the MARTA network, enough trips converge on the core links and stations of the network for some degree of central concentration in average cost recovery to exist. However, given the dispersion in MARTA travel patterns, the centralized pattern of cost recovery is significantly less for MARTA relative to BART, as evident by the small range in link and station cost recovery profiles.

Apart from the spatial finding that travelers who go to or from outlying stations pay a lesser share of costs, the influence that trip length has on OD cost recovery is negative in both networks; those who travel further on the systems pay less of a share of their costs, on average. But the influence trip length has on OD cost recovery on the MARTA system is about 32 times the magnitude relative to the BART system (Tables 11 and 12) — though, as explained above, this magnitude of difference is partly explained by the interaction between trip length and station distance from the core station at MARTA. Regardless of this caveat, this general finding suggests both that flat rate fares greatly increase trip subsidies on a per-mile basis, and this effect is greatly reduced, though not eliminated, with distance-based fares. This second implication is partly explained by the fact that those who travel longer distances tend to begin or end their trips in outlying areas of both networks, where there tend to be fewer riders, making them more expensive to operate on a per-rider basis (Mallett, 2022). A distance-based fare structure does not account for this variability in per-mile and per-rider costs throughout the network, so does not correct for the spatial variability in cost recovery patterns. BART's distance-based fare structure rewarding distance through its stepwise formula amplifies this. The net result in both systems is that the spatial cost recovery pattern leads to the incidence of transport subsidies being geographically concentrated in suburban and exurban areas of the respective region; persons who begin or end their trips in outlying areas of the networks are disproportionately subsidized.

The finding that off-peak travel is more subsidized than weekday peak period travel runs counter to findings from past research. Cervero (1981), Parody et al. (1990), and others show that the peak period is the costliest to serve in net terms — meaning, even after cost recovery is considered. This is partly because these other studies include fixed asset costs in their cost allocations, whereas I do not. However, as shown in the cost allocation research (Mallett, 2022), this divergence in findings is also explained by the fixed headway schedule that BART and MARTA operate throughout the day. Because of this practice, the weekday peak period has a low marginal cost of operating compared to what is typical of traditional commuter rail systems and bus networks. At BART, the additional cost of serving the weekday peak period is almost exclusively driven by capital because the agency resizes the length of its trains without compromising frequency of service. By comparison, MARTA has a nominal number of additional trains in service during weekday peak period relative to weekday base period and does not resize its trains, leading to an even smaller marginal gross cost difference — so small that the additional operating hours associated with base service makes it slightly more expensive than the weekday peak period.

The correlations between time period cost recoveries and travel pattern variables associated with different time periods indicates that the number of trips and passenger-miles generated across time periods explains much of the difference in time period cost recoveries.

Given the relatively fixed headway schedules of both agencies, this makes sense; holding all else equal, more trips and trip-miles will result in more revenues, which will increase cost recovery. This is especially true for BART, given its fare structure being distance-based. On the other hand, the distance that origin and destination stations are from the core of the respective system does not have any statistically significant influence on temporal cost recovery variation. Indeed, there is little variability in these terms across time periods. Considering that spatial patterns of cost recovery are negatively associated with the distance origin and destination stations are from the core in both networks, this suggests that outlying areas are subsidized regardless of the time of travel. Thus, while there is geographic incidence of subsidies, it is not temporally variable. Even so, research that explicitly interacts space and time could more decisively test this.

While this research shows that core stations and links and the weekday peak period are less subsidized than suburban stations and links and other time periods, there may be bases for charging premiums for travel in these areas and during these times. My cost allocation study (Mallett, 2022) does not account for externalities, including the costs of congestion beyond expenditure costs of serving it. Core areas of these networks and peak times of travel may generate so much crowding and inconvenience costs for passengers — for example, BART passengers departing Downtown San Francisco during the evening commute often backtrack to secure a seat — that a premium is warranted for use in these areas or during these times to internalize delay time costs and manage demand. In addition, minimum fares or trip length restrictions are often used to manage capacity or product differentiate one service from others — for example, a regional service relative to a local service. Among other examples, BART's minimum fare assumes a trip of at least six miles to distinguish its regional focus from peer agencies' local focus, and New York City's Metro-North Railroad does not sell fares for travel between Harlem/125th Street and Grand Central Terminal Stations to distinguish its commuter rail-orientation from New York City Transit's local travel focus. In this analysis, I do not control for these policies or objectives, which can inflate the cost recovery level found for peak period travel and travel in core areas of each network.

Future research can also benefit by interacting these findings with the socioeconomic makeup of riders to evaluate whether cost recovery equity patterns in these networks has a disparate impact. Do different socioeconomic groups of riders consume more subsidized OD pairs than others, such that they are more subsidized, on average? While both BART and MARTA survey riders to create socioeconomic profiles of stations, different socioeconomic groups consume different OD pairs from each station. So, although there is a sufficient sample size of riders at each station to devise station profiles, there is not a sufficient sample of riders of different OD pairs to evaluate the incidence of trip subsidies.

Finally, this analysis uses rail transit as a case study of how transport subsidies more broadly have a spatial and temporal dimension. It is conceivable that the findings from this research similarly apply to other modes of travel, such as highways. Future research should explore if other modes of transport have spatial and temporal dimensions of subsidies. In addition, these findings and findings of research on other modes can be used to evaluate a bigger question this research contributes: To what extent are suburban and exurban location choices enabled through transport subsidies?

Bibliography

Altshuler, Alan. "Equity, pricing, and surface transportation politics." *Urban Affairs Review* 46.2 (2010): 155-179.

Altshuler, Alan A., and David E. Luberoff. Mega-projects: The changing politics of urban public investment. Brookings Institution Press, 2004.

Brown, Anne E. "Fair fares? How flat and variable fares affect transit equity in Los Angeles." Case Studies on Transport Policy 6.4 (2018): 765-773.

Cervero, Robert. "Flat versus differentiated transit pricing: what's a fair fare?." Transportation 10.3 (1981): 211-232.

Fielding, Gordon J. "Transit in American Cities." The Geography of Urban Transportation. (1995): 287-303.

Frisk, Mikael, et al. "Cost allocation in collaborative forest transportation." European Journal of Operational Research 205.2 (2010): 448-458.

Hodge, David C. "Fiscal equity in urban mass transit systems: A geographic analysis." Annals of the Association of American Geographers 78.2 (1988): 288-306.

Iseki, Hiroyuki. "Equity in regional public transit finance: Tradeoffs between social and geographic equity." Journal of Urban Planning and Development 142.4 (2016)

Jones Jr, David W. "Urban transit policy: An economic and political history." (1985).

Garrett, Mark, and Brian Taylor. "Reconsidering social equity in public transit." Berkeley Planning Journal 13.1 (1999).

Giuliano, Genevieve. "Low income, public transit, and mobility." Transportation Research Record 1927.1 (2005): 63-70.

Guajardo, Mario, and Mikael Rönnqvist. "A review on cost allocation methods in collaborative transportation." International transactions in operational research 23.3 (2016): 371-392.

Mallett, Zakhary. "Spatial and Temporal Variability of Rail Transit Costs and Cost Effectiveness." Transportation Research Record (2022): 03611981221104807.

Meyer, John R., and Jose A. Gomez-Ibanez. "Autos, transit, and cities." Autos, Transit, and Cities. Harvard University Press, 1981.

Morales Sarriera, Javier, and Frederick P. Salvucci. "Rising costs of transit and Baumol’s cost disease." Transportation Research Record 2541.1 (2016): 1-9.

Parody, Thomas E., Mary E. Lovely, and Poh Ser Hsu. "Net Costs of Peak and Offpeak Transit Trips Taken Nationwide by Mode." Transportation Research Record.1266 (1990): 139-145

Pickrell, Don H. "Rising deficits and the uses of transit subsidies in the United States." Journal of Transport Economics and Policy (1985): 281-298.

Pucher, John. "Discrimination in mass transit." Journal of the American Planning Association 48.3 (1982): 315-326.

Reilly, J. M. Transit Costs During Peak and Off-Peak Hours. Transportation Research Record 625 (1977): 22–26.

Rosenthal, Edward C. "A cooperative game approach to cost allocation in a rapid-transit network." Transportation Research Part B: Methodological 97 (2017): 64-77.

Rubensson, Isak, Yusak Susilo, and Oded Cats. "Is flat fare fair? Equity impact of fare scheme change." Transport Policy 91 (2020): 48-58.

Savage, Ian. "The analysis of bus costs and revenues by time period. II. Methodology review." Transport Reviews 9.1 (1989): 1-17.

Schweitzer, Lisa. “Mass Transit.” The Geography of Urban Transportation (2017): 187.

Serebrisky, Tomás, et al. "Affordability and subsidies in public urban transport: What do we mean, what can be done?." Transport reviews 29.6 (2009): 715-739.

TAYLOR, BRIAN D. "The Geography of Urban Transportation Finance." The Geography of Urban Transportation (2017): 247.

Taylor, Brian D., Mark Garrett, and Hiroyuki Iseki. "Measuring cost variability in provision of transit service." Transportation Research Record 1735.1 (2000): 101-112.

Taylor, Brian D., and Eric A. Morris. "Public transportation objectives and rider demographics: are transit's priorities poor public policy?." Transportation 42.2 (2015): 347-367.

Taylor, Brian D., and Alexandra Tassiello Norton. "Paying for transportation: What's a fair price?." Journal of Planning Literature 24.1 (2009): 22-36.

Taylor, Brian D., and Kelly Samples. "Jobs, jobs, jobs: Political perceptions, economic reality, and capital bias in US transit subsidy policy." Public Works Management & Policy 6.4 (2002): 250-263.

United States Census Bureau, 2017. Table 1. State and Local Government Finances by Level of Government and by State: 2017. [XLS] Available at: <https://www2.census.gov/programs-surveys/gov-finances/tables/2017/summary-tables/17slsstab1a.xlsx> [Accessed 16 June 2022].

Wachs, Martin. "US transit subsidy policy: In need of reform." Science 244.4912 (1989): 1545-1549.

Zhao, Pengjun, and Yixue Zhang. "The effects of metro fare increase on transport equity: New evidence from Beijing." Transport Policy 74 (2019): 73-83.