Does *Uber* reduce public transit ridership? Evidence and impacts in the San Francisco Bay Area

Laura Grant* and Lianne Sturgeon†‡

September 2024

Abstract

Many public transit authorities believe ride-hailing reduces transit ridership and threatens solvency. Municipalities may face rideshare-induced increases in congestion and pollution. We develop methods to measure metro-specific, hourly impacts using a decade of route-by-hour data from San Francisco Bay Area transit. We leverage weather and traffic as exogenous, hourly route-specific shocks, interacting them with ride-hailing availability. Our analysis reveals large short-run decreases in ridership. Comparing ridership trends before and after ride-hailing, we find long-term year-over-year losses exceeding 10 percent. Fare revenue loss and social costs are *each* $100 million in 2013, growing to $200 million by 2018.

Keywords: Public transportation; Ride-hailing; Disruptive technology; First/last mile; Externalities of driving.

JEL Codes: R40; H42; O33; Q51

*Claremont McKenna College, Robert Day School of Economics and Finance: `lgrant@cmc.edu`.

†Scripps College 2019: `liannerenees@gmail.com`.

‡This research grew from Sturgeon's undergraduate thesis: *The Impact of Transportation Network Companies on Public Transit: A Case Study at the San Francisco International Airport* (Sturgeon, 2019). The advisee-advisor relationship evolved into co-authorship. We appreciate feedback from Sturgeon's co-advisor Roberto Pedace, Mary Evans, Joe Price, seminar attendees at Claremont McKenna and Arizona State University, and participants at the Association of Environmental & Resource Economists 2021 Summer Conference and the ASSA-AEA 2022 Meeting. Any errors are our own. Appendix available on request. Generous research support provided by the Lowe Institute at CMC. Declarations of interest: none.

1 Introduction

Uber Technologies Inc. pioneered the first widely available ride-hailing app, revolutionizing taxi-like services. The arrival of its main competitor, *Lyft*, cemented this new mode of transportation in the market. Founded in March 2009 and June 2012, respectively, these companies have since become dominant players, reshaping consumer transportation habits.[1] Their services disrupted the traditional taxi industry, which is struggling to compete with their affordability and convenience (Berger et al., 2018; Cramer and Krueger, 2016; Wallsten, 2015). Many public transit authorities are concerned about potential negative impacts on ridership. Anecdotal evidence and studies suggest that *Uber* and *Lyft* have lured passengers away from public transit (Clewlow, 2019).[2] However, others argue that ride-hailing complements public transit by delivering passengers to transit-serviced areas, offering on-demand rides that coordinate with public transit schedules (Feigon and Murphy, 2016; Rayle et al., 2016). To assess ride-hailing's effects on transit ridership, we develop methods for measuring its metro-specific, hourly, and daily impacts.

The emerging body of empirical studies disagrees on whether and to what extent *Uber* and *Lyft* compete with or complement public transportation. Hall et al. (2018) provide an important starting point by using variation in *Uber*X entry across U.S. cities to estimate a difference-in-differences model of monthly public-transit ridership. Analyzing two years of data post-*Uber*X market entry, they find a complementary relationship, with public transit ridership increasing by 5 percent on average. However, their results differ by city type: cities with above-median ridership show a decline, while cities with above-median populations see an increase. Thus, for a place like San Francisco, which is above both medians, the implications are ambiguous. Similarly, Babar and Burtch (2017) find complementarity between ride-hailing companies and rail transit, adding covariates to explain further aspects of the relationship. In contrast, Nelson and Sadowsky (2018) use a regression discontinuity design to show that *Uber*X initially complements transit, but after *Lyft* enters the market, ride-hailing becomes a substitute. They argue that competition between the two companies drives down ride-hailing prices, making it a more attractive option than public transit.

[1] These companies, often referred to as Transportation Network Companies (TNCs), provide an internet-based mobile app that connects passengers with nearby drivers. In most urban areas, drivers can arrive within minutes, offering convenient transport to a chosen destination.

[2] See Ride-hailing is pulling people off public transit and clogging up road and MTA Blames Uber for Decline in New York City Subway, Bus Ridership, accessed October 2024.

Diao et al. (2021) extend this work, corroborating the finding of substitution, reinforcing competitive dynamics between ride-hailing and transit. On the other hand, Coogan et al. (2018) estimate cross-price elasticities between rail and ride-hailing, finding no significant relationship, suggesting ride-hailing and public transit are neither strong complements nor substitutes. In Pittsburgh, Grahn et al. (2020) find that *Uber*'s surge pricing correlates with bus boardings, but they cannot detect complementarity due to a shared "pie" of consumers. Similarly, Erhardt et al. (2021) study San Francisco's Muni ridership but find limitations in detecting complementarity due to aggregated data that masks nuanced patterns.

Our work addresses gaps in these studies by analyzing detailed route-by-hour data over ten years rather than using monthly aggregates over a few years. This allows us to capture both long-run trends and short-run variations. We also split the data into subsets such as commuters and airport travelers, estimating who is affected, when, and for how long. More importantly, we identify the causal mechanisms driving trade-offs between public transit and ride-hailing, directly addressing how ride-hailing affects public transit usage.

Figure 1 demonstrates the short-run fluctuations, month-to-month, and also the long-run trends, supporting approaches for with-in hour and extended time period estimations. Before the arrival of *Uber*X and *Lyft*, BART ridership grew by 0.2 percent per month. In the period from mid-2012 to mid-2014, BART ridership maintained an upward trend, though less steep. After ride-hailing became more established in 2014, BART ridership reversed trend, declining by 0.2 percent monthly. By the end of the period, BART ridership had dropped by 60 to 85 thousand daily passenger trips compared to the pre-ride-hailing trend.

In the next section, we introduce a conceptual framework that allows for both substitution and complementarity between public transit and ride-hailing. Each mode has distinct characteristics that vary by time of day and location. Public transit is affordable, and rail systems avoid traffic congestion, but transit is rigid in timing and location, prone to delays, and can be uncomfortable. Ride-hailing, while more expensive, offers point-to-point service on-demand, although it too can be subject to traffic. Consumers weigh these trade-offs to make optimal choices. Factors like inclement weather or fluctuating traffic speeds affect the costs and benefits, making weather and travel time key variables. Additionally, ride-hailing may reduce the interdependence between morning and evening commutes, as consumers are no longer dependent on public transit for both trips. This framework forms the basis for our hypotheses on consumer preferences.

Figure 1. BART ridership through time versus ride-hailing growth, by month-year

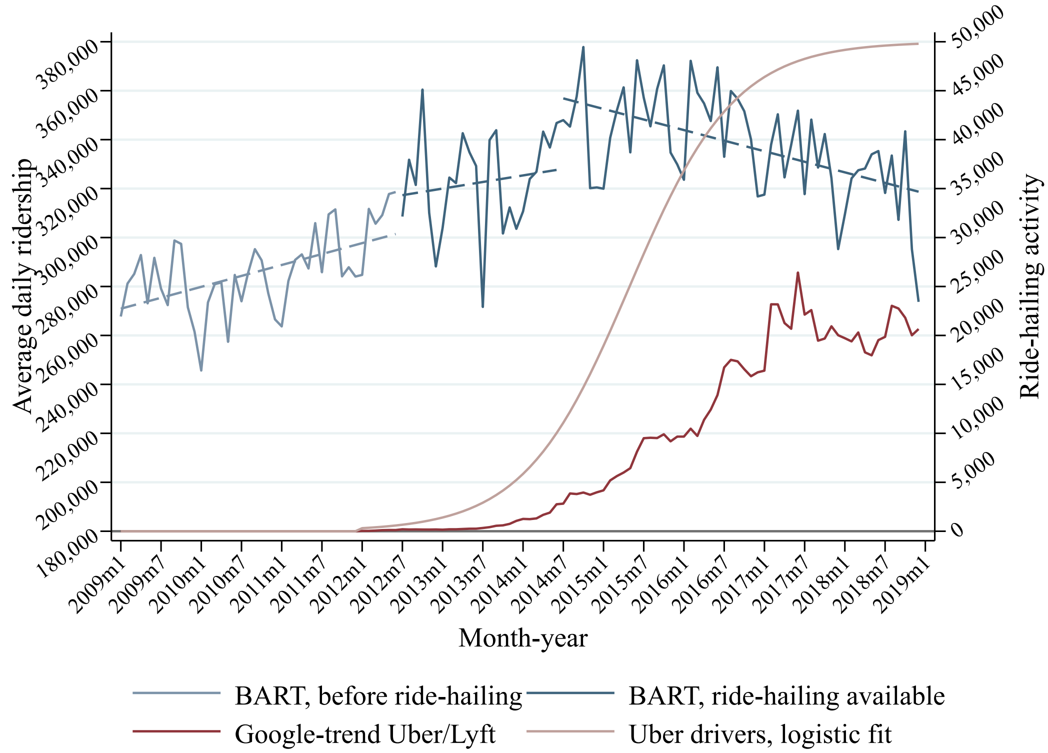

Notes: Daily count of BART ridership on the left vertical-axis; truncated. Ride-hailing activity on the right vertical-axis is represented by (1) Google searches and (2) a logistic curve fit to the count of active Uber drivers. Trend lines before ride-hailing, during early adoption, and post.

We estimate ride-hailing's impact on San Francisco's Bay Area Rapid Transit (BART) system. Using a difference-in-differences approach, we first measure heterogeneous treatment effects by hour-route, with outdoor temperature and traffic speed as exogenous shocks. This variation allows us to capture how fluctuations in weather and congestion affect transit ridership once ride-hailing is available. We also examine whether ride-hailing affects the relationship between morning and evening commutes, with *Uber* potentially decoupling these modes. A second analysis compares daily ridership trends before and after ride-hailing, controlling for covariates, and examines heterogeneity across time of day, day of week, and location, as proxies for travel purpose.

We find that inclement weather and faster driving speeds reduce BART usage more when ride-hailing services are available than without. Additionally, 1 in 100 daily commuters that use public transit in the morning now seem to switch to ride-hailing after work. In the long run, ride-hailing reduces BART ridership by 9.8 to 10.7 percent year over year. We also find impacts during the first two years, that ride-hailing substitutes for public transit in San Francisco, which provides an update to prior studies. Commuter hours see only modest declines, likely due to inelastic demand, while the nighttime and airport transit sectors experience large drops in ridership.

Each public transit authority must understand ride-hailing's effects to manage short-term demand changes and plan for long-term operational and capital funding. Accurate estimates of ridership variation are essential for informed public policy. This research also has implications for social costs, as ride-hailing contributes to road degradation, congestion, collisions, and pollution, and can inform regulatory decisions. Our methodology can help municipalities evaluate ride-hailing's effects on revenue, congestion, and environmental impacts, assisting in policy development and coordination with ride-hailing companies.

2 A conceptual framework for transit demand

A rich literature in urban economics provides transit-demand models that guide our hypotheses. Many studies estimate discrete choices where a rider weighs the characteristics of their transit options and selects the optimal one. We build on this literature to identify hourly and daily causal mechanisms, which we apply in our panel-estimation methods, along with other variables available to municipalities. For instance, McFadden et al. (1977) forecast the demand for BART before its construction, using a discrete-choice framework.

A consumer trip from point A to B may involve one or more modes, such as driving, ride-hailing, carpooling, public transit, walking, biking, and parking. Each option has several negative impacts. Disutility and budget arise from both direct and indirect costs, including: fuel or fare; travel time; waiting time; discomfort due to weather, climate, and crowding; risk of delay and associated stress; risk of harm from collision or assault; time spent finding parking; and parking fees. By minimizing these costs, consumers make specific transit choices. In this analysis, we focus on consumers choosing between public transit (via rail) and ride-hailing, assuming these are the two best options. As in previous literature on transit choice, e.g., McFadden et al. (1977), we assume that a consumer's choice is independent of irrelevant alternatives, thereby ruling out other transit modes.

Ride-hailing complements public transit when passengers combine the two modes. For example, public transit users often face challenges in getting from their location to a transit station or from the station to their final destination. This is referred to as the "first/last mile" problem. Ride-hailing can bridge this gap, providing geographic flexibility to the fixed routes of public transit (Wang and Odoni, 2014). In some cases, a person might use public transit one way and ride-hailing the other, amplifying the utility of both modes.

Substitution occurs when one transit option outweighs the other in overall utility. Public transit generally offers lower direct costs compared to ride-hailing.[3] However, ride-hailing outperforms public transit in several key dimensions. Its innovations include greater convenience, as rides are on-demand, hailed via smartphone apps, and offer point-to-point service with real-time updates on arrival and travel time. Compared to public transit, ride-hailing also may reduce the overall time spent both in the outdoor elements while waiting and in transit, conditional on traffic congestion. Ride-hailing provides more comfort, with better seating, privacy, and climate control. As a result, passengers substitute public transit for ride-hailing when they perceive that these added amenities outweigh the higher price.

If ride-hailing significantly alters consumer disutility compared to public transit, we will observe changes in ridership.[4] We hypothesize that these effects vary by time of year, day of the week, and time of day, as well as by origin and destination.

3 Empirical context and data descriptions

Our empirical setting is the San Francisco Bay Area. We gather data spanning a decade, from 2009 to 2018, with source descriptions provided in Table 1. The primary outcome

[3] Polzin (2016) estimates public transit costs at \$0.26 per mile, while ride-hailing ranges from \$0.65 to \$2.00 per mile, depending on location and time of day. For example, in 2017, a weekday 4 p.m. trip from Embarcadero Station to San Francisco Airport costs \$8.95 via BART and approximately \$29.73 via *Uber*X. Similarly, a trans-bay trip from Downtown Berkeley Station to San Francisco Airport costs \$9.55 via BART and approximately \$47.92 via *Uber*X; `www.sfoconnect.com/sites/default/files/document/TNCWhitePaper.pdf`, accessed October 2024.

[4] Some consumers are infra-marginal, meaning they always choose public transit. For these riders, ride-hailing and public transit neither compete nor complement. This can dampen our estimates, as we are measuring average treatment effects.

of interest is Bay Area Rapid Transit (BART) ridership, measured as hourly counts of passengers traveling from each origin station to any destination station. The key treatments are the introduction and prevalence of ride-hailing services. Additionally, we compile covariates that influence BART ridership and demand for ride-hailing.

Figure 2. BART route map, effective September 2018

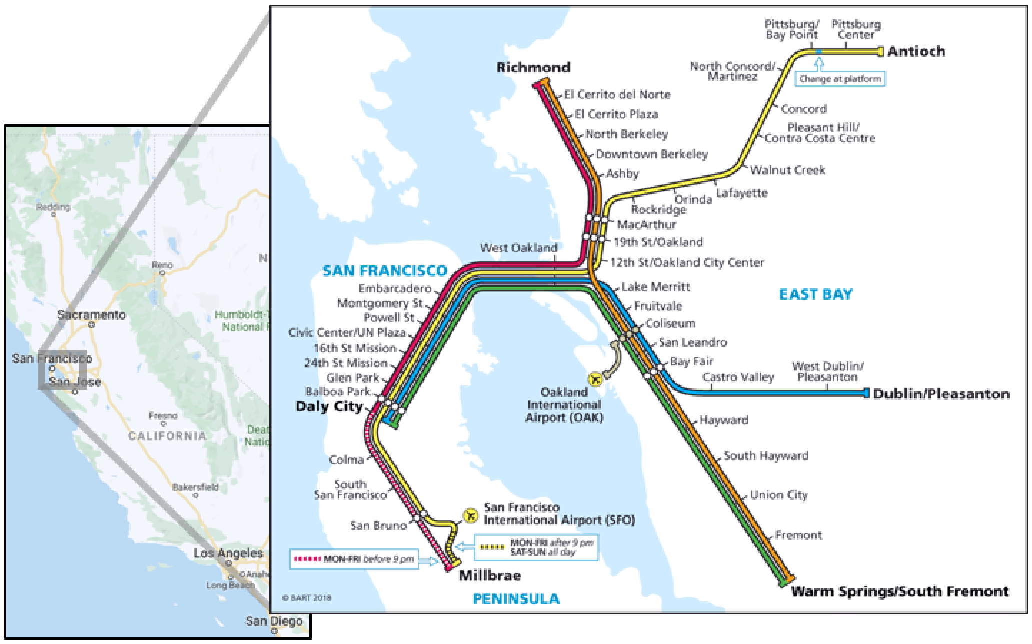

BART is an electric railway system serving San Francisco and the greater East Bay; see Figure 2. The system has 48 stations, connecting eight endpoints. Each pair of stations, defined as an origin and a destination, represents a distinct route. For our analysis, we focus on the 44 stations in operation by 2014, resulting in 1,892 distinct routes in our dataset. BART regularly collects and publishes data on ridership, performance, safety, and other factors, which are available at `data.bart.gov`.

Our primary dataset consists of 137.4 million hourly observations (Table 1). For some empirical models, we aggregate these data to daily counts of trips-by-route, yielding over 6.6 million observations. On average, BART serves nearly 310,000 passenger-trips per day. As shown in the introductory Figure 1, ridership increased from 2009 through 2013, peaking at approximately 380,000 daily passenger-trips in early 2014, before declining in recent years. BART ridership also exhibits seasonal fluctuations: November, December, and January see fewer daily passenger-trips, while ridership tends to be higher during June and from August through October. The number of daily passenger-trips per route vary widely, from zero to over 6,000, with an average of 170 trips per route-day.

Table 1. Summary statistics

Continuous variables	Obs.	Mean	S.D.	Min.	Max.
BART ridership					
Daily total count	3,638	308,205	109,730	2,670	548,731
Route-by-day count	6,622,374	169.31	305.69	0	6,374
Route-by-hour count	137,418,912	8.16	25.91	0	1,826
Route-by-hour count, **by group**:					
SFO airport, to	3,231,169	6.13	12.20	0	429
SFO airport, from	3,195,787	6.86	14.72	0	530
Commuters, AM	2,358,119	87.66	98.69	0	1,317
Commuters, PM	2,360,327	84.54	81.26	0	1,463
Recreational, day	49,165,317	6.70	14.19	0	1,826
Recreational, night	18,548,655	4.81	13.53	0	1,366
All other route-hours	62,029,327	4.55	12.88	0	1,539
Avg temperature F°, route-hour	137,418,912	59.209	8.958	25	109.9
Rain ″, hour	87,648	0.002	0.012	0	0.380
Highway speed, mph, route-hour	136,004,391	62.0	6.4	12.9	78.3
Population, 1000s, month	120	4,417.1	154.5	4,188.1	4,673.2
Employment, 1000s, month	120	979.8	101.9	845.0	1,156.6
VMT, 1000s, day	6,622,374	59,125	5,663	32,620	71,088
Gas price, real $, week	520	3.738	0.587	2.34	5.155

Indicator variables, as proportions

BART characteristics		Other characteristics	
Workday	0.743	Ride-hailing in SF	0.663
Weekend	0.257	Ride-hailing widely available	0.466
BART holiday	0.019	Federal holiday	0.027
Busy SFO travel date	0.069	Recreational holiday	0.011
		Rainy hour	0.067

Note: Data from January 1, 2009 to December 31, 2018. The time-unit of observation for each continuous variable is listed after the variable name. VMT is vehicle miles traveled, as measured on surface streets in the SF-OAK metro areas. Proportions calculated from hourly data.

The hourly ridership average, pooled over all data, is 8 passengers per route. However, this varies widely, from none to 1,826 on a single route. We split ridership into groups based on hourly patterns in our data and using consumer surveys. We establish the following seven groups: 1) airport transit to SFO, 2) airport transit from SFO, 3) morning commuters, 4) evening commuters, 5) daytime recreation, 6) nighttime recreation, and 7) all other route-hours; details provided in the appendix, with summaries in Table A-2 and Figure A-2.[5] *SFO Airport* and *Recreational* route-hours and *All other* ridership each have 5 to 7 passengers per route-hour; Table 1. *Commuter* route-hours have 12 to 19 times higher BART ridership, with 88 and 85 passengers per route-hour, AM and PM respectively; observations are the three-hour windows, morning and evening, and 314 routes heading into or out of the downtown core.

Several indicator variables account for BART schedules and rider demand, given as proportions in Table 1. *Workdays* indicate Monday through Friday, which have more ridership than *Weekends*. Three types of holidays, *BART holiday*, *Federal*, and *Recreational*, allow differential responses to each. In addition, we create an indicator for *Busy SFO travel dates*, based on transit to/from the airport, which tend to be days around holidays.

Uber Technologies Inc formed in 2009 and began operating in San Francisco in 2011, offering luxury, black-car services. Uber's popularity rose in mid-2012 with the introduction of *Uber*X, a cheaper service; *Lyft* entered simultaneously in San Francisco. Sixty-six percent of our observations occur after the two services launched in San Francisco on July 1, 2012 and are indicated by *Ride-hailing in SF*. Yet, adoption of ride-hailing technology by drivers and consumers is gradual. Thus, we divide our decade of data into three phases: before ride-hailing, adoption of ride-hailing, and mature. The before ride-hailing phase occurs before *UberX* and Lyft launch in SF. The ride-hailing adoption phase is when the number of drivers begin to take off and consumers start using the service more regularly. The data on ride-hailing activity indicate that two years pass before Uber and Lyft are prominent and rooted in SF. Therefore, the period after July 1, 2014 is the mature phase, when ride-hailing has a large supply and demand.

We use two exogenous variables — weather and highway traffic speeds — to identify mechanisms by which ride-hailing changes transit ridership. We collect these data by hour and at numerous locations, to create cross sectional variation. These aspects directly affect transit and interact with the choice to use ride-hailing, once the services become available. Our route-by-hour weather measurement is the temperature at the weather station nearest

[5] BART opened a station at Oakland Airport after ride-hailing entered the market, data which we include in 'all other route-hours' and which contributes to estimation of rider trends.

the BART entry, lagged by one hour, averaged with the temperature nearest the BART exit. Average temperatures hover near 60° Fahrenheit, with a standard deviation of 9°. Around seven percent of all date-hours have rainfall. Finally, we create a measure for driving speed, for each date-hour by route, in miles per hour (mph). Speed proxies for its inverse, driving travel time, which affects choice of transit. We collect traffic speed data measured by highway sensors near a route's origin and destination, then average the speeds that match direction of travel for the route. Additional details on these variables are in A.1.

Additional covariates capture important trends that affect transit, including monthly data on population, employment, and vehicle miles traveled, and local gas price by week, converted to 2018$; see Summary Statistics, Table 1. These are necessary for our long-run analysis, which compares time trends before and after ride-hailing. For robustness, we also test numerous other variables, which we find do not impact our estimated impacts on ridership; these include train on-time performance metrics, BART fares, BART fatalities, crime, air traffic, the #deleteUber twitter movement, and amount of daylight. For our short-run analysis, we drop the monthly time-series variables and exploit several forms of time fixed effects, which capture constants-by-time, including unobserved trends.

4 Empirical analysis

In this section, we analyze the short-run impacts of ride-hailing on BART ridership and then assess long-run trends. For each, we lay out our specification and determine the effects for the full sample of data. We follow with results by ridership group — commuters, recreational, and airport transit — to assess heterogeneous effects on trips by time of day, purpose, and location.

4.1 Short-run: heterogeneous-treatment, difference-in-differences by cause

We determine how ride-hailing affects BART ridership, as identified by variation in two exogenous mechanisms — weather and traffic speeds. Both weather and traffic speed have panel-level variation by route-hour. We implement heterogeneous-treatment, difference-in-differences to test for the effects on public transit ridership, before and with the availability of ride-hailing. Variations in these factors each cause fluctuations in ridership by changing the trade-offs between modes of transit. We estimate how a factor affects transit before ride-hailing.

We use a panel regression model:

$$\mathrm{bart}_{ih} = [\mathrm{F}_{ih}]' * \boldsymbol{\alpha} + [\mathrm{F}_{ih}]' * \mathrm{RH}_h * \boldsymbol{\Gamma} + X'_h * \boldsymbol{\zeta} + \boldsymbol{\gamma_i} + \boldsymbol{T}_h + \epsilon_{ih}.$$

The outcome variable, $bart$, is hourly BART ridership, indexed by route i and hour of date h. The specification is difference-in-differences, before versus during ride-hailing availability. The heterogeneous treatments are the non-parametric indicators for binned values of our two main factors of interest: weather and speed. The $[F]$ is a placeholder for each factor; these variables are continuous from which we construct bins as binary variables. We estimate a separate specification for each causal factor. The dummy variable RH indicates ride-hailing is available, thus RH takes the value 1 if the observation occurs July 1, 2012 or later. $\boldsymbol{X}_h'$ is a collection of covariates. The $\boldsymbol{\gamma_i}$ is a vector of route fixed effects. Time fixed effects, $\boldsymbol{T}_t$, account for month-year, day of week, and hour of day. The coefficients in vector $\boldsymbol{\alpha}$ are the effects of the $[F]$-bins on $bart$ ridership before ride-hailing entered the transit market. Our primary estimates are the $\boldsymbol{\Gamma}$ coefficients on the RH-$[F]$-bin interaction terms, as they measure how BART ridership changes when ride-hailing is available. Positive value implies that ride-hailing complements public transit, whereas, a negative value signals substitution. We use ordinary least squares estimation and cluster the standard errors by route.

Our first factor is weather. People using public transit are more exposed to the weather than if driving or riding in a car — facing bigger repercussions from precipitation and variation in temperature, in particular, when getting to and from transit stations. We expect that extreme temperatures or rain reduce BART-ridership demand. Ride-hailing offers shelter from the storms and temperature swings, thus we hypothesize that when these undesirable weather factors are present, ride-hailing is a substitute for public transit. Furthermore, weather is exogenous, providing a plausible causal mechanism of how ride-hailing and public transit interact. We use interactions of weather variables with ride-hailing availability to identify this mechanism. Specifically, we create indicator bins for every two-degrees, by each hour: $\sum_{d=38}^{78} \text{bin}_{d,ih}$. We exclude the bin centered on 73°F for before ride-hailing. We include an indicator variable for a rainy hour, in the $[F]$ matrix; therefore we measure the interaction of rain with ride-hailing availability as well. Covariates include the proportion of hour day-lit, hourly route-specific travel speed, and indicators for holidays and busy travel days. Time fixed effects for each month-year, day of week, and hour of day control for underlying trends, seasonal and weekday transit patterns, and intra-day demand that are common across all routes.

Figure 3. BART ridership responses to temperature,
before and with ride-hailing availability

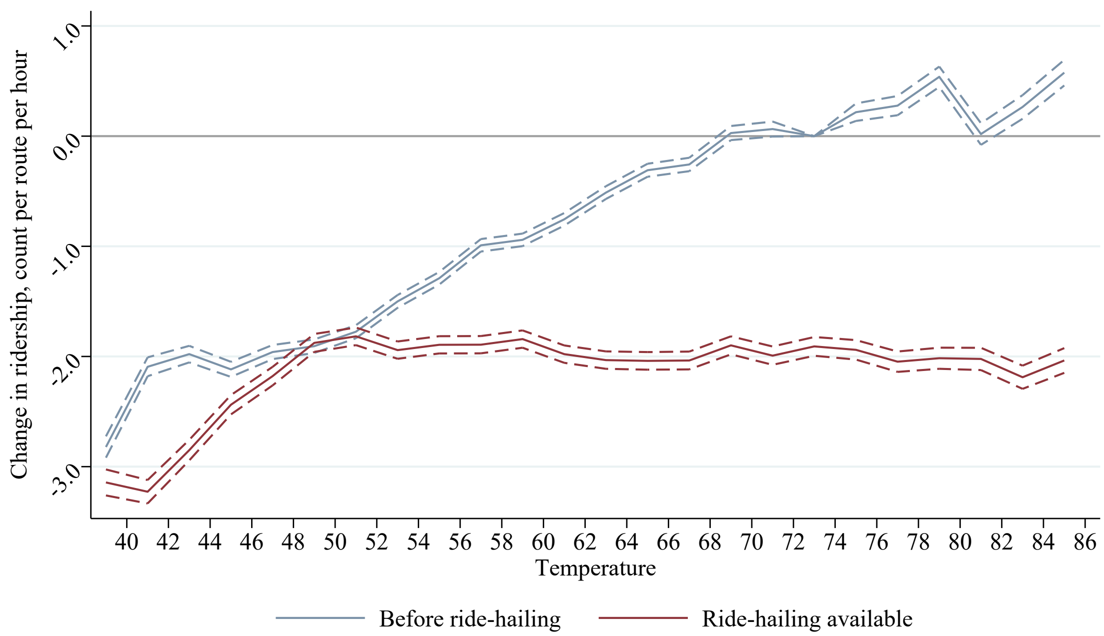

Notes: Two-degree temperature bins, centered on odd values. Adoption period July 2012 to June 2014 omitted. Dashed lines represent 95% confidence intervals.

We estimate results using three time windows: (1) omitting the first two years of ride-hailing, (2) for the entire decade, and (3) breaking the period of ride-hailing into phases, setting two years for adoption. We provide the results from (1) because they are the most conservative; others available on request. We summarize the findings with figures, one for the full data set and three additional for our groups, *Commuter*, *Recreational*, and *Airport* transit. The overall pattern is clear: before ride-hailing is established, extreme temperatures and rain reduce demand. When ride-hailing is available, demand is further

Figure 4. BART Commuters' responses to temperature, before and with ride-hailing availability

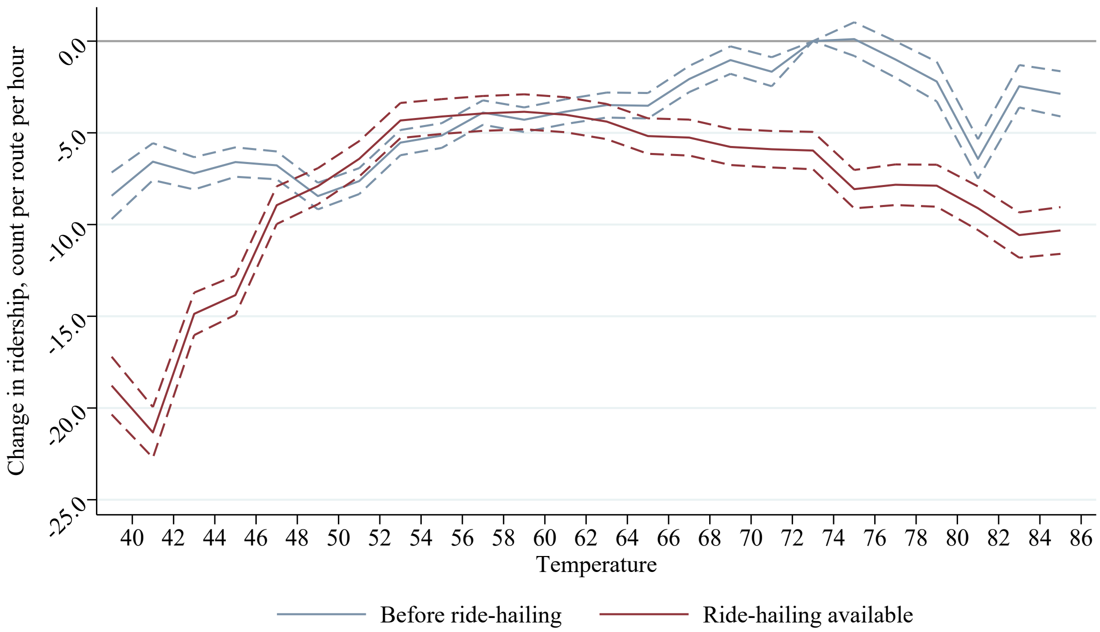

Notes: Two-degree temperature bins, centered on odd values. Adoption period July 2012 to June 2014 omitted. Dashed lines represent 95% confidence intervals.

reduced for most temperatures and for rainy hours, compared to before the availability of ride-hailing. Figure 3 provides the temperature results for all ridership. For the period before ride-hailing, ridership is highest in the range of 68 to 76 F°, which coincides with comfortable ambient temperatures. Whereas, passengers per hour decrease linearly as temperatures drop. The line for "Ride-hailing available" represents the impact. When ride-hailing is present as an alternative mode to BART, ridership falls for all temperatures except for bins 48 to 52 F°. The Uber-effect is around one passenger per route-hour in the ranges of 40 to 46 F° and 52 to 56 F°. Above 56 F° the impact increases to two passengers lost per route-hour. These translate to ride-hailing inducing losses of 0 to 25 percent, depending on the temperature. A rainy hour decreases ridership by 1.6 passengers per route-hour, or 19.7 percent. Once ride-hailing is available, rain does not significantly change the magnitude of lost ridership.

The results by public-transit ridership group demonstrate heterogeneity and confirm our

Figure 5. Recreational responses to temperature,
before and with ride-hailing availability

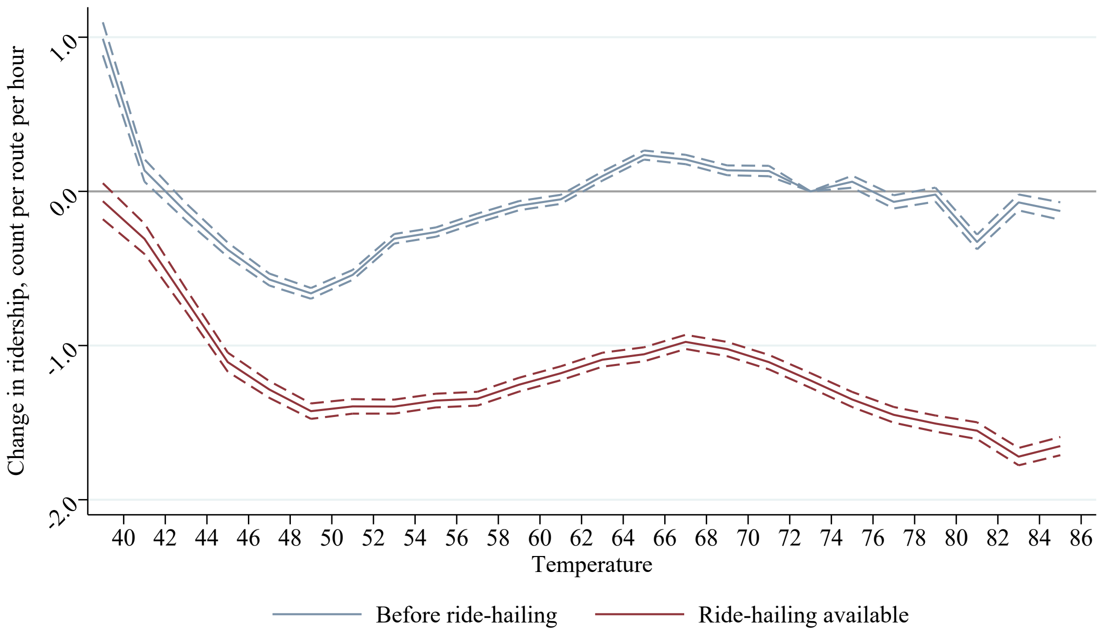

Notes: Two-degree temperature bins, centered on odd values. Adoption period July 2012 to June 2014 omitted. Dashed lines represent 95% confidence intervals.

intuition and hypotheses. Even with availability of ride-hailing, commuters are fairly insensitive to changes for moderate ambient temperatures, between 48 and 64 F°; Figure 4. However for cooler temperatures, the commuter ridership before ride-hailing drops an additional rider per route-hour for every two-degree bin, with 15 fewer passengers per route-hour in the 40-42 F° bin. In percent terms, these changes values are smaller than the overall effect above because commuter route-hours have 10 times more passengers than the pooled average. The impact once ride-hailing is available during extreme temperatures is a loss of about 8.2 people in the coolest weather and about 4.7 passengers during warmer weather. In addition, a rainy hour causes 5 fewer commuters per route-hour, before ride-hailing, and 0.2 fewer commuters when ride-hailing is an available substitute, which is a 0.023 percent loss; however, the result is not statistically significant. Figure 5 provides results for recreational trips: the substitution effect when ride-hailing is present for the temperature is about 1 to 1.5 passengers per route-hour, or 16.4 to 24.6 percent. With ride-hailing availability, rain reduces hourly recreational passengers by 2.1 percent. Figure 6 provides

Figure 6. SFO responses to temperature,
before and with ride-hailing availability

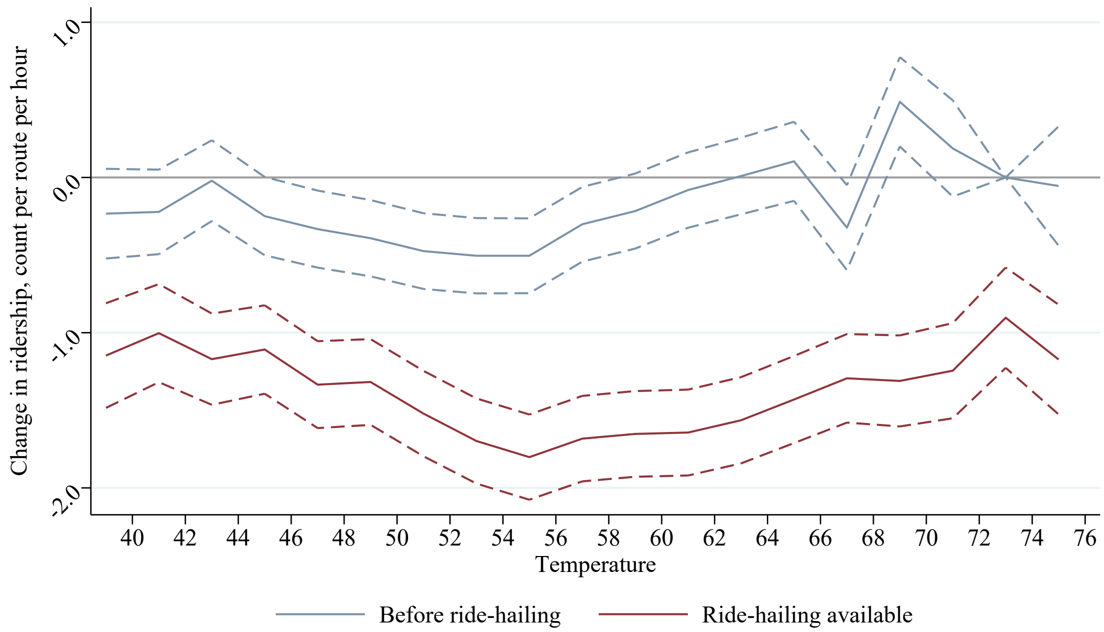

Notes: Two-degree temperature bins, centered on odd values. Adoption period July 2012 to June 2014 omitted. Dashed lines represent 95% confidence intervals.

results for transit to and from the San Francisco International Airport; again, ride-hailing causes losses to ridership. The gap ranges from 1 to 2 passengers per route-hour, or 20.8 to 41.6 percent. During a rainy hour that has ride-hailing availability, airport transit is 2.9 percent lower. According to the weather results, ride-hailing crowds out public transit.

Our second causal factor is highway traffic speeds. Time saving is one of the main reasons for choosing to take ride-hailing over public transit (Feigon and Murphy, 2018; Rayle et al., 2016). In the Bay Area, driving is usually faster than public transit. Before the availability of ride-hailing, we posit that during hours with slower-than-normal vehicles speeds, those who drive may shift to BART, causing BART ridership to increase. Now that ride-hailing is available, we hypothesize that ride-hailing will reduce BART ridership for speeds in the range of free-flowing traffic. However, there are slow speeds at which point ride-hailing no longer prevails in the dimension of travel-time cost. We call this speed the "cross-over" point and expect that at speeds lower than this, the availability of ride-hailing increases public transit as more people opt to avoid traffic congestion.

We approximate speed, in mph, for every hour of each BART route by averaging origin and destination speed-by-direction of travel.[6] We create indicator bins, defined by 5-mph speeds for each route-hour, ih: $\sum_{d=40}^{75} \text{bin}_{d,ih}$. We exclude a before-ride-hailing bin, from 65 up to 70 mph, centered on 67.5 mph. Covariates include hourly weather variables, the proportion of hour day-lit, and indicators for holidays and busy travel days. Time fixed effects for each month-year, day of week, and hour of day control respectively for underlying long run trends, seasonal and weekday transit patterns, and intra-day demand outcomes that are common across all routes.

Figure 7. BART ridership responses to highway speeds

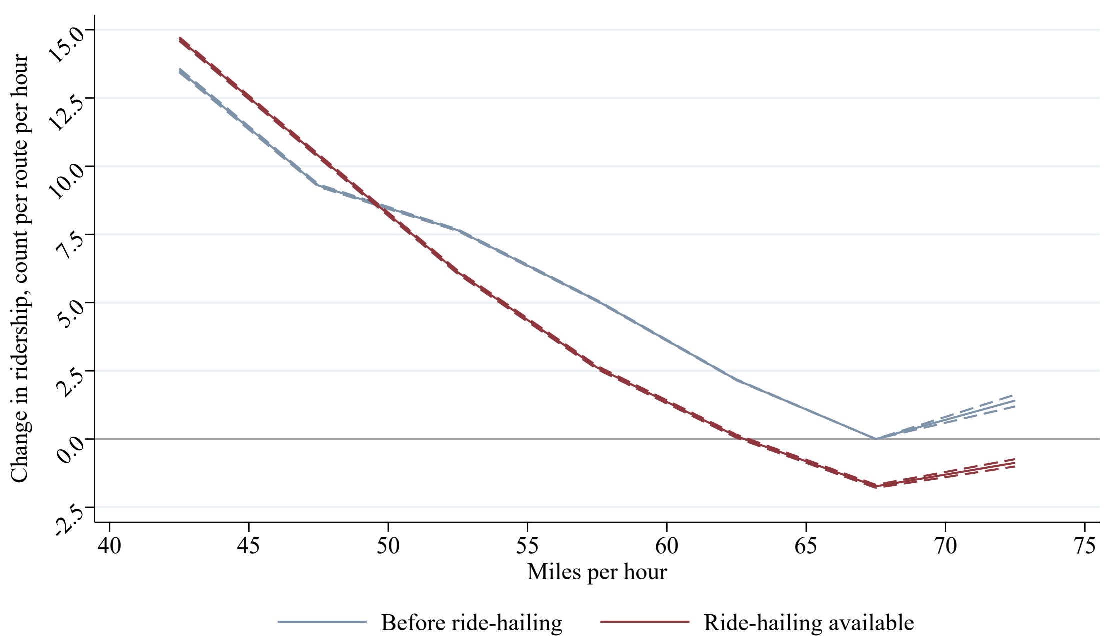

Notes: Each bin is 5 mph: 40-45 mph, 45-50 mph, ..., and so on. We exclude the before-ride-hailing bin 65-70 mph, centered on 67.5 mph.

As shown in Figure 7, before ride-hailing exists, more consumers choose BART as speeds decrease away from free-flowing traffic of 65 to 70 mph. As highway speeds slow to around

[6]For robustness, we also test two other time measures: an approximation of travel time and occurrence of bottlenecks, defined as unexpected traffic jams. These data are poorer in quality and require more assumptions than our speed measures, but the results are similar.

Figure 8. BART ridership responses to highway speeds, by group

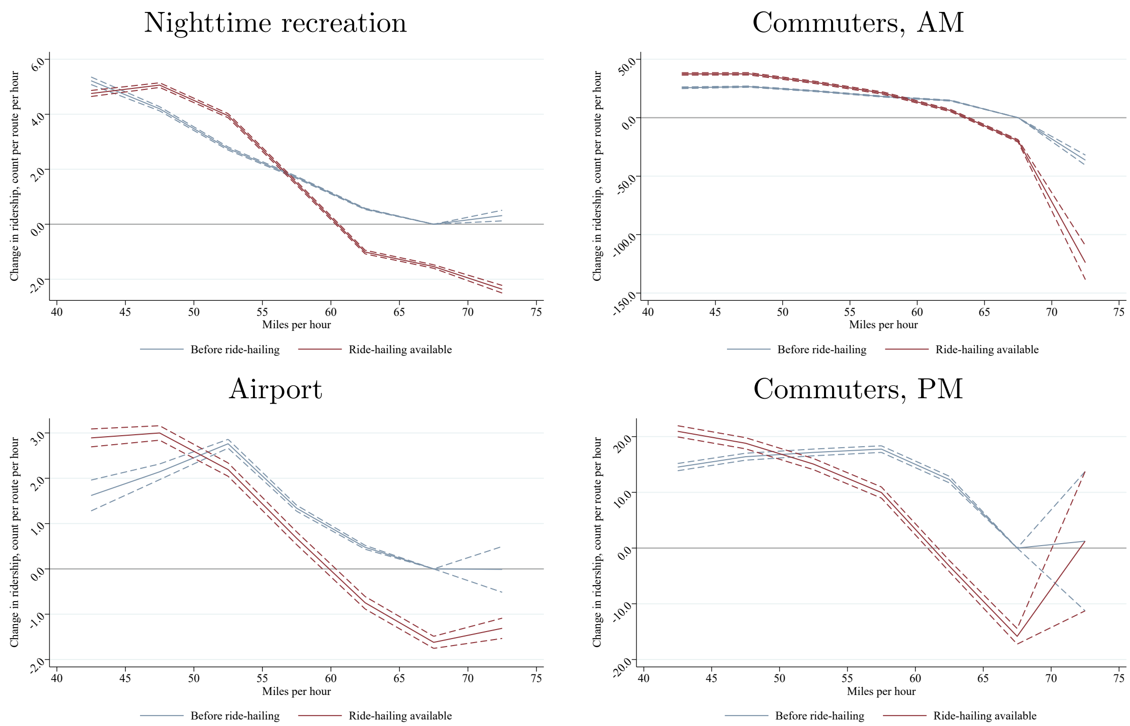

Notes: Available in the appendix, results for daytime recreational, to/from SFO separately, and other.

40 mph, about 13 more people choose public transit per route, per hour. With ride-hailing available, fewer passengers choose BART at high speeds because of reduced time costs, but this wedge narrows as traffic speeds decrease and at speeds below 50 mph, more passengers choose BART. This cross-over point may indicate consumers' change in willingness to pay for time savings — a relatively high speed may indicate that consumers are more elastic to travel-time costs and vice-versa. Or the point might reflect the availability of outside options for distinct travel segments. Cross-over at a lower speed implies less elastic demand with respect to travel time.

Decomposing these results by group provides additional features of relationship between highway speed, public transit and ride-hailing; Figure 8. There is large variation in magnitudes between and within groups: Very little change of ridership occurs for airport transit; the spread over speeds is less than twenty percent of the standard deviation. Whereas there are enormous changes for commuters, particularly in the morning, which ranges from

25 more riders at slow driving speeds to almost 50 fewer at fast speeds — representing 75 percent of the standard deviation. Furthermore, if the cross-over point occurs at a lower speed, the group seems more inelastic to changes in travel time. In other words, airport and evening commuters appear more tolerant of longer travel times, whereas morning commuters and nighttime recreation passengers are more responsive and shift to transit at lower road speeds.

We also analyze the impacts on commuter round-trips, i.e., how ride-hailing affects public transit to work versus home, using difference-in-differences analysis. Once ride-hailing is available, transit riders have a new outside option for either direction of a round trip: they can split the two directions into different modes of travel. Commuters can more easily decouple how they get to work from how they return home. Thus, we hypothesize that public transit in the morning complements evening ride-hailing and vice-versa. We test for this relationship using commuter route-hour subset of data. We collapse these data to average ridership per route for all hours of morning commuting and merge those to the same day's evening commutes. We regress morning ridership and its interaction with the availability of ride-hailing on evening ridership by date and route, per the specification at the beginning of this section, 4.1. We find that once commuters can 'take an Uber' in the evening, about 1 in 100 do so each day. This result is robust to several variations of covariates, including using non-rush-hour ridership to control for a secular trend in ridership. An alternative explanation is that more people use ride-hailing in the morning, with the expectation of using public transit home. However our former interpretation is more consistent with greater inelasticity of morning commuters and anecdotes about having and desiring more flexibility in the evening.

4.2 Long run: comparing trends before ride-hailing and during its availability

In our long-run estimate, we compare daily trends in BART ridership, before ride-hailing companies and then with ride-hailing availability. We estimate the change in daily ridership by route, for each additional day that ride-hailing is available. The long-run results are an average over all years, and the findings provide the cumulative effects of ride-hailing, year over year. Our regressions are fixed-effects panel specifications:

$$\text{bart}_{it} = \begin{cases} \beta_0 d_t + \beta_1 \text{RH}_t + \beta_2 d_t * \text{RH}_t + X'_t * \boldsymbol{\Gamma} + \boldsymbol{\gamma_i} + \boldsymbol{T}_t + \epsilon_{it} \\ \beta_0 d_t + \beta_1 \text{RH}_t + \beta_2 d_t * \text{RH}_t + \beta_3 \text{RH\_m}_t + \beta_4 d_t * \text{RH\_m}_t + X'_t * \boldsymbol{\Gamma} + \boldsymbol{\gamma_i} + \boldsymbol{T}_t + \epsilon_{it} \end{cases}$$

As above, the outcome variable, *bart*, the count of BART passenger-trips for route i , but now aggregated to date t. The top equation represents a simple comparison of before ride-hailing and when ride-hailing is available. d_t is the sequential date, providing the time trend for before ride-hailing is available. The dummy variable RH is defined as before. $\boldsymbol{X_t'}$ is the collection of covariates: daily temperature low and high, a rainy day indicator; holiday indicators for BART scheduling, Federal, and recreational days; and the 25 busiest travel days by year to/from SFO. $\boldsymbol{X_t'}$ also includes monthly time-series for SF-area population, employment, and traffic congestion, and weekly gas prices. The $\boldsymbol{\gamma_i}$ is a vector of route fixed effects. We account for the month-of-year and day-of-week with time fixed effects, $\boldsymbol{T}_t$. β_0 is the linear time trend in *bart* ridership, conditional on other controls, before ride-hailing companies entered the transportation market. The β_1 estimate is a constant intercept without meaningful interpretation. The coefficient, β_2, on the interaction, $d_t * RH_{it}$, is our fundamental estimate as it measures how the ridership changes for each additional day of ride-hailing availability. Again, a positive value implies that ride-hailing complements public transit, whereas, a negative value signals substitution.

The bottom equation splits the availability of ride-hailing into two periods, modeling the trend-lines in Figure 1. The dummy variable, *RH-m* indicates that ride-hailing is widely available and in use in the SF area, i.e., "mature." We use a period of two years for these main results and test robustness of results to varying lengths, in the appendix. Thus, the other result of primary interest is β_4, which provides the change in trend, for each additional day, after ride-hailing reaches maturity compared with the early phase of ride-hailing. The interpretation of the β_4 estimate depends on the sign of β_2. Thus, we perform a linear t-test of $\beta_2 + \beta_4$ to determine the joint effect of ride-hailing availability: If the joint result is greater than 0, the net effect of ride-hailing, in the long run, is complementary to public transit, less than 0 implies that the two modes are substitutes.

Again, we find the entrance of ride-hailing companies reduces ridership. We use ordinary least squares estimation, and we follow the now-standard protocol of clustering standard errors by route to account for potential serial correlation. Table 2 reports the results. Column (1) omits the first two "adoption years" of ride-hailing availability. The effect of

Table 2. Long run: Comparison of daily ridership trends, pre- and during ride-hailing availability

Variables	(1)	(2)	(3)
Date	-0.0011	-0.0228**	0.0053**
	(0.00204)	(0.00181)	(0.00160)
Date*Ride-hailing		-0.0466**	-0.0156**
		(0.00205)	(0.00125)
Date*Ride-hailing*Mature	-0.0447**		-0.0271**
	(0.00195)		(0.00148)
Overall effect of ride-hailing	-0.0447	-0.0466	-0.0427
availability, (*t-test*)	(0.00195)	(0.00205)	(0.00189)
Post-estimation calculations			
Daily Δ in trips, at year's end	-30,849	-32,193	-29,509
%-Δ in trips, at year's end	-10.25	-10.69	-9.803
Annual Δ in passenger-trips	-5,529,943	-5,875,223	-5,385,393
Population	-0.1793**	0.2908**	-0.2988**
	(0.03948)	(0.02383)	(0.02957)
Employment	0.4117**	0.4386**	0.4009**
	(0.01997)	(0.02053)	(0.01930)
ln(Gas price, real $)	12.4344**	6.2072**	12.6172**
	(1.26143)	(1.32123)	(1.22979)
Daily VMT (1000s), 3-day moving avg	0.0023**	0.0020**	0.0020**
	(0.00010)	(0.00009)	(0.00009)
Temperature F°, low	0.4643**	0.6941**	0.5477**
	(0.03673)	(0.04145)	(0.03850)
Temperature F°, high	0.0872**	0.1338**	0.1173**
	(0.01743)	(0.01666)	(0.01674)
Rainy day	-3.1426**	-3.3602**	-3.0873**
	(0.18208)	(0.18751)	(0.17982)
Observations	5,317,808	6,622,374	6,622,374
R-squared	0.161	0.161	0.161
Route FEs	1,892	1,892	1,892

Note: Outcome is daily BART ridership per route; the treatment is an additional day of ride-hailing availability. Column (1) omits the first two years of ride-hailing availability, whereas column (2) keeps this period. Our social cost estimates are based off of column (1). Column (3) estimates the impact of ride-hailing availability separately for the first two years relative to when ride-hailing is widely available. Robust standard errors clustered by route; ** p<0.01, * p<0.05.

ride-hailing is 5.5 million trips lost system-wide annually or 5.4 percent of average annual ridership. We use this value for our calculations of long-run social costs in the next section. Column (2) estimates the simple before/after-trend model and does not distinguish between the adoption and mature period. The overall effect of ride-hailing availability is similar. The *Date* trend in ridership, prior to ride-hailing availability, is negative and significant, -0.02 trips lost per route-day — a decrease of 43 BART trips per additional day over all routes. Our results concur with Manville et al. (2018a), which shows that population growth in Southern California outpaces the proportion who use transit. In column (3), we estimate the impact of ride-hailing during the first two-years versus the mature period. The impact of ride-hailing is 30 fewer trips per additional day. Once ride-hailing is widely available, the loss of BART ridership further decreases by 51 trips per additional day. We use a t-test to determine the *Overall effect of ride-hailing availability*, -0.0427 trips per route per additional day. The results for the substitution effect are similar across all specifications: overall, ride-hailing is associated with a *year-over-year loss* of 9.8 to 10.7 percent.[7]

Table 2 also reports the estimates for the covariates; the signs conform to our hypothesized effects. Ridership is correlated with population increase, except for the in (3) where the *Date*-trend switches sign too. As employment rises in the region, transit demand increases. When gas price rises, the cost of driving increases and consumers substitute away from automobiles, increasing BART ridership. BART ridership and vehicle miles traveled (VMT), which is a measure of automobile activity, increase together in response to regional demand. We find more BART trips in warmer weather and fewer on rainy days.

In Table 3, we decompose the results of column (3), by ridership group, to obtain the respective effects for airport travel, commuters, and recreational transit. The results for before and during ride-hailing availability follow the pattern from the complete data set. However, the magnitudes vary by group, and mostly follow our hypotheses. We first assess routes to and from SFO Airport, with public transit to and from SFO falling by 10.5 percent, year-over-year, about the same as the overall average. Next, we estimate effect

[7] To determine percent change, the numerator is the estimated marginal effect per additional day on a route's ridership times 365 days in a year. The denominator is the average ridership per route, before ride-hailing, multiplied by the number of routes. Dividing these gives the average annual percent change. We test for sensitivity by making several modifications to the specification: we include lag and lead variables (Table A-4), vary the length of the adoption period (Table A-3, and add several other month-varying independent variables).

Table 3. Long run: Comparison of daily ridership trends, pre- and during ride-hailing availability, by group

Variables	SFO	Comm, AM	Comm, PM	Rec, day	Rec, night	Other
Date	0.0071**	-0.0564**	-0.0391**	0.0010	0.0004	0.0012
	(0.00246)	(0.00514)	(0.00378)	(0.00066)	(0.00041)	(0.00065)
Date*Ride-hailing	-0.0011	-0.0690**	-0.0405**	-0.0025**	-0.0054**	-0.0083**
	(0.00219)	(0.00592)	(0.00465)	(0.00071)	(0.00042)	(0.00050)
Date*Ride-hailing*Mature	-0.0113**	0.0110**	-0.0003	-0.0092**	-0.0051**	-0.0050**
	(0.00230)	(0.00384)	(0.00331)	(0.00063)	(0.00032)	(0.00044)
Overall effect of ride-hailing	-0.0125**	-0.0580**	-0.0408**	-0.0117**	-0.0106**	-0.0133**
availability, (*t-test*)	(0.00241)	(0.00574)	(0.00445)	(0.000537)	(0.000545)	(0.000520)
Post-estimation calculations						
Daily Δ in trips, at year's end	-382.7	-6,653	-4,679	-7,683	-6,961	-8,768
%-Δ in trips, at year's end	-10.53	-9.405	-6.710	-8.666	-17.35	-11.72
Annual Δ in passenger-trips	-69,834	1,214,173	-853,918	1,402,148	1,270,383	1,600,160
Population	-0.1376**	1.1248**	0.7837**	-0.0253	0.0210*	0.0036
	(0.04396)	(0.12345)	(0.09372)	(0.01552)	(0.00828)	(0.01265)
Employment	0.0912**	0.4534**	0.2927**	0.0488**	0.0456**	0.0730**
	(0.01674)	(0.03522)	(0.02745)	(0.00424)	(0.00326)	(0.00354)
ln(Gas price, real $)	1.2294	8.1642**	3.1359	1.0198**	-0.7283**	-0.4522
	(0.78856)	(2.25736)	(2.19619)	(0.33963)	(0.23874)	(0.31438)
Daily VMT in 1000s	0.0011	-0.0190**	-0.0129**	-0.0003*	-0.0006**	-0.0009**
	(0.00059)	(0.00175)	(0.00133)	(0.00015)	(0.00010)	(0.00014)
Temperature F°, low	-0.0631	-3.5352**	-2.1822**	-1.3763**	0.1536**	-0.2130**
	(0.03175)	(0.19510)	(0.10527)	(0.11405)	(0.02415)	(0.02012)
Temperature F°, high	0.0934**	4.5130**	2.3685**	1.0460**	-0.0211	0.4551**
	(0.03506)	(0.24568)	(0.11187)	(0.07828)	(0.02145)	(0.02431)
Rainy day	-0.9293*	1.3522**	-0.9695*	-0.8958**	-1.2035**	-0.6195**
	(0.42727)	(0.31772)	(0.38832)	(0.14166)	(0.06819)	(0.05985)
Observations	303,760	800,867	801,488	6,486,612	3,697,009	6,312,300
R-squared	0.266	0.162	0.144	0.023	0.126	0.188
Route FEs	84	314	314	1,806	1,806	1,808

Note: Outcome is daily BART ridership per route; the treatment is an additional day of ride-hailing availability. The estimates parallel column (3) of Table 2, which assumes a 2-year adoption period. We estimate each group's results separately. Robust standard errors clustered by route; ** $p<0.01$, * $p<0.05$.

on commuters. We hypothesize that ride-hailing has less influence on commuters because demand is inelastic with respect to getting to work each day. Indeed, both morning and evening rush-hour routes have smaller losses than average, 9.4 percent and 6.7 percent annually. We expect recreational transit to have the most flexibility and exhibit larger effects. However, the daytime reduction is 8.7 percent, indicating this group may include a less-elastic segment of passengers. The nightlife results bear out our intuition — ridership has been dropping precipitously, by an average of 17.4 percent per additional year. The 'Other' category is too broad for a practical explanation. Adding up all groups' annual change in trips gives a total loss of 6,410,530 trips, which squares with the overall estimates above.

A concern with these long-run results is the inherent challenge of before-and-after comparisons: there is always the risk that an unobservable or unaccounted trend could be influencing the outcome simultaneously with the treatment of interest. For example, if an omitted factor is driving down BART ridership during the same period that ride-hailing availability is increasing, we may have incorrectly attributed the decline in ridership to ride-hailing. Fortunately, several aspects alleviate this concern. First, we examine demographic trends. Over our study period, population and employment opportunities have increased in the Bay Area. Air traffic to and from SFO and OAK has also trended upward throughout the decade (`data.sfgov.org`). In addition, Bay Area residents own fewer cars in the latter half of the decade, which should increase demand for both public transit and ride-hailing services. These prevailing factors would likely drive up demand for BART, suggesting that we may be underestimating the long-run effects of ride-hailing. Second, BART provides ridership forecasts and sets monthly ridership goals (`data.bart.gov/dataset/customer-ridership`). Actual ridership has consistently fallen below these projections in recent years, indicating the presence of an external and unaccounted-for shock, such as the rise of ride-hailing. We also review BART's schedule changes over the decade and find little evidence of reduced service. Third, the results by group are distinct and intuitive. For an unobserved trend to explain our findings, it would have to affect BART market segments in a very specific and unlikely manner. Finally, as discussed in the previous section, we provide additional evidence that ride-hailing substitutes for public transit, pointing to plausible causal mechanisms for the long-run decline in ridership.

5 Costs of transit riders switching to ride-hailing

Evidence shows that *Uber*X and *Lyft* reduce BART ridership in the San Francisco Bay Area. As consumers are choosing ride-hailing over high-occupancy public transit, we now consider two categories of impacts: a decrease in BART's fare revenue and several negative externalities from the additional vehicle miles. The externalities include social costs from pollution, increased travel times, and auto-related collisions. We first determine average values for several parameters: BART fares, the mileage of BART trips, the occupancy of riders in ride-hailing vehicles, and the social costs per mile. We then convert change in BART ridership to dollar values. We calculate separate totals for the long-run daily trend and the short-run hourly effects. We then discuss the connection between these.

We use 2018 data as our base year and covert dollar values from other years. The average 2018 BART fare was $4.[8] To calculate the social costs, we need an estimate of additional vehicle miles traveled (VMT), which requires assumptions about their distance of travel and how many passengers are in each ride-hailing vehicle. The average BART trip distance, 14.8 miles, is a good proxy for distance per trip.[9] To provide a range of costs, we assume either 10 or 20 miles per trip. These distance measures are conservative: by ignoring the travel between the origin and destination and entry and exit BART stations, we underestimate the total trip length. The occupancy rate for San Francisco, from Castiglione et al. (2016), is 1.66 passengers per vehicle.

We use two approaches to determine the values of social costs per VMT: (1) Parry et al. (2007) provides values for four impacts, carbon dioxide (CO_2) pollution, local pollution, congestion, and traffic collisions. (2) We collect information that is aligned with our study period, and specific to the three countires on the Bay Area, to determine alternative values for each impact. Our methods are as follows. For pollution cost per VMT, we use emission cost per metric ton and pollutant grams per mile.[10] To calculate the average Bay Area congestion cost per VMT, for each year from 2013 to 2018, we divide reported annual congestion cost in 2018$ by total annual VMT in the three counties BART operates. We then average the congestion cost per VMT across the six years ride-hailing available in the

[8] See 2019 BART Fact Sheet, accessed October 2024.
[9] *Ibid.*
[10] See Benefit-Cost Analysis Guidance and Motor Vehicle Emission Simulator, accessed October 2024.

Bay Area. We calculate the collision cost as the rate of injury per VMT weighted by the probability and cost of injury type.[11] The average costs for congestion and collisions further the underestimation of the impact of the marginal vehicle because both these impacts rise in marginal cost per vehicle mile added. In sum, the average social cost per additional vehicle miles traveled is $0.150 to $0.574; all values in Table 4. Ours are higher than in Parry et al. (2007) because they reflect contemporary costs and pertain to the populous and congested region in which BART operates.

From Column 1 of Table 2, our results show that ride-hailing induces an annual loss of 5,529,943 trips. Thus, the estimated lost BART revenue resulting from ride-hailing is about $22.1 million year-over-year. By 2018, we estimate a loss of $132.7 million. For comparison, we project a counterfactual 2018 revenue of $536.1 to $581.8 million and BART's operating budget for the same year was $920.6 million.[12] The lost revenue is a private transfer, from BART to the ride-hailing companies and drivers. Moreover, we acknowledge that private welfare increases because these consumers reveal their preference for ride-hailing over BART. In fact, Cohen et al. (2016) estimate consumer surplus using *Uber*X surge-price data for individual trips; for the four cities with largest markets in 2015; they find a valuation of $2.9 billion.

Table 4. Social costs of additional ride-hailing trips: Average year-over-year trend

	$/VMT		**Total Costs**		
Impact	Parry et al. (2007)	Authors' Calculations	Low	Moderate	High
CO_2 Pollution	$0.004	$0.017	$291,175	$574,657	$1,149,314
Local Pollution	$0.029	$0.005	$1,941,167	$165,980	$331,960
Congestion	$0.05	$0.433	$4,852,919	$14,418,335	$28,836,670
Collisions	$0.044	$0.109	$2,911,751	$3,632,317	$7,264,634
Sum	$0.150	$0.574	$9,997,012	$18,791,289	$37,582,578

Note: Values in 2018$. Derivations of parameters, methods, and ranges described in main text.

We monetize the year-over-year social costs using estimates of the costs per VMT, a standard benefits-transfer approach. The change in BART ridership combined with ride-hailing occupancy and distance per trip provides a range of 33,312,910 to 66,625,819 additional

[11] See 2015 Motor Vehicle Crashes: Overview and Tiger Benefit-Cost Analysis (BCA) Resource Guide, accessed October 2024.
[12] See BART Fact Sheet, accessed October 2024.

VMT induced by the availability of ride-hailing. Table 4 provides the year-over-year social costs. 'Low' uses the low VMT estimate and costs from Parry et al. (2007); 'Moderate' uses the low VMT estimate and our calculated costs; and 'High' column uses the high VMT estimate and our own costs. We find that total negative externality costs increase year over year by $10 to $37.6 million annually.

We now calculate costs based on our hourly difference-in-differences estimates. We focus on the two factors with the largest impacts on ridership: temperature and congestion.[13] We first find the weighted-average change in ridership over all the binned estimates, based on Figures 3 and 7. Then we multiply by the average number of system-wide operational hours per day, inferring an average daily loss of 30,864 riders due to congestion and 66,586 riders due to temperature.

Using these values, BART is potentially losing up to $248,000 per day in fare revenue due to a preference for ride-hailing over BART when temperatures are uncomfortable *or* ride-hailing is faster. We translate ridership into additional VMT using these ridership changes, the occupancy rate of 1.66, and the same range of 10 to 20 additional VMT per ride-hailing trip. Table 5 shows the social costs from our short-run analysis. As above, 'Moderate' uses the low VMT and 'High' uses the high VMT estimate; both apply our calculations of cost/VMT. The daily sum of impacts is — $103,000 to $207,000 from congestion shocks and $230,000 to $461,000 from temperature shocks.

Table 5. Social costs of additional ride-hailing trips: Average daily values due to hourly congestion and temperature fluctuations

Impact	**Authors' $/VMT**	**Percent /Impact**	**Congestion Costs** Moderate	High	**Temperature Costs** Moderate	High
CO_2 Pollution	$0.017	2.96%	$3,202	$6,405	$6,909	$13,818
Local Pollution	$0.005	0.87%	$925	$1,850	$1,996	$3,991
Congestion	$0.443	77.18%	$82,446	$164,892	$177,866	$355,732
Collisions	$0.109	18.99%	$20,242	$40,485	$43,670	$87,341
Sum	$0.574	100%	$106,816	$213,632	$230,441	$460,882

Note: Values in 2018$. Derivations of parameters, methods, and ranges described in main text.

Combining the effects of temperature and congestion on ridership is challenging because the

[13] We disregard the short-run results from rain and AM-PM commuters as these are small in both magnitude and duration.

two factors interact. We choose to use temperature as the primary variable to estimate total average daily costs, as it more effectively captures the combined impact of both temperature and congestion. Congestion costs, on the other hand, are less sensitive to temperature fluctuations because congestion tends to be highest during commute hours when demand is relatively inelastic. Bad weather, such as rain, often exacerbates congestion, as more people opt to drive instead of taking public transit, and driving conditions worsen. As a conservative estimate of short-run impacts, we use the lower bound of average daily temperature costs and multiply it by 365 days to calculate an annualized perpetual cost of $84 million.

The daily fluctuations in ridership — and respective costs — are four times larger than the year-over-year losses. The changes in hourly and daily ridership occur within the downward, long-run annual trend. Therefore we cannot simply add the long- and short-run costs together. We can surmise that congestion and temperature displace numerous transit riders temporarily, yet according to our estimates, about 80 percent of these passengers return to BART. We add the year-over-year 'moderate' value of $19 million to account for the long-run trend. Given our calculations rely on many assumptions, we refrain from an exact figure of total social costs. We also note that these measures are for a single metropolitan area and one form of alternative transit. Our measure is on the order of $103 million dollars in 2013 growing to $197 million in 2018, for a total impact over this period of nearly one billion dollars.

6 Discussion

We are unaware of any other studies that discern whether ride-hailing is a complement or substitute for public transit at the metropolitan level, and across both hourly and daily timescales. We are the first to differentiate outcomes for distinct types of riders — commuters, airport travelers, and recreational passengers. Previous studies have reached conflicting conclusions regarding complementarity or substitution, which can be attributed to varying levels of data aggregation, differing identification strategies, and diverse time spans. The fundamental issue is that different municipalities respond to ride-hailing in distinct ways. Additionally, we examine multiple factors that may lead riders to switch between public transit and ride-hailing. Public transit authorities and municipalities require these finer details to inform decision-making, which is why we provide a template for replicating our approach in other regions.

We present methods and findings for two timescales. In the short run, we capture hourly fluctuations, which allow us to identify the mechanisms driving ridership declines due to ride-hailing. We focus on two key factors that may push transit riders toward ride-hailing: increased protection from weather and reduced travel time. In contrast, our long-run analysis compares time-series trends before and after the introduction of ride-hailing, shedding light on how hourly ride-hailing impacts accumulate into annual ridership changes. This provides a broader view of the factors influencing ridership, including how the availability of ride-hailing affects daily transit patterns.

We apply our methods using data from the San Francisco Bay Area. Our findings reveal hourly substitution: BART users temporarily deviate from public transit in response to worsening weather, decreased daylight, or faster travel speeds, with significant declines in ridership when ride-hailing is available. While many riders return to public transit once these short-term conditions improve, some appear to habituate to ride-hailing, likely contributing to the long-term decline in ridership. This leads to a permanent loss of 5.5 million BART trips, on average, for each additional year of ride-hailing availability. The decline persists despite population growth and an increase in employment opportunities in the city center. Naturally, unobservable trends across time and space could affect our results. For instance, surge pricing practices by ride-hailing companies, which are not fully captured by route and time fixed effects, might introduce bias. However, the patterns we observe across different rider groups bolster our confidence, with small elasticities for commuters and high elasticities for recreational and airport travelers.

Our results are also supported by related facts and literature. BART ridership has been significantly lower than forecasts by the transit authority, indicating that existing models may not account for certain factors. The 2018 SF Mobility Report suggests that while San Francisco residents are driving their own cars less, overall driving and congestion have increased, with vehicle speeds declining by 23 percent over the past decade.Several recent studies find that ride-hailing contributes to an increase in vehicle miles traveled (VMT), even after accounting for other growth factors (Erhardt et al., 2019; Henao and Marshall, 2019; Tirachini and Gomez-Lobo, 2020), with Tirachini (2019) offering a comprehensive review. This increase implies that ride-hailing induces additional trips by current drivers or causes transit riders to shift to ride-hailing. In particular, Rayle et al. (2016) surveyed ride-hailing passengers in central San Francisco and found that over 30 percent chose ride-hailing over their next-best option of public transit. Similarly, Schaller (2021) analyzed ride-hailing consumers in four major U.S. cities, including San Francisco, concluding that VMT increases due to empty ride-hailing vehicles searching for passengers and because many consumers shift from public transit to ride-hailing.

7 Conclusions

Uber and *Lyft* have been incredibly successful in developing app-based technologies that match passengers with drivers for a taxi-like service (Castiglione et al., 2016; Cramer and Krueger, 2016). In addition to improving informational efficiencies, ride-hailing offers other private benefits. For drivers, ride-sharing gigs provide flexible compensation without the commitments of traditional employment or taxi leasing (Angrist et al., 2021; Chen et al., 2020, 2019). For passengers, ride-hailing increases access to areas not served by public transit, arrives on-demand, and often provides greater comfort and safety than public transit. Research shows that ride-hailing improves personal mobility and increases consumer surplus (Christensen and Osman, 2021; Cohen et al., 2016). As these technologies continue to evolve, transport as a service is expected to occupy a larger share of the market, influencing demand for public transport and altering its provision and operations (Hörcher and Tirachini, 2021; Webb, 2019).

Concurrently, public transit ridership across the U.S. has seen significant declines. Using data from San Francisco, we estimate that ridership in 2019 would have been 11.2 to 21 percent higher in the absence of ride-hailing.[14] Various factors contribute to this decline (Graehler et al., 2019; Grisby et al., 2018; Manville et al., 2018b).[15] Is ride-hailing displacing public transit? Initial studies yielded mixed results, but growing evidence supports our findings that ride-hailing contributes to the decline in public transit use (Diao et al., 2021; Erhardt et al., 2021; Hall et al., 2018; Nelson and Sadowsky, 2018; Tirachini, 2019). This decline has several consequences. First, a reduction in fare revenue hinders the ability of

[14] Values are from an unconditional linear projection of the data in Figure 1, from 2011 to 2014.

[15] A 2018 article in *The Economist* attributes the decline to "remote working, Uber, cheap car loans, and the internet." Public transport is in decline in many wealthy cities, accessed October 2024.

transit agencies to maintain infrastructure and make necessary capital investments. We estimate that BART fare revenue is 10.3 to 17.4 percent lower than the counterfactual. Second, more traffic results from increased ride-hailing, as fewer people use mass transit. Although pooled ride-hailing options exist, Tirachini and Gomez-Lobo (2020) find that the average occupancy rate for ride-hailing vehicles does not offset the rise in vehicle miles traveled (VMT), contributing to increased congestion.

More traffic and congestion lead to further societal costs, such as lost personal time for all drivers, as well as amplifying other externalities. Using San Francisco-specific parameters, we estimate travel-time costs at \$0.443 per VMT, accounting for 77 percent of social impacts. Modeling by Erhardt et al. (2019) suggests that ride-hailing has tripled congestion delays in the metro area compared to a counterfactual without these services. Anderson (2014) demonstrated that shutting down Los Angeles's metro system resulted in a sharp increase in congestion costs, and conversely, Gu et al. (2021) found that new subway lines reduced congestion in China. We adjust their valuation to U.S. conditions, estimating a traffic speed decline of 0.25 to 0.5 percent, which translates into annual external costs of \$9.5 to 35.3 million — remarkably similar to our estimates. Furthermore, increased traffic raises collision rates, leading to property damage, injuries, and fatalities. Barrios et al. (2020) find that increased ride-hailing activity correlates with a rise in quarterly fatalities in U.S. cities. They argue that ride-hailing contributes to more vehicles roving for passengers and more driving on unfamiliar roads, increasing the risk of collisions. Valuing the cost of these fatalities at \$15 to \$40 per capita, we estimate the annual collision cost for the San Francisco metro area to be \$49 to \$133 million. Thus, the Bay Area's 2018 collision costs likely total around \$40 million — a reasonable figure when combining property damage, injuries, and fatalities. Although smaller, environmental costs also arise from increased ride-hailing. At less than 4 percent of the total social impacts, the pollution cost, including local emissions and CO_2, is \$0.22 per VMT.

These findings, combined with other research, reinforce a critical takeaway: a robust public mass transit system is essential for mitigating the social costs of driving. Municipalities and transit agencies need specific, detailed data to create policies that address the crowding out of transit by ride-hailing. While our results may lack external validity for other regions, as public transit systems differ in their structure and demographics, we provide a methodological tool that can be adapted to local conditions. Public agencies should improve data collection to capture detailed measurements of ride-hailing's impacts and respond with appropriate policies.

How should public transit adapt as ride-hailing continues to transform the transportation industry? Unlike private services, public transit agencies focus on maximizing societal welfare and ensuring distributional equity under budget constraints. As ride-hailing places increasing pressure on public transit budgets, it may be more strategic for agencies to concentrate on the segments where they perform best. However, focusing resources on high-performing routes might require cutting services in areas with low-income or transit-dependent populations, raising equity concerns.

Government policies can play a vital role in supporting public transit authorities and regulating ride-hailing companies to mitigate social costs. Governments already subsidize public transit, and such subsidies have proven net beneficial (Parry and Small, 2009). Subsidies can also target routes that divert drivers from congested roads. Public-private partnerships have been explored as a way to foster complementarity, where governments subsidize ride-hailing trips to and from transit hubs.[16] In terms of regulation, an ideal solution would be a vehicle miles tax or congestion pricing, which would charge ride-hailing companies according to their marginal social impact. Based on our estimates, the appropriate Pigouvian tax would be $0.57 per mile. However, Dudley et al. (2017) highlight the challenges of enforcing congestion pricing, as ride-hailing companies have successfully avoided such policies in cities like London. Recent work by Almagro et al. (2024) uses Chicago to model optimal policies, assessing the welfare consequences of road pricing and public transit schedules and fares. An alternative, as shown by Hanna et al. (2017), is the use of high-occupancy vehicle (HOV) lanes to mitigate traffic. As a positive incentive, pooled ride-hailing vehicles often can use these lanes or be exempted from congestion taxes. Meanwhile, California has mandated that 90 percent of ride-hailing fleets must be electric vehicles (EVs) by 2030.[17] While this addresses environmental concerns, it overlooks the largest social costs of ride-hailing — congestion and collisions, which account for 96.2 percent of the total social costs. As ride-hailing causes more people to shift from high-occupancy modes, reducing vehicle miles traveled is critical to mitigating these rising costs.

[16] See Transit and TNC Partnerships, accessed October 2024.
[17] See California regulator adopts EV mandate for Uber, Lyft ride-hail fleets, accessed October 2024.

References

Almagro, M., Barbieri, F., Castillo, J. C., Hickok, N. G., and Salz, T. (2024). Optimal urban transportation policy: Evidence from Chicago. *National Bureau of Economic Research*. No. w32185; found at https://www.nber.org/papers/w32185.

Anderson, M. L. (2014). Subways, strikes, and slowdowns: The impacts of public transit on traffic congestion. *American Economic Review*, 104(9):2763–96.

Angrist, J. D., Caldwell, S., and Hall, J. V. (2021). Uber vs. taxi: A driver's eye view. *American Economic Journal: Applied Economics*, 13(3):272–308.

Babar, Y. and Burtch, G. (2017). Examining the impact of ridehailing services on public transit use. *Available at SSRN 3042805*. Forthcoming in *Information Systems Research*.

Barrios, J. M., Hochberg, Y., and Yi, H. (2020). The cost of convenience: Ridehailing and traffic fatalities. *National Bureau of Economic Research*. No. w26783; found at https://www.nber.org/papers/w26783.

Berger, T., Chen, C., and Frey, C. B. (2018). Drivers of disruption? Estimating the Uber effect. *European Economic Review*, 110:197–210.

Castiglione, J., Chang, T., Cooper, D., Hobson, J., Logan, W., Young, E., Charlton, B., Wilson, C., Mislove, A., Chen, L., et al. (2016). TNCs today: a profile of San Francisco transportation network company activity. *San Francisco County Transportation Authority*. Details at https://www.sfcta.org/projects/tncs-today, accessed 06/20/2021.

Chen, K.-M., Ding, C., List, J. A., and Mogstad, M. (2020). Reservation wages and workers' valuation of job flexibility: Evidence from a natural field experiment. *National Bureau of Economic Research*. No. w27807; found at https://www.nber.org/papers/w27807.

Chen, M. K., Rossi, P. E., Chevalier, J. A., and Oehlsen, E. (2019). The value of flexible work: Evidence from Uber drivers. *Journal of Political Economy*, 127(6):2735–2794.

Christensen, P. and Osman, A. (2021). The demand for mobility: Evidence from an experiment with Uber riders. *IZA Discussion Paper*.

Clewlow, R. R. (2019). Disruptive transportation: The adoption, utilization, and impacts of ride hailing in the United States. *Transfers Magazine*, pages 7–14.

Cohen, P., Hahn, R., Hall, J., Levitt, S., and Metcalfe, R. (2016). Using big data to estimate consumer surplus: The case of Uber. *National Bureau of Economic Research*. No. w22627; found at https://www.nber.org/papers/w22627.

Coogan, M., Spitz, G., Adler, T., McGuckin, N., Kuzmyak, R., and Karash, K. (2018). Understanding changes in demographics, preferences, and markets for public transportation. Technical Report 201, Transportation Research Board.

Cramer, J. and Krueger, A. B. (2016). Disruptive change in the taxi business: The case of Uber. *American Economic Review*, 106(5):177–82.

Diao, M., Kong, H., and Zhao, J. (2021). Impacts of transportation network companies on urban mobility. *Nature Sustainability*, 4(6):494–500.

Dudley, G., Banister, D., and Schwanen, T. (2017). The rise of Uber and regulating the disruptive innovator. *The Political Quarterly*, 88(3):492–499.

Erhardt, G. D., Mucci, R. A., Cooper, D., Sana, B., Chen, M., and Castiglione, J. (2021). Do transportation network companies increase or decrease transit ridership? Empirical evidence from San Francisco. *Transportation*, pages 1–30.

Erhardt, G. D., Roy, S., Cooper, D., Sana, B., Chen, M., and Castiglione, J. (2019). Do transportation network companies decrease or increase congestion? *Science Advances*, 5(5):eaau2670.

Feigon, S. and Murphy, C. (2016). Shared mobility and the transformation of public transit. Available at https://www.apta.com/wp-content/uploads/Resources/resources/reportsandpublications/Documents/APTA-Shared-Mobility.pdf, accessed 04/26/2020. Project J-11, Task 21.

Feigon, S. and Murphy, C. (2018). Broadening understanding of the interplay among public transit, shared mobility, and personal automobiles. Technical Report 195, Transportation Research Board.

Graehler, M., Mucci, R. A., and Erhardt, G. D. (2019). Understanding the recent transit ridership decline in major US cities: Service cuts or emerging modes. In *Transportation Research Board 98th Annual Meeting, Washington, DC, January*.

Grahn, R., Qian, S., Matthews, H. S., and Hendrickson, C. (2020). Are travelers substituting between transportation network companies (TNC) and public buses? a case study in Pittsburgh. *Transportation*, pages 1–29.

Grisby, D., Dickens, M., and Hughes-Cromwick, M. (2018). Understanding recent ridership changes: Trends and adaptations. *American Public Transportation Association.*

Gu, Y., Jiang, C., Zhang, J., and Zou, B. (2021). Subways and road congestion. *American Economic Journal: Applied Economics*, 13(2):83–115.

Hall, J. D., Palsson, C., and Price, J. (2018). Is Uber a substitute or complement for public transit? *Journal of Urban Economics*, 108:36–50.

Hanna, R., Kreindler, G., and Olken, B. A. (2017). Citywide effects of high-occupancy vehicle restrictions: Evidence from "three-in-one" in Jakarta. *Science*, 357(6346):89–93.

Henao, A. and Marshall, W. E. (2019). The impact of ride-hailing on vehicle miles traveled. *Transportation*, 46(6):2173–2194.

Hörcher, D. and Tirachini, A. (2021). A review of public transport economics. *Economics of Transportation*, 25:100196.

Manville, M., Taylor, B. D., and Blumenberg, E. (2018a). Falling transit ridership: California and Southern California. Available at https://scag.ca.gov/sites/main/files/file-attachments/its_scag_transit_ridership.pdf, accessed 04/26/2020.

Manville, M., Taylor, B. D., and Blumenberg, E. (2018b). Transit in the 2000s: Where does it stand and where is it headed? *Journal of Public Transportation*, 21(1):11.

McFadden, D., Talvitie, A., Cosslett, S., Hasan, I., Johnson, M., Reid, F., and Train, K. (1977). Demand model estimation and validation. *Urban Travel Demand Forecasting Project, Phase 1*.

Nelson, E. and Sadowsky, N. (2018). Estimating the impact of ride-hailing app company entry on public transportation use in major U.S. urban areas. *The BE Journal of Economic Analysis & Policy*, 19(1).

Parry, I. W. and Small, K. A. (2009). Should urban transit subsidies be reduced? *American Economic Review*, 99(3):700–724.

Parry, I. W., Walls, M., and Harrington, W. (2007). Automobile externalities and policies. *Journal of Economic Literature*, 45(2):373–399.

Polzin, S. E. (2016). Implications to public transportation of automated or connected vehicles. *National Center for Transit Research, University of Florida*.

Rayle, L., Dai, D., Chan, N., Cervero, R., and Shaheen, S. (2016). Just a better taxi? A survey-based comparison of taxis, transit, and ridesourcing services in San Francisco. *Transport Policy*, 45:168–178.

Schaller, B. (2021). Can sharing a ride make for less traffic? Evidence from Uber and Lyft and implications for cities. *Transport policy*, 102:1–10.

Sturgeon, L. R. (2019). The impact of transportation network companies on public transit: A case study at the San Francisco International Airport. *Claremont Scholarship*.

Tirachini, A. (2019). Ride-hailing, travel behaviour and sustainable mobility: An international review. *Transportation*, 47(4):2011–2047.

Tirachini, A. and Gomez-Lobo, A. (2020). Does ride-hailing increase or decrease vehicle kilometers traveled (VKT)? A simulation approach for Santiago de Chile. *International Journal of Sustainable Transportation*, 14(3):187–204.

Wallsten, S. (2015). The competitive effects of the sharing economy: How is Uber changing taxis? *Technology Policy Institute*, 22:1–21.

Wang, H. and Odoni, A. (2014). Approximating the performance of a "last mile" transportation system. *Transportation Science*, 50(2):659–675.

Webb, J. (2019). The future of transport: Literature review and overview. *Economic analysis and policy*, 61:1–6.