

Analyzing and Optimizing Shared Mobility Fleet Impacts

Submitted in partial fulfillment of the requirements for

the degree of

Doctor of Philosophy

in

Engineering and Public Policy

Matthew Bremer Bruchon

B.A., English, North Carolina State University

B.S., Computer Engineering, North Carolina State University

B.S., Electrical Engineering, North Carolina State University

M.S., Technology and Policy, Massachusetts Institute of Technology

M.S., Electrical Engineering and Computer Science, Massachusetts Institute of Technology

Carnegie Mellon University

Pittsburgh, PA

May 2022

© Matthew Bremer Bruchon, 2022

All rights reserved.

Abstract

Passenger vehicles enable activity, but they generate unpriced negative externalities such as air emissions and traffic. Those externalities constitute a market failure that may justify policy intervention. Passenger vehicle travel, especially within urban areas, is being transformed by vehicle electrification and by shared mobility options offered by ridesourcing services such as Uber and Lyft. These transformations' impacts on externalities are unclear *a priori*, as is the role of policy to influence them. To investigate these externalities and what options can address them, I use a mixture of simulation and empirical analysis.

One study asks how much an efficiently priced Pigovian tax on unpriced air emissions externalities would incentivize a ridesourcing service to increase vehicle fleet electrification and reduce its emissions. Applying a mathematical optimization of fleet size and powertrain mix with and without a tax, the study finds the tax increases the optimal usage of battery electric vehicles by 5% to 156% and reduces emissions externalities of the optimal fleet by 10% to 22%.

The second study assesses the potential of ridesplitting services—that is, services that combine multiple rides in one car at the same time (e.g., UberPool or Lyft Line)—in a similar manner, asking how much a Pigovian tax on air emissions and traffic (congestion, collisions, and noise) would incentivize increased ridesplitting. Using a dynamic fleet dispatch mathematical optimization problem with and without a tax, the study finds the tax increases use of ridesplitting by 2 percentage points and reduces externalities by 1%.

The third study uses the case study of Chicago's ridesourcing congestion charge policy to ask whether taxes can be practically effective at discouraging solo ridesourcing rides and encouraging ridesplitting. This study conducts an empirical analysis, using a difference-in-differences model to find a downtown zone surcharge reduced total (solo plus ridesplitting) ridesourcing rides by a mean [95% confidence interval] of 8.2% [7.6%, 8.7%], increased ridesplitting opt-in rates by 4.2 [4.0,4.3] percentage points, and increased ridesplitting rides that were successfully matched with another ride by 3.4 [3.3,3.5] percentage points. Special attention is given to spatially autocorrelated outcomes and potential spillover effects, and the direction of results is robust across modeling choices.

To enable these insights, each study considers challenges and applies methods for use of travel data in policy analysis. Specifically, all three apply unsupervised machine learning to make modeling datasets resemble a representative or average set of travel demand, and the first two use supervised learning to estimate varying road network conditions. These steps mitigate avoidable bias in policy findings when using high-volume trip-level travel data.

These studies can inform transportation companies, policymakers, and other stakeholders on the market failures involved with passenger vehicle travel and what policies can help correct for those failures. They also help characterize how effective different transformation pathways in passenger transportation (i.e., electrification or shared mobility) are as solutions to those negative impacts.

ACKNOWLEDGEMENTS

I would like to acknowledge and thank the sponsors of my research. This work was supported by the Carnegie Mellon Center for Climate and Energy Decision Making through a cooperative agreement (Awards SES-0949710 and SES-1463492) between the National Science Foundation and Carnegie Mellon University. This work was also developed as part of the Center for Clean Air Climate Solution (CACES), which was supported under Assistance Agreement No. R835873 awarded by the U.S. Environmental Protection Agency. It has not been formally reviewed by the EPA. The views expressed in this document are solely those of authors and do not necessarily reflect those of the Agency. EPA does not endorse any products or commercial services mentioned in this publication. This project is funded in part by Carnegie Mellon University's Mobility21 National University Transportation Center, which is sponsored by the US Department of Transportation. This material is based upon work supported by Carnegie Mellon University's College of Engineering Moonshot Award for Autonomous Technologies for Livability and Sustainability (ATLAS). This work was also supported by the Department of Engineering & Public Policy at Carnegie Mellon University. This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number ACI-1548562. Specifically, it used the Bridges and Bridges-2 systems, which are supported by NSF award number ACI-1928147, at the Pittsburgh Supercomputing Center (PSC). (Towns et al., 2014; Nystrom et al., 2015)

I thank my advisor and committee members for their assistance developing this work. In particular, I am grateful to my advisor, Jeremy Michalek, for providing the initial vision for this work, helping me learn how to define meaningful research questions and carry out studies to answer them, and consistently giving substantive and useful advice through challenging

model-building processes. I also appreciate his diligence in reviewing numerous revisions to manuscripts and verbose appendices. Inês Azevedo has been a valuable source of input and guidance throughout my research progression, first in Pittsburgh and now from across the country. Nick Muller has provided his expertise on environmental modeling, economics, and econometrics during the development of each study, and it has improved each study substantially. I also thank my external committee member, Kara Kockelman, for providing her expertise in the domain of transportation, responding quickly with thoughtful feedback, sanity checks on my models and results, and many helpful contributions as this dissertation took shape.

Many other people have helped build this work. My colleagues in the Vehicle Electrification Group have actively collaborated on these studies and provided frequent feedback, including Connor Forsythe, Jacob Ward, Lance Chen, Charlotte Andreasen, David Cohen, Aaron Burns, Adam Koling, Arthur Yip, Jiahui Chen, and Johnathan Vicente. Faculty at Carnegie Mellon and elsewhere have also provided valuable input, including Kate Whitefoot, Costa Samaras, Julian Marshall, Alex Davis, Destenie Nock, Daniel Armanios, Corey Harper, Gérard Cornuéjols, Akshaya Jha, and John Helveston.

I also appreciate the help and support of the Engineering and Public Policy department staff, who facilitated my PhD through multiple application cycles, close deadlines, and departmental events.

Finally, I deeply appreciate all of the family and friends who have helped me along the way. Thank you to my parents for providing an environment that helped me learn and grow. Thanks to my sisters and brother for being there when needed. And thank you to my partner, Kristen, for always providing a sounding board, unwavering support, and constant encouragement.

*Matthew B. Bruchon
Pittsburgh, PA
February 2022*

TABLE OF CONTENTS

Abstract	i
Acknowledgements	ii
List of Tables	vii
List of Figures	viii
I. Introduction	1
1.1 Motivation	2
1.2 Overview of contributions	3
1.2.1 First study	3
1.2.2 Second study	4
1.2.3 Third study	5
1.2.4 Joint findings	5
1.2.5 Methodological contributions	6
II. Effects of air emission externalities on optimal ridesourcing electrification	8
2.1 Introduction	8
2.1.1 Prior work	9
2.1.2 Contribution	10
2.2 Data	11
2.3 Methods	13
2.4 Results	15
2.5 Findings	18
2.6 Data availability	19
2.7 Acknowledgements	19
III. Effects of traffic and air emission externalities on optimal ridesplitting	20
3.1 Introduction	20
3.1.1 Prior work	21
3.1.2 Contribution	23
3.2 Methods and Materials	23
3.2.1 Solution method	23
3.3 Data for Chicago case study	31
3.3.1 Cost data	31
3.3.2 Travel data	33
3.4 Results and Discussion	35

3.4.1	“Status quo” use of ridesplitting	35
3.4.2	Changes induced by a Pigovian tax on externalities	37
3.4.3	Variation in results across times and areas	37
3.4.4	Sensitivity to rebalancing strategies	38
3.4.5	Discussion	38
3.5	Data availability	39
3.6	Acknowledgements	39
IV.	Effects of Chicago’s congestion price on solo and shared ridesourcing use	41
4.0.1	Prior work	42
4.0.2	Policy defining treatment	43
4.1	Materials and methods	45
4.1.1	Modeling strategy	45
4.1.2	Data	49
4.2	Results	52
4.2.1	Outcome 1: total TNC ride volumes	52
4.2.2	Outcome 2: percent of riders opting in to pooled ride services	53
4.2.3	Outcome 3: percent of riders matched into a shared ride pool	53
4.2.4	Sensitivity to handling of geographic relationships	53
4.2.5	Discussion and future work	55
4.3	Data availability	57
4.4	Acknowledgements	58
V.	Conclusions and future work	59
5.1	Summary of findings	59
5.1.1	First study’s findings	59
5.1.2	Second study’s findings	61
5.1.3	Third study’s findings	62
5.1.4	Comparison across studies	63
5.2	Policy implications	64
5.3	Methodological findings	66
5.3.1	Optimization methods	66
5.3.2	Machine learning to enable representative policy findings	67
5.4	Modeling lessons learned	69
5.4.1	Importance of benchmarking on realistic models	69
5.4.2	Value of high performance computing resources	70
5.4.3	Limitations of parallelization	71
5.5	Future work	71
5.5.1	Direct extensions to this work	72
5.5.2	Adjacent research areas	73
5.5.3	Data improvements	74
A.	Chapter 3 Algorithms	76
B.	Chapter 2 Manuscript	79

LIST OF TABLES

3.1	Mathematical notation	25
3.2	Optimization parameters	28
3.3	Vehicle assignment problem	28
3.4	Rebalancing for anticipated demand (efficiency-centric version)	30
3.5	Rebalancing for anticipated demand (equity-centric version)	30
3.6	Unmet trip rebalancing	31
3.7	Cost components included in each test case.	32
4.1	Summary of Chicago’s January 6, 2020 TNC ride fees changes, which raised fees, made them time-varying, and introduced a discount for opting into shared ride pooling.	44
4.2	Results for models that group origin-destinations by geographic proximity and include a spatial autocorrelation term using merge distance (main specification).	54
4.3	Multi-way partitioning problem for same-volume clusters	54
4.4	Results for models that group origin-destinations with no consideration of geographic proximity and do not include a spatial autocorrelation term.	56
4.5	Results for models that group origin-destinations by geographic proximity and include an alternative spatial autocorrelation term (average origin-to-origin distance).	56

LIST OF FIGURES

2.1	Illustration of time-space graph showing passenger trips (red), charging arcs (yellow), dispatch arcs (gray), and relocation arcs (blue). Some arcs are omitted for simplicity.	14
2.2	Summary of changes to the optimal ridesourcing fleet when air emission externalities are internalized, including share of fleet-wide vehicle-distance traveled (VDT) from BEVs (left) and total air emissions externality costs per trip-mile (right) in three cities for the optimal fleet technology mix and routing to serve exogenous travel demand. Each measure’s relative change induced by a Pigovian tax (expressed as a percentage of the “no tax” case) is annotated. All cases use a 7% real private firm discount rate, no labor costs, \$50/tonne CO ₂ externality price, the AP3 damage model, \$9.41 million (2018) value of statistical life, and the Pope et al. (2019) concentration-response function.	16
2.3	Vehicle purchases (x-axis) and average utilization (y-axis) by powertrain type for cost-minimizing fleets when excluding (left) and including (right) a Pigovian tax on air emissions. All cases use a 7% real private firm discount rate, no labor costs, \$50/tonne CO ₂ externality price, the AP3 air emissions damage model, \$9.41 million (2018) value of statistical life, and the Pope et al. (2019) concentration-response function.	17
3.1	Private, external, and total costs associated with ridesourcing, with and without a Pigovian tax on external costs. Dollar values are summed across all phases of ridesourcing operations (including “deadheading” travel of empty vehicles) then divided by the total number of trips served.	35
3.2	Time series view of riders per 15-minute block who have a private ride (“Shared” = 0) or who overlap for some duration with one or more other rides in the same car (“Shared” > 0). “Change” represents each ride’s amount of sharing under a Pigovian tax on external costs minus their amount of sharing in the case with no Pigovian tax.	36
4.1	Pre- and post-treatment trends for three outcome variables: daily TNC rides, the percentage of TNC rides for which the rider opted into a shared ride, and the percentage of TNC rides for which the car was shared by two or more pooled ride requests for some period of time. Outcomes are shown as weekly averages, and the treatment date (January 6, 2020) is marked with a dashed line.	48

- 4.2 Distribution (log scale) of ride volumes for each travel corridor. With no grouping step, ride volumes vary drastically across observational units, which may lead to misleading effect estimates. Grouping by geographic proximity (main model specification) nearly matches ride volumes across all units. With geography-agnostic grouping (alternative specification), sizes are equal. 51

CHAPTER I

Introduction

The mobility offered by passenger vehicles enables a wide range of activities, but that travel also imposes substantial negative impacts on society. Most of those negative impacts, including air emissions and traffic, constitute unpriced external costs that are not paid not by the travelers, but rather by society at large. Those unpriced external costs are a form of market failure that can be addressed through public policy, but determining effective policy interventions requires first investigating these complex transportation systems and their underlying demand.

Ridesourcing services from transportation network companies (TNCs) such as Uber and Lyft may exhibit many of the same attributes and (positive and negative) impacts as other modes of passenger vehicle travel. Ridesourcing use warrants special attention due to its rapid growth in absolute and relative terms, and because of uncertainty regarding its overall impacts on cities versus other travel modes. It also warrants special attention because unlike private vehicle travel, the full traveler population can be readily analyzed with trip-level data: every trip is requested electronically, logged by the service provider, and made public (in a small number of cities and with privacy precautions). Findings from ridesourcing use analysis do not directly transfer to broader passenger vehicle travel but may offer suggestive evidence to inform future research directions.

This dissertation investigates the market failures that underlie and aggravate the external costs of ridesourcing, including air emissions and traffic. It investigates those market failures' size, the public policy "levers" that can address them, and the technological, operational, and behavioral changes that can improve outcomes. To yield these insights, it

uses a combination of simulation (via mathematical optimization), empirical estimation (via econometric modeling), and machine learning, extending existing methods and developing custom approaches when necessary.

1.1 Motivation

The societal impacts of the U.S. transportation sector are massive, and transportation systems are evolving rapidly. It has been the largest contributor of greenhouse gases in the nation since 2016, it contributes as much as \$101 billion in annual congestion costs (\$605 per auto commuter), and it results in over 36,000 annual deaths (more than 20% of whom are non-motorists such as pedestrians and cyclists). (United States Environmental Protection Agency, Office of Air and Radiation, Climate Change Division, 2018; Schrank et al.; National Highway Traffic Safety Administration, 2021) The trends of shared transportation and electrification, in particular, are rapidly growing and may be transformational. Ridesourcing services from transportation network companies (TNCs) like Uber and Lyft have more than doubled the for-hire vehicle sector’s share of trips from 2009 to 2017, and that sector is nearing the ridership levels of our urban bus and rail systems. (Conway et al., 2018; Schaller, 2021) Meanwhile, electrification has been called necessary to mitigate the climate crisis. (IPCC, 2018), and policymakers are considering or (in the case of California) enacting ridesourcing electrification targets (George and Zafar, 2018).

The direction of ridesourcing’s impacts going forward are uncertain, including not only emissions but also traffic outcomes (congestion, collisions, and noise). (Schaller, 2021; Ward et al., 2021a) It seems likely that electrification may reduce emissions, and that increased ridesplitting (also referred to as ride pooling, e.g., the use of services such as UberPool/Uber X Share or Lyft Line that combine multiple rides into one car at the same time) may reduce emissions and traffic. However, batteries are constrained by range limits, and many rides cannot be pooled without inducing extra travel and long delay times for pickup detours. Operational modeling can estimate the potential the “levers” of electrification and ridesplitting have to improve outcomes and estimate whether public policies are needed (as implied by the market’s failure to price external costs). This thesis’s first two studies conduct that estimation.

That operational simulation-based estimation is necessitated in part by the relative dearth of empirical data on policy interventions and their impacts on ridesourcing. The largest body of empirical studies in this area has primarily considered TNC effects upon entry and how they vary across regions. (Ward et al., 2021b, 2019; Hall et al., 2018; Babar and Burtch, 2020) Of the studies considering “mature” TNC services and their impacts, many do not use observational travel data (instead using surveys or other data) and do not consider interventions by policymakers. (Hampshire et al., 2017; Rayle et al., 2016; Zgheib et al., 2020) However, a recent ridesourcing congestion charge instituted in Chicago provides a basis for empirical analysis of a policy intervention for ridesourcing impacts. This dissertation’s third study combines highly granular within-city travel data and outcome data for a case study of Chicago and assesses the realized efficacy of a real-world policy.

1.2 Overview of contributions

Each study considers a market failure, a policy lever to address it, and a pathway to improve outcomes. These studies are described in turn.

1.2.1 First study

This study, found in Chapter II, considers electrification of ridesourcing fleets as a means to reduce air emissions, considering those emissions as a market failure of unpriced externalities. It asks: *how might ridesourcing fleets currently use electric vehicles to minimize private costs, and how would electrification and air emissions change if emissions costs were priced in?* This question is investigated using an “oracle”-style joint optimization of ridesourcing fleet vehicle purchases, routing, and—for battery electric vehicles (BEVs)—charging. The study finds that a Pigovian tax shifts fleets away from conventional vehicles in favor of hybrid and battery electric vehicles (BEVs). externality reductions vary from 10% in New York City (where externality costs for both gasoline and electricity consumption are high and a tax induces a partial shift to BEVs), to 22% in Los Angeles (where high gasoline and low electric grid externalities lead a tax to induce a near-total shift to BEVs). The latter number equals \$29 million of annual environmental and health outcomes in Los Angeles alone.

This implies unpriced externalities substantially distort ridesourcing operations, public policy has a clear role to play, and electrification meaningfully lowers emissions. However, the resulting emissions reductions are not extreme. Furthermore, electrification does not address other vehicle travel externalities from increased traffic.

1.2.2 Second study

The second study, found in Chapter III, still considers air emissions and their external costs, but primarily focuses on those externalities resulting from increased traffic (congestion, collisions, and noise). Those traffic costs are, by some estimates, a larger overall source of ridesourcing external costs than air emissions. Ward et al. (2021a) Due to its focus on traffic externalities, this study focuses on ridesplitting rather than electrification as a pathway to reducing external costs. Specifically, it asks: *how might ridesourcing fleets currently use ridesplitting to minimize private costs, and how would ridesplitting and traffic and emissions change if their external costs were priced in?*

That focus on ridesplitting makes dispatch to serve passenger trips of realistic demand density more central than optimal fleet acquisition, so the model optimizes for successive short-term batches of ride requests at a much more granular scale (city block-level routing rather than trip-to-trip chaining), grouping riders into pooled trips where optimal and assigning each trip to a car. The problem formulation is otherwise similar to the first study: the fleet operator's decision-making is guided by minimizing costs, with and without a Pigovian tax on external costs.

Operational changes and external cost reductions are smaller in this study than in the first. A tax increases the use of ridesplitting by 2 percentage points and decreases external costs by 1%. This demonstrates the existence of a market failure pertaining to traffic externalities, but suggests that market failure distorts behavior related to ridesplitting use less substantially. In other words, ridesplitting often does reduce external costs in a meaningful way, but many of those improvements would already be realized by a service operating to minimize its private costs of operation.

1.2.3 Third study

The third study, found in Chapter IV, also estimates how ridesourcing operations and behavior may be influenced by price signals. However, rather than simulating how an optimal decision-maker would minimize a ridesourcing service’s costs given an economically efficient price signal, this study assesses how a real-world ridesourcing ecosystem responded to a price signal chosen by municipal policymakers. More specifically, it considers a congestion charge Chicago instituted in January 2020—an additional \$1.75 in fees on ridesourcing use in the central business district during specific times of weekdays—and the changes that resulted. It asks: *did Chicago’s congestion charge meet its stated goals of discouraging solo ridesourcing rides and encouraging ridesplitting in the central business district?* To answer this question, it conducts difference-in-differences estimation, including the effects of spatially autocorrelated outcomes and possible policy spillover to mitigate potentially biased estimates.

This study finds the congestion charge did move the needle on its policy objectives. Relative to other areas of the city, the increased surcharge in high-charge areas caused a mean [95% confidence interval] reduction of 8.2% [7.6%,8.7%] in total (solo plus ridesplitting) rides, a 4.2 [4.0,4.3] percentage point increase in ridesplitting opt-in rates, and a 3.4 [3.3,3.5] percentage point increase in the rate of riders who not only opted into ridesplitting, but also were successfully matched with another rider.

1.2.4 Joint findings

Taken together, these studies demonstrate that the unpriced externalities of ridesourcing do constitute substantial market failures, they do (to varying degrees across studies) distort decision-making, and that there is a role for policy to play. The first study demonstrates that unpriced air emissions external costs distort the use of electric vehicles, since a Pigovian tax induces substantially more electrification (in terms of both vehicle acquisition and utilization levels) and reduces air emissions meaningfully. The second study estimates the distortion in ridesplitting caused by unpriced external costs to be smaller. The third study cannot say whether behaviors were more “distorted” from an economically efficient outcome before or after the congestion charge policy, but does show that the policy was effective at altering

behavior and operations.

Separately from those primary findings, these studies may provide a broader reason for optimism in that electrification and ridesplitting both can improve societal outcomes with comparatively small changes to the private costs of fleet acquisition and operation. Regulators may take this as suggestive evidence that price-based regulations can nudge ridesourcing operations in a meaningful manner. However, ridesourcing services themselves may also take some heartening news from this: the studies suggest that even with no tax being implemented, if services wish to improve their societal impacts, they have pathways to do so (electrification and ridesplitting) without costing them much more, or even saving them money, depending on assumptions.

It is worth considering in more depth what we might learn from comparing the size and direction of findings across studies, an exercise left for Chapter V. The differences are open to interpretation and may arise from several distinct underlying mechanisms.

1.2.5 Methodological contributions

A notable secondary contribution of this dissertation is to extend methods for modeling transportation fleet operations and impacts. These contributions pertain to computation and to policy relevance.

Computationally, the first study formulates (including custom heuristics) a larger-scale joint acquisition, routing, and battery charging problem than existing formulations would solve. The second study adds computational shortcuts to an existing problem formulation and makes it handle different objectives. In both cases, the code is open-sourced.

To improve policy applicability, all three studies use machine learning to make modeling datasets more representative of average outcomes. The second study uses unsupervised learning to choose a subset of dates that closely resemble a representative subset of travel demand, since optimizing a full year (tens or hundreds of millions of trips) would be intractable on most readily available computing environments. The third study uses unsupervised learning to create observational units of similar travel volumes that maintain geographic relationships, necessary to avoid biased or misleading treatment effect estimates. The first and second papers also apply supervised learning to infer how road network conditions in

between trips vary across time and space.

CHAPTER II

Effects of air emission externalities on optimal ridesourcing electrification

This chapter was developed with coauthors Jeremy Michalek and Inês Azevedo. It is based on work published in:

Matthew B. Bruchon, Jeremy J. Michalek, and Inês L. Azevedo. 2020. [Effects of Air Emission Externalities on Optimal Ridesourcing Fleet Electrification and Operations](#). *Environmental Science & Technology*, 55(5):3188–3200.

An abridged version is provided here, and the full manuscript is included as an appendix.

2.1 Introduction

Vehicle electrification could drastically reduce ridesourcing emissions while also lowering operating costs. Electricity is often cleaner and cheaper than gasoline per vehicle distance traveled (VDT), and for heavily used vehicles lower fuel costs and operation emissions may offset higher upfront costs and manufacturing emissions. Policymakers are considering or (in the case of California) creating ridesourcing electrification targets (George and Zafar, 2018). TNCs have set goals to fully electrify in order to reduce emissions, including EV driver subsidies (Uber, 2018).

However, the premise that full fleet electrification is viable or desirable warrants investigation. Battery electric vehicles (BEVs, which plug in to charge and rely entirely on electricity stored in large battery packs) have a higher upfront cost than conventional vehicles (CVs); battery production emissions are non-trivial (Tessum et al., 2014; Michalek

et al., 2011); and, depending on region and design, BEVs do not always reduce emissions externalities relative to CVs (Weis et al., 2016; Yuksel et al., 2016; Holland et al., 2019, 2016). Furthermore, BEVs have limited range, must detour to recharge (increasing VDT), and cannot serve demand while charging (increasing fleet size needed to satisfy demand). In contrast, gasoline hybrid electric vehicles (HEVs) have no additional range constraints but do burn gasoline and emit tailpipe pollution.

In general, it may be that the lowest-cost or lowest-emission fleet does not use a single homogeneous technology but, rather, a mixture of technologies, with different duty cycles (e.g., peak versus off-peak) being served by different technologies. Operational modeling is needed to the constraints and tradeoffs of BEVs relative to HEVs and CVs.

2.1.1 Prior work

A body of literature considers operations and outcomes of electrified vehicle fleets, but the question of electrification’s role within a ridesourcing fleet’s optimal technology mixture and its impact on resulting emissions is relatively unexplored.

Some studies use agent-based modeling (ABM) to simulate fleet routing, often in tandem with a second model, in order to explore operational impacts of shared all-EV fleets. In two studies, Bauer et al. find that an all-EV fleet operating could reduce private costs and emissions relative to an all-CV or all-HEV, and that coordinated charging would allow BEVs to meet the same level of service as CVs at lower cost (Bauer et al., 2018, 2019). Chen et al. (2016) find EVs can meet ridesourcing demand with minimal increased empty VDT, but only if the fleet size increases (Chen et al., 2016); a follow-on study found that an all-EV taxi fleet could reduce cumulative GHG emissions 60% in the base case (Gawron et al., 2019). Sheppard et al. (2019) estimate that 12.5 million vehicles could replace the fleet of 276 million personally owned vehicles (Sheppard et al., 2019). Chen & Kockelman (2016b) estimate an all-EV fleet could capture 14% to 39% of all passenger trips within the Austin, Texas region, depending on pricing (Chen and Kockelman, 2016).

Our study differs in that it optimizes a mixture of fleet technologies under different objectives and considers multiple cities. Also, to avoid the potential of ABM to introduce unquantifiable algorithmic biases (and to address similar limitations of heuristic approaches),

we pair heuristics with mathematical optimization to understand heuristic quality, to gain intuition on their biases, and to compare fairly across cases.

There is a separate stream of optimization research formulating the problem of range-constrained all-EV fleet routing. These formulations are often tested on small benchmarks (exact solutions for 100-200 trips or heuristic solutions for several hundred more), rarely consider external costs, and rarely jointly optimize purchases and routing even in when a fixed mixture of powertrains is assumed (Pourazarm et al., 2016; Sassi et al., 2015). Optimizing fleet size and mix at scale requires careful formulation and problem-specific heuristics, which our study contributes for its problem (applied to an instance of 5,000 trips).

In the grey literature, a 2019 report examined powertrain choice from the perspective of a TNC driver’s costs of vehicle ownership (Pavlenko et al., 2019), finding HEVs are financially favorable and BEVs may become favorable around 2023-2028, using assumptions for factors such as the total distance traveled per year that, in practice, vary across vehicles in the fleet. A later analysis by the same group found that a per-trip fee, indexed to tailpipe emissions, between \$0.58-\$1.12 would make BEVs financially favorable to HEVs (Slowik et al., 2020). In the related context of car-sharing fleets, in which the user pays for short-term rental of a car and drives it themselves, Zoepf (2015) finds BEVs reduce private costs when 20-40% of a CV fleet is electrified but increase private costs beyond that threshold (Zoepf, 2015).

2.1.2 Contribution

We investigate the optimal technology mix and operations of a ridesourcing fleet whose operator has perfect foresight of exogenous (inflexible) passenger trip requests and perfect control over fleet acquisition and routing. Centralized acquisition may represent TNCs that have owned or leased vehicles (Toyota Motor Corporation, 2018; Lyft, Inc., 2020a), a future with autonomous fleets (Lyft, Inc., 2020b; Uber Technologies, Inc., 2020), or purpose-built ridesourcing fleet vehicles (Welch, 2020; Sun and Shirouzu, 2019). Centralized routing may grow with autonomous technology (whereas today’s services centrally nudge routing via human drivers responding to price signals). Also, fleet-wide regulations (e.g., California’s regulation of fleet-wide CO₂/passenger-mile (California State Senate, 2018)) increase the role of centrally coordinated acquisition and routing.

We assess the potential of electrification by comparing costs and emissions of CV, HEV, BEV, and optimally mixed fleets across a range of scenarios. By comparing cases that include or exclude emission externality costs in the objective, we assess the degree to which unpriced emissions externalities bias fleet outcomes away from socially optimal solutions and consider whether policy intervention may be therefore justified on economic efficiency grounds.

We contribute to the prior literature by (1) constructing a mixed-integer optimization model with heuristics that make meaningfully-sized problems tractable and provide near-optimal solutions for fair comparisons across scenarios and (2) applying the model to characterize how the optimal technology mix, operations, and life cycle air emissions externalities of a TNC fleet change across scenarios representing geographic and temporal variation, uncertainty, and the internalization of air emissions externalities (as a Pigovian tax passed through to the fleet operator). Our model is also unique in its treatment of vehicle costs, incorporating into the optimization the effect of vehicle usage on period of use, resale value at end of use, and the resulting discounted future cash flow.

2.2 Data

Trips

We instantiate the model using a 2017 dataset of 1.5 million trips served by RideAustin (a nonprofit ridesourcing service in Austin, Texas). We use these Austin trips to also model Los Angeles and New York City (varying private and external costs by region but not travel demand). We sample down to 5,000 trips across seven representative days using weekday-season and special event categories. We sample each category proportionally to its average daily demand, but costs and distance values of travel arcs contained within each representative day (which affect capital costs of each vehicle) are scaled up to annual quantities based on the number of days per year each category entails. Our model also requires a distance and duration for relocations from trip to subsequent trip (or to the charge station), which we estimate using k-nearest neighbors regression (Wang et al., 2016) on the trip data.

Vehicles and fuel

We model a typical present-day ridesourcing vehicle with otherwise-identical CV, HEV, and BEV counterparts; of model year 2018 options in the United States market, there are five light-duty passenger vehicles with BEV and CV variants. Of the five model year 2018 vehicles with BEV and CV variants, the Kia Soul has backseat passenger space best suited for ridesourcing, so we adopt its price and energy efficiency. To capture that efficiency varies with driving conditions (Yuksel et al., 2016), we use each travel arc’s estimated driving speed to interpolate efficiency between standard EPA test city and highway drive cycles (US Department of Energy Office of Energy Efficiency and US Environmental Protection Agency, 2018).

Given each vehicle’s annualized distance traveled d as determined by the optimization, we assume the fleet resells the vehicle at 12 years or 170,000 miles (i.e., at N years, where $N = \min(N_{\text{MAX}}, \frac{170,000}{d})$). Given private firm discount rate r and vehicle purchase price p , the private costs of each vehicle investment are:

$$\kappa = \left(1 - \frac{v(N, d)}{(1 + r)^N}\right) p \times f_{\text{CR}}(r, N) \quad (2.1)$$

The term $f_{\text{CR}}(r, N)$ is a capital recovery factor (dependent on resell year). The age- and mileage-dependent resale value function $v(N, d)$ was estimated by regressing Kelley Blue Book used market values on age and miles driven (separately for each powertrain).

Each region’s electricity and gasoline prices come from EIA data (US Energy Information Administration, 2018b,a).

Air emissions externalities

To compute external costs per unit of greenhouse gas emissions, we adopt the social cost of carbon \$50 per ton of CO₂ equivalent estimated by the Interagency Working Group on the Social Cost of Carbon (United States Government, 2016). For conventional air pollutants, external costs depend on emission location, and we use the AP3 reduced complexity model (Muller, 2014) to compute and monetize estimated health damages associated with these emissions, assuming a value of statistical life of \$9.41M (2018) and the Pope et al. (2019)

concentration-response relationship (Pope et al., 2019). We use two alternative reduced complexity models, InMAP and EASIUR, in sensitivity cases (Tessum et al., 2017; Heo et al., 2016).

Emissions factors for tailpipe combustion, fuel refining, materials mining, and vehicle manufacture are taken from Argonne National Laboratory’s GREET model. (Argonne National Laboratory, 2019) Each analysis region’s refining emissions are sited based on refining capacity of counties in nearby states. Mining and manufacture emissions are sited using U.S. county employment counts with relevant industry codes.

To model grid emissions, we use marginal AP3 emissions factors estimates generated by replicating an existing Siler-Evans et al. (2012) regression methodology. (Siler-Evans et al., 2012; Azevedo et al.)

2.3 Methods

We optimize fleet composition (mix of CVs, HEVs, and BEVs) and operations (vehicle routing and BEV charging) in order to minimize the cost of satisfying exogeneous demand (origin and destination location and time) under a range of scenarios. Figure 2.1 illustrates our modeling framework with an example. Vehicle purchase choices determine the vehicles available to dispatch (left). Routing options, jointly optimized with purchases, are represented using a graph, where each vertex (dot) represents a specific place and time, and the arcs connecting them include available options for:

- Trip arcs: passenger trip requests that must be served
- Charging arcs: spending time parked (divided into 15-minute charging increments) at a charging location while recharging a battery or waiting for the next trip
- Dispatch arcs: “Deadheading” from a vehicle’s home base to the first passenger trip request
- Return arcs: “Deadheading” from a vehicle’s final passenger trip back to its home base
- Relocation arcs: ”Deadheading” from the end of one passenger trip to the beginning of a next passenger trip or between passenger trips and recharge locations

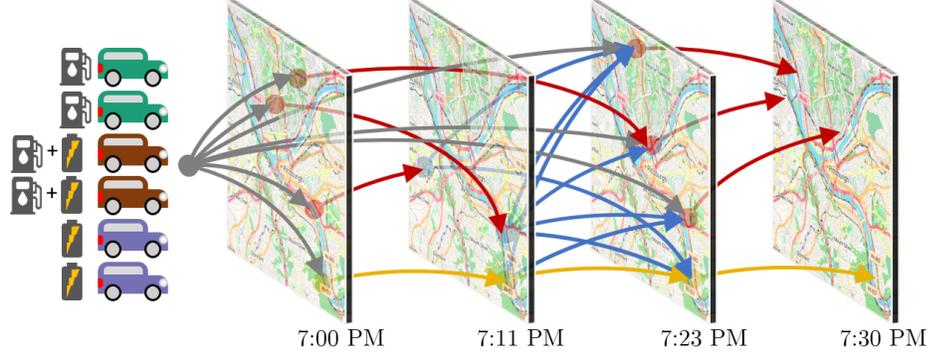


Figure 2.1: Illustration of time-space graph showing passenger trips (red), charging arcs (yellow), dispatch arcs (gray), and relocation arcs (blue). Some arcs are omitted for simplicity.

Formulation

Our full optimization formulation is provided in the appendix. In all test cases, the objective function, Eq.(2.2), sums vehicle acquisition costs κ_k , gasoline and per-mile maintenance costs $c_{k,i,j}$, and battery charging costs c_t (which do not vary with time in the base case). In cases where air emissions externalities are internalized ($\tau = 1$), the fleet also considers a Pigovian tax on externalities from manufacturing emissions δ_k , tailpipe and fuel refining emissions $d_{k,i,j}$, and marginal grid emissions d_t (which vary with time in the base case):

$$\begin{aligned} \text{minimize}_{\mathcal{X}} \quad & \sum_{k \in \mathcal{K}} \kappa_k + \sum_{(i,j) \in \mathcal{A}} c_{k,i,j} a_{k,i,j} + \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}_+} c_t \Delta q_{k,t}^{\text{CHG}} + \\ & \tau \left[\sum_{k \in \mathcal{K}} \delta_k + \sum_{(i,j) \in \mathcal{A}} d_{k,i,j} a_{k,i,j} + \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}_+} d_t \Delta q_{k,t}^{\text{CHG}} \right] \end{aligned} \quad (2.2)$$

The set of decision variables \mathcal{X} includes the number of vehicles n_k of each powertrain type k purchased, assignments $a_{k,i,j}$ of vehicles k to arcs (i,j) , charge level $q_{k,t}$ and energy charged from the grid $\Delta q_{k,t}^{\text{CHG}}$ for each vehicle k at each discrete time point t , and total annualized capital cost κ_k for each vehicle type (determined by vehicle utilization levels) for all vehicle types $k \in \mathcal{K}$, arcs $(i,j) \in \mathcal{A}$, and times $t \in \mathcal{T}$.

Constraints ensure vehicle flow conservation, enforce demand satisfaction and vehicle purchases, track battery charge levels, limit charging speeds, and endogenously compute vehicle acquisition costs κ_k and manufacturing externalities δ_k .

Heuristics

For problems similar to this, the optimization state of art for exact solutions is around 200 trips, beyond which solution times become prohibitive. Particularly due to BEV charge constraints, our formulation cannot be solved by a commercial solver alone.

To improve scalability, we introduce a set of problem-specific heuristics that find near-optimal solutions quickly for our sample size of 5,000 trips. These heuristics rely in part on minimum-cost network flow formulations, which can solve extremely quickly (Klein, 1967). In short, they eliminate less optimal relocation arcs, coordinate fleetwide charging in an iteratively more realistic manner, and use mathematical optimization to bring our heuristically-developed “starting point” solution closer to optimality. Details are provided in the appendix.

2.4 Results

Pigovian tax impact on costs

With no Pigovian tax, private costs range from 45.5¢ to 49.0¢ and external costs, shown in the right-hand portion of Figure 2.2, range from 10.1¢ to 14.8¢ per trip-mile (that is, total annualized life cycle costs divided by the number of annualized miles of passenger trips served). Depending on variability of regional costs, a tax leads the fleet to increase its usage of BEVs by 5% to 156% and dispatch these vehicles in a manner that reduces emissions externalities per trip-mile by 10% to 22%.

In absolute terms (x-axis of the right-hand side of Figure 2.2), these reductions range from 1.3¢ to 2.3¢ per trip-mile. A recent Fehr & Peers consulting report estimated that Uber & Lyft drive 104 million monthly trip-miles in Los Angeles (Balding et al., 2019). Multiplying those trip-miles by the 2.3¢ per trip-mile decrease in externalities, we can roughly estimate external cost reductions of \$29 million per year in Los Angeles (\$24 million in reduced criteria pollutant emissions and the remainder in reduced GHG emissions).

External cost reductions are greatest in percentage and absolute terms in Los Angeles, where fuel emissions externalities are high and electricity generation externalities are low relative to the other cities modeled (a larger relative difference for criteria pollutant external

costs than for GHGs). In percentage terms, they are smallest in New York City, where the external costs of electricity generation are highest of the three cities. Austin sees the largest increase in BEV usage, partially because of lower gas prices that lead a private cost-minimizing fleet to use many CVs and few BEVs. However, due to lower health damages per unit of tailpipe emissions and a less “clean” grid than Los Angeles, Austin’s external cost reductions fall between the other two cities’ in percentage terms, and are smallest in absolute terms.

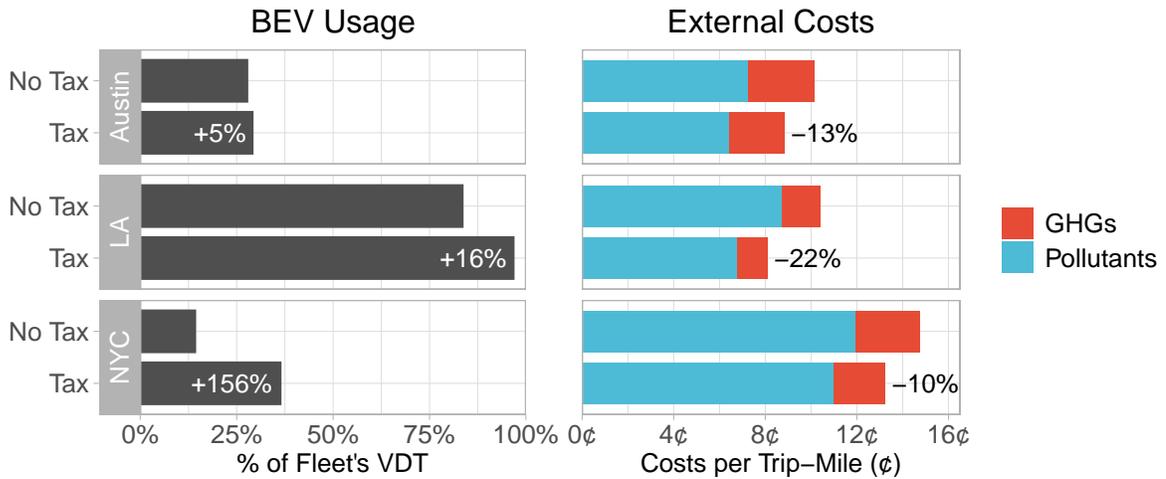


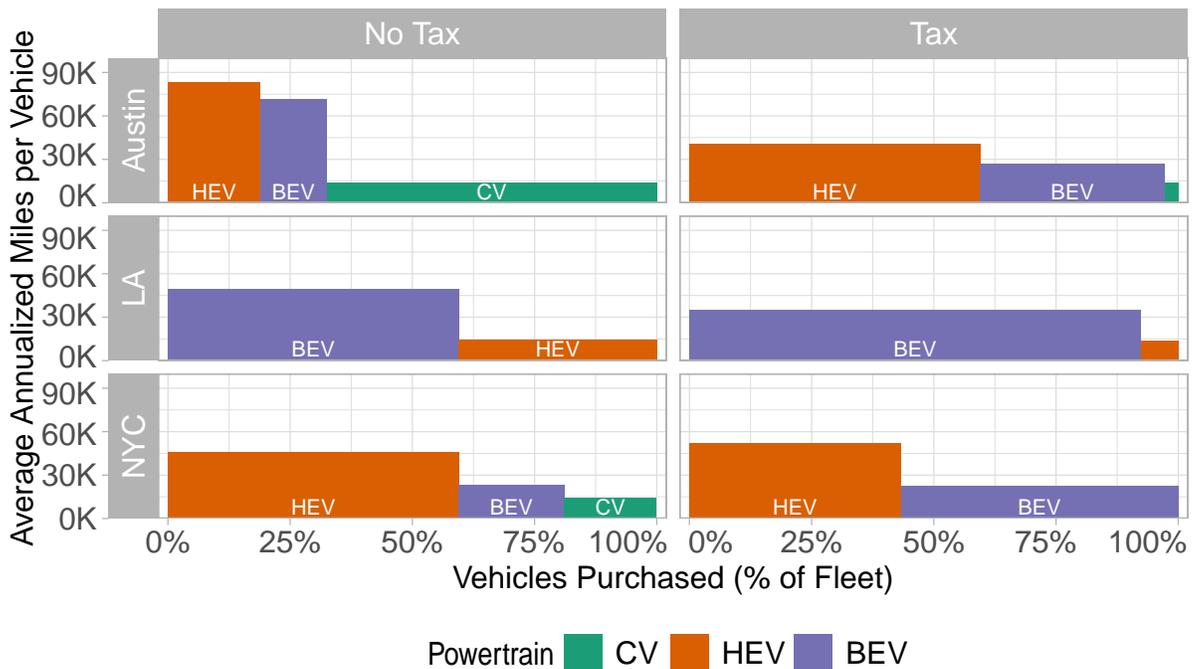
Figure 2.2: Summary of changes to the optimal ridesourcing fleet when air emission externalities are internalized, including share of fleet-wide vehicle-distance traveled (VDT) from BEVs (left) and total air emissions externality costs per trip-mile (right) in three cities for the optimal fleet technology mix and routing to serve exogenous travel demand. Each measure’s relative change induced by a Pigovian tax (expressed as a percentage of the “no tax” case) is annotated. All cases use a 7% real private firm discount rate, no labor costs, \$50/tonne CO₂ externality price, the AP3 damage model, \$9.41 million (2018) value of statistical life, and the Pope et al. (2019) concentration-response function.

Pigovian tax impact on powertrains

These external cost reductions are accomplished in each city not only by shifting VDT away from gasoline usage (in CVs and HEVs) and towards electricity usage (in BEVs), but also by a corresponding change in vehicle purchases. Figure 2.3 illustrates for each city, with and without the Pigovian tax, the share of vehicle purchases for each powertrain (out of an optimal fleet size ranging from 37 to 39 vehicles in the base case) and the annual miles driven per car of each powertrain type. Across the three cities, the number of BEVs in the

optimal fleet increases by 63% to 180% when a Pigovian tax is imposed on the fleet, and BEVs' total vehicle-distance traveled increases by 5% to 156%. HEVs serve virtually all of the remaining demand in these three Pigovian tax cases, while CVs are at or near 0% of the fleet's purchases and distance traveled.

Figure 2.3: Vehicle purchases (x-axis) and average utilization (y-axis) by powertrain type for cost-minimizing fleets when excluding (left) and including (right) a Pigovian tax on air emissions. All cases use a 7% real private firm discount rate, no labor costs, \$50/tonne CO₂ externality price, the AP3 air emissions damage model, \$9.41 million (2018) value of statistical life, and the Pope et al. (2019) concentration-response function.



Value of mixing powertrains

Across cities, a fleet that optimally determines the mixture of powertrains to purchase and dispatch substantially reduces its private costs and the air emissions externalities it produces. Relative to an all-CV fleet, a fleet optimized for private costs reduces private costs by 5% to 14% and, in doing so, also reduces emissions externalities by 14% to 66%. The best homogeneous fleet does not depend on a Pigovian tax: it is all-HEV in Austin and New York City and all-BEV in Los Angeles regardless of whether a tax is included. Relative to the best homogeneous fleet, the mixed fleet optimized without a Pigovian tax reduces

private costs by 1% to 4%, and the mixed fleet with a Pigovian tax reduces social costs by 1% to 4%.

Sensitivity cases

We tested sensitivity to parameters including: battery capacity, battery price, discount rate, resale value, labor costs, electricity rates, grid emissions, emissions damage model, and the social cost of carbon. Sensitivity case results are in the Supporting Information of the published manuscript. Bruchon et al. (2020)

2.5 Findings

Across a wide range of scenarios, these results consistently suggest that internalizing air emissions externalities results in a greater degree of electrification (shift from CV to HEV and BEVs and shift from HEV to BEV) as well as operational changes that together reduce air emissions externality costs (by 10% to 22% in the base case and 4% to 75% across sensitivity cases), and lower social costs (by 2% to 3% in the base case and 0% to 18% across sensitivity cases). This suggests a potential role for policy because when emissions externalities are unpriced, firms have incentives to lower private cost in ways that increase air emissions, implement a lower degree of electrification, and charge BEVs when the grid is less clean than socially optimal. While the change in social cost is fairly small across most of the scenarios examined, the change in who bears the cost (private versus external costs) can be significant—as estimated above for Los Angeles, as high as \$29 million of annual environmental and health outcomes (a number which will likely grow with increased ridesourcing usage).

Pigovian taxes offer efficiency and flexibility, but in the absence of such an option, other policies that encourage similar outcomes, such as policies encouraging increased electrification, could potentially improve economic efficiency. However, any such policy should be designed with care. A blunt instrument favoring one technology over others may not be desirable because (1) the optimal fleet is generally a mixed fleet; (2) beyond fleet composition, it is important how intensively each vehicle type is used; and (3) factors that vary with location and over time, like energy prices, vehicle cost, population density, and grid emission factors,

can dramatically change the degree of electrification that is optimal.

2.6 Data availability

Data and code supporting the findings of this study are available at

<https://github.com/mbbruch/FleetElectrification>.

2.7 Acknowledgements

This work was supported by the Carnegie Mellon Center for Climate and Energy Decision Making through a cooperative agreement (Awards SES-0949710 and SES-1463492) between the National Science Foundation and Carnegie Mellon University. This work was also developed as part of the Center for Clean Air Climate Solution (CACES), which was supported under Assistance Agreement No. R835873 awarded by the U.S. Environmental Protection Agency. It has not been formally reviewed by EPA. The views expressed in this document are solely those of authors and do not necessarily reflect those of the Agency. EPA does not endorse any products or commercial services mentioned in this publication. It was also supported by the Department of Engineering & Public Policy at Carnegie Mellon University. This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number ACI-1548562. Specifically, it used the Bridges system, which is supported by NSF award number ACI-1928147, at the Pittsburgh Supercomputing Center (PSC). (Towns et al., 2014; Nystrom et al., 2015)

CHAPTER III

Effects of traffic and air emission externalities on optimal ridesplitting

This chapter was developed with coauthors Jeremy Michalek and Connor Forsythe.

3.1 Introduction

Chapter 1 optimized ridesourcing fleets with and without a Pigovian tax on air emission externalities and found that electric vehicles are consistently used more when the market failure of unpriced externalities is fixed. However, that study found that even if electrification improves outcomes, it does not cause massive emissions reductions or “deep” decarbonization. Furthermore, electrification does nothing to reduce traffic externalities (congestion, collisions, and noise) resulting from vehicle miles traveled; these traffic externalities may be more costly than those from emissions. (Ward et al., 2021a)

For that reason, this study considers whether market failures lead to distorted behavior along a different dimension of operations: the use of ridesplitting (also referred to as ride pooling, e.g., the use of shared ride services such as UberPool, Uber Share X, or Lyft Line that combine multiple rides into one car at the same time). This study repeats the high-level process described in Chapter 1—simulating an optimal fleet’s operations with and without internalized externalities—but adds the external costs of traffic (more affected by ridesplitting than by powertrain choice) and the decision variable of ridesplitting. While this study asks a parallel question to the first study, the underlying optimization and solution steps are almost entirely different.

This study will assess whether market failures of unpriced traffic and air emissions

externalities lead to sub-optimal usage of ridesplitting services and how much potential ridesplitting holds to improve outcomes of interest.

3.1.1 Prior work

It seems likely that some amount of ridesplitting may reduce the externalities imposed on society from ridesourcing, but it is an open question how viable ridesplitting services are and how widespread they can become. Henao & Marshall (2018) conducted a quasi-natural experiment, driving for ridesourcing services and manually collecting ride data; in their period of data collection in Denver, Colorado, only 15% of UberPool and Lyft Line rides were actually matched with another rider ($n=54$). (Henao and Marshall, 2019) Schweiterman et al. (2019) scraped a large number of prices and wait times for Uber, UberPool, Lyft, and Lyft Line. (Schwieterman, 2019) Using USDOT travel time valuations for average and business travel, that study found that UberPool discounts exceeded on average the value of time lost relative to Uber, but that Lyft Line’s efficacy depended on neighborhood and on personal vs. business travel.

Attempts have been made to mathematically characterize the parameters that influence the viability of ridesplitting services. Tachet et al. (2017) develops and validates in four cities a “scaling law” for sharability (the proportion of trips that can be shared in a neighborhood). (Tachet et al., 2017) It finds that quantity depends only on four quantities (assuming uniformly distributed demand): tolerable delay times, road speeds, geographic area, and frequency of trip requests. Similarly, Bilali et al. (2019) build an analytical model calibrated on trip data, finding that delay times substantially alter shareability and that (particularly in areas with low tolerance of delay times or low trip density) reservation systems can improve shareability.

The question of how many shared rides a ridesplitting service actually leads to has a sizeable impact on outcomes. For example, a recent parametric analysis found that while switching a private vehicle ride to a private ridesourcing ride may increase externalities on average by about 30¢/trip, switching a private vehicle ride to a ridesplitting service may on average increase or decrease externalities, depending on the ride’s likelihood of actually being pooled with another rider. (Ward et al., 2021a) Even when rides are shared, it is

not obvious that VMT reductions always result. A 2018 California Air Resources Board ridesourcing emissions inventory, using highly granular trip data across the state, found that time-weighted occupancy rates are only slightly higher for pooled rides than for private rides—partially because parties who request a ridesplitting service include fewer people on average. (California Air Resources Board, 2021)

An increasingly large stream of literature, which we extend, uses optimization to model “online” ridesplitting services responding to successive batches of ride requests by altering existing cars’ trajectories and matching cars to rides (some of which may be shared). This stream of literature began with Santi et al. (2014), which demonstrated how “shareability networks” can make this optimization tractable, was demonstrated with dynamic ridesourcing fleets in Alonso-Mora et al. (2017), and was subsequently extended to questions of fleet composition and optimal pricing. (Santi et al., 2014; Alonso-Mora et al., 2017a,b; Wallar et al., 2019b,a) A separate model by Simonetto et al. solves a slightly more myopic, but more scaleable version of the ridesplitting problem using linear programming. (Simonetto et al., 2019) This stream of studies demonstrates that ridesplitting problems can be solved in “real-time” (a near-optimal solution returned in runtime less than or equal to the optimized batch of requests).

However, this stream of literature generally focuses on optimizing quality of service metrics (e.g., delay times and wait times) and does not consider policy interventions to reduce VMT. One exception, by Liu et al., considers a secondary test case in which private rides are charged a flat \$2 fee and the ridesourcing fleet finds (via Bayesian optimization) a new equilibrium pricing and operational scheme, finding the fee would increase transit mode share but also increase ridesplitting mode share (from 14% to 51% of all trips), reducing VMT by 10.5%. (Liu et al., 2019) This study suggests incentives have power to alter the use of ridesourcing—further motivating this study—but does not explicitly consider the size of external costs that may exist, compute an efficient level of tax specific to each ride, or identify the effects of market failure caused by unpriced externalities.

3.1.2 Contribution

We investigate the optimal operations of a ridesourcing fleet, reacting dynamically to new batches of requests, whose operator must minimize costs while meeting (inflexible) passenger trip requests. This informs the empirical literature in part by providing a simulation-derived estimate of how and when ridesplitting is viable as a service, and how that differs within a city (rather than only in aggregate).

Relative to the stream of literature beginning with Santi et al. (2014) and continuing to Liu et al. (2019), which enabled optimization of the dynamic ridesplitting problem, our study uses the same modeling approach (with some modifications) but differs in an important way. (Santi et al., 2014; Liu et al., 2019) Unlike most of those studies, we consider a fleet optimizing not for quality of service metrics, but for total costs (including quality of service as measured through the value of travel time), which (arguably) could more closely resemble an efficiently operated fleet. Importantly, this enables us to estimate the difference in outcomes when the fleet must face a Pigovian tax on externalities, which allows us to answer a different research question: how large of a market failure exists for policy to address? We also differ in that we consider how pooling is used differently across neighborhoods in a city and whether different neighborhoods are affected differently by policy interventions.

3.2 Methods and Materials

We construct an optimization that centrally dispatches a fleet of ridesourcing vehicles to serve exogenous ride requests, minimizing varying cost functions across a representative set of days and comparing outcomes. This section describes our method of defining and solving this problem and our approach to building a representative case study.

3.2.1 Solution method

We construct an optimization that centrally dispatches a fleet of ridesourcing vehicles to serve exogenous ride requests, minimizing a cost function $\gamma(v, t)$ across a representative set of days and comparing outcomes. To optimize vehicle dispatch and trip matching, we adapt the concept of trip shareability networks described in Santi et al. (2014), applied to dynamic fleet operations in Alonso-Mora et al. (2017). (Santi et al., 2014; Alonso-Mora et al.,

2017a) This method optimizes ridesourcing vehicle assignment in sequential batches of ride requests. For each batch, we execute the following solution steps, considering each vehicle’s in-progress trips and current location:

1. **Request-to-vehicle enumeration:** Build a graph describing the costs to optimally route each vehicle to serve each single-rider *request* it can feasibly serve (“RV graph”).
2. **Request-to-trip enumeration:** Starting from the RV graph, list all groupings of *requests* into k -request *trips* that a perfectly located, empty vehicle could feasibly serve (“RT list”). Find rules for which $k+1$ -request trips would be certainly infeasible for a specific vehicle if it could not serve a given k -request trip (“dependency map”).
3. **Trip-to-vehicle enumeration:** Using the RV graph, RT list, and dependency map, build a graph describing the costs of optimally routing each vehicle to serve each trip it can feasibly serve (“RTV graph”).
4. **Problem reduction:** Heuristically reduce the number of options in the RTV graph (“reduced RTV graph”) to improve scalability.
5. **Vehicle assignment:** Find minimum-cost matchings in the reduced RTV graph.
6. **Vehicle rebalancing:** Heuristically route available vehicles toward any remaining unserved requests.

We subsequently describe each step using notation summarized in Table 3.1.

Step 1: Request-to-vehicle enumeration

As in Alonso-Mora et al. (2017),(Alonso-Mora et al., 2017a) we first determine which vehicles $v \in \mathcal{V}$ are able to serve which new ride requests $r \in \mathcal{R}$ at what cost. This step (Algorithm 1 in Appendix A) loops in parallel over each r - v combination. For each combination, a function $\gamma(v, r)$ computes the cost of v ’s minimum-cost turn-by-turn vehicle route meeting three requirements: (1) r ’s total delay (pick-up delay plus en-route detours) is below the maximum allowed, δ^{MAX} ; (2) any passengers already in v from a previously optimized request batch still reach their destination with delay at or below δ^{MAX} ; and (3) the chaining of multiple

Table 3.1: Mathematical notation

Label	Type	Description
\mathcal{V}	Set	All vehicles
\mathcal{R}	Set	All requests in current optimization time frame
\mathcal{T}	Set	All trips (combinations of requests) an ideal car could serve
G_{RV}	Graph	Matchings of request nodes r to vehicle nodes v with edge cost $c_{v,r}$
G_{RTV}	Graph	Matchings of trip nodes t to vehicle nodes v , with edge cost $c_{v,t}$, with additional arcs linking each t to member requests r
\mathcal{I}_r^R	Set	Trips t for which an edge exists in G_{RTV} linking request r to t
\mathcal{I}_t^T	Set	Vehicles v for which an edge exists in G_{RTV} linking trip t to v
\mathcal{V}^U	Set	Vehicles that are unoccupied with and unassigned following an optimization
\mathcal{R}^U	Set	Requests not yet assigned a vehicle following an optimization
$a_{v,t}$	Variable	Assignment of vehicle v to trip t
o_v^{MAX}	Parameter	Maximum occupancy of vehicle v
ρ	Parameter	Arbitrarily large reward per request served
n_t	Parameter	Number of requests in trip t
$\gamma(v, r)$	Function	Computed cost of vehicle v 's minimum-cost vehicle routing for solo ride request r
$\gamma(v, t)$	Function	Computed cost of vehicle v 's minimum-cost vehicle routing for trip t

pickups and dropoffs does not violate the vehicle’s maximum passenger occupancy o_v^{MAX} . The option of the route serving r with v is stored at reduced dimensionality as an arc in an RV graph G_{RV} , linking r ’s node to v ’s node with edge cost $\gamma(v, r)$. If no route can meet the three requirements, then no corresponding arc is added.

Our implementation of $\gamma(v, r)$ is built upon a depth-first search method implementation by Haojia Zuo that uses a graph partitioning tree framework and the open-source METIS software for shortest-path finding. (Zuo, 2013; Karypis and Kumar, 1999) We use this graph partitioning framework in an “offline” manner, precomputing in advance a lookup list of lowest-cost travel times and (private and external) costs between each pair of nodes in the road network. We then use this lookup list for quick retrieval during the “online” optimization stage. The cost components considered by $\gamma(v, r)$ vary by modeling scenario as described in Section 3.3.1, but maximum travel delay δ^{MAX} is taken as a constraint in all scenarios. This represents a fleet’s need to provide a minimally acceptable level of service across different incentive structures.

Step 2: Request-to-trip enumeration

Unlike prior works using this approach to dynamic vehicle assignment, before enumerating each vehicle’s feasible groupings of requests into trips $t \in \mathcal{T}$, we first do so for an imaginary ideal vehicle v^* that has no pre-existing en-route passengers and is sited perfectly to serve each potential combination of requests. This reduces the number of vehicle-trip pairs that

need computing without compromising optimality. Rather than repeating the majority of enumeration steps once per each vehicle, it instead does them once up front. It also determines shortcuts to make each v 's enumeration faster, improving the total time needed to build \mathcal{T} .

This step (Algorithm 2 in Appendix A) proceeds successively from trips of size $k = 2$ to k^{MAX} . The value k is not necessarily the number of riders in v at any specific moment but rather is equivalent to n_t , the number of requests in trip t , arranged in a feasible vehicle route involving a chain of pickups and dropoffs. Two steps are carried out for each k : all trip candidates T^* of size k to test for feasibility are exhaustively enumerated, then each trip in T^* is tested for feasibility with ideal vehicle v^* . We parallelize each step and take computational shortcuts enabled by updating forward- and reverse-dependency lists of trips (i.e., lists of k -trips that are only feasible if a specific $k - 1$ -trip is). Each time a trip is added to \mathcal{T} , it is also added as a dependency of each size $k - 1$ antecedent trip in D , a map of trips to “exclusion rules” based on set membership.

For enumerating 2-trips (that is, trips composed of two requests), if any vehicle was fully connected to all $r \in \mathcal{R}$ in the RV graph G_{RV} , then all pairwise combinations of requests are enumerated; otherwise, only those pairs that are both feasible by a specific $v \in \mathcal{V}$ are.

For $k \geq 2$, we build increasingly large trips based on lower-cardinality trips already found, relying on the property that in order to be feasible, each k -trip must be a complete clique of k requests in a shareability graph, where each smaller subclique is itself a complete clique (as explained fully in Santi et al. (2014)(Santi et al., 2014)). More specifically, each feasible trip clique of size k must include k unique $k - 1$ -request subcliques (each differing only by which individual request it lacks), and for $k \geq 3$, each $k - 1$ request subclique is itself the union of $k - 1$ subcliques of size $k - 2$. To more quickly identify which size $k - 1$ trip candidates differ by only one trip, and thus should be unioned to form a trip candidate of size k , we iterate over $k > 2$ -trips with at least 2 dependent $k - 1$ trips, enumerating size k pairwise unions of those dependent $k - 1$ -trips and counting how many times each size k union is identified.

Using $k = 3$ as an example, the trip $\{r_1, r_2, r_3\}$ need only be checked for feasibility if it is formed by the unions of k distinct pairs of feasible $k - 1$ -trips: $\{r_1 \cup r_2\}$, $\{r_1 \cup r_3\}$,

and $\{r_2 \cup r_3\}$. This implies that r_1 , r_2 , and r_3 must each be found in at least a pair of two feasible $k - 1$ -trips differing by only one request. So we would find that the trip $\{r_1\}$ has two dependencies differing by one trip— $\{r_1, r_2\}$ and $\{r_1, r_3\}$ —and identify the trip $\{r_1, r_2, r_3\}$ as the union of those two dependencies.

For each union identified in this manner $k*(k-1)/2$ times—that is, having the precondition that all $k*(k-1)/2$ possible subcliques of size $k-2$ and their k parent cliques were identified—we call $\gamma(v^*, r)$, where v^* is an ideally located vehicle with no en-route riders, to check whether its delay times are potentially feasible. If so, the trip is added to \mathcal{T} and is also added as a dependency of each size $k - 1$ antecedent trip in D , a map of trips to “exclusion rules” based on clique membership.

Step 3: Trip-to-vehicle enumeration

We use the list of ideally feasible trips \mathcal{T} as a starting point to build each individual vehicle v 's feasible trips of size k from 2 to k^{MAX} . In this step (Algorithm 3 in the SI), we begin from the request-vehicle graph G_{RV} , which defines for each vehicle the individual requests that are (and are not) feasible. G_{RV} , in tandem with trip-to-trip exclusion mappings D , determines which trips of size $k = 2$ can be ruled out and need not be checked for feasibility with each v ; we proceed similarly for larger trips, looping over each vehicle in parallel until all trips $t \in \mathcal{T}$ have been ruled out or checked for feasibility. Each feasible matching $v \in \mathcal{V}$ to $t \in \mathcal{T}$ is added as an arc to the request-trip-vehicle graph G_{RTV} with additional arcs added linking t to its member requests as described in Alonso-Mora et al. (2017). (Alonso-Mora et al., 2017a)

Step 4: Option elimination

To make the routing assignment possible to solve, a necessary heuristic step reduces the number of vehicle-to-trip combinations considered. Each combination of vehicle v to trip t is retained that is either one of the n lowest-cost options for v to serve or one of the n lowest-cost options serving t , using the value of n shown in Table 3.2. The remaining combinations are pruned from the graph G_{RTV} . If n is sufficiently large, this step can drastically reduce the size of the problem, making it tractable without drastically altering the optimal solution.

We choose $n = 15$ so that all scenarios are tractable given our computational platform limits (48 hour model runtime limit and memory constraints).

Step 5: Routing assignment

Having constructed and pruned the graph G_{RTV} , we use a binary optimization to assign vehicles to trips. Table 3.3 describes our problem formulation, adapted from Alonso-Mora et al. (2017) as revised by Liu et al. (2019). (Alonso-Mora et al., 2017a; Liu et al., 2019) Equation 3.1a, the objective, minimizes total costs (sometimes including a Pigovian tax), including a large reward term to place priority on serving as many requests as feasible. Equation 3.1b ensures no vehicle is assigned more than one trip grouping, and Equation 3.1c ensures no request is included in multiple trip groupings. Finally, Equation 4.3e constrains assignments to zero or one.

Table 3.2: Optimization parameters

Parameter	Label	Base value	Sensitivity cases
Request batch time interval (min)	τ	5	1 min, 10 min, 800 rides
Fleet size (fraction of 30-min demand level)	$ \mathcal{V} $	0.6	-
Vehicle passenger capacity	o_v^{MAX}	2 for all cars	1, 4
Maximum trip chain size (including non-overlapping passengers)	k^{MAX}	10	-
Max. rider wait+detour time (min)	δ^{MAX}	15	TBD
Minimum options retained per trip and vehicle	-	50	TBD
Travel speed (mph)	-	15	-
Solver runtime stopping condition (min)	-	15	“real-time” runtime
Integer optimality gap stopping condition	-	0.5%	0.1%

Table 3.3: Vehicle assignment problem

$$\underset{a_{v,t}}{\text{minimize}} \sum_{(v,t) \in G_{\text{RTV}}} (\gamma(v,r) - \rho n_t) a_{v,t} \quad \text{Minimize costs including reward for served requests} \quad (3.1a)$$

subject to

$$\sum_{t \in G_{\text{RTV}}} a_{v,t} \leq 1 \quad \forall v \in G_{\text{RTV}} \quad \text{No vehicle assigned multiple trips} \quad (3.1b)$$

$$\sum_{t \in \mathcal{I}_r^R} \sum_{v \in \mathcal{I}_t^T} a_{v,t} \leq 1 \quad \forall r \in G_{\text{RTV}} \quad \text{No request included in multiple assigned trips} \quad (3.1c)$$

$$a_{v,t} \in \{0, 1\} \quad \forall (v,t) \in G_{\text{RTV}} \quad \text{Assignments are binary} \quad (3.1d)$$

To improve solution quality for a given runtime, we provide the solver a “warm start” using a heuristic described in Alonso-Mora et al. (2017) that ranks matchings by decreasing size and increasing cost then greedily assigns them. (Alonso-Mora et al., 2017a)

Step 6: Vehicle rebalancing

Following each optimization over a batch of new ride requests, there may be some requests that were not assigned a vehicle, and potentially some vehicles that were not assigned a trip. Possible reasons include shifting distribution of demand across a city over time (fleet vehicles not well-distributed for demand), an excess of demand, suboptimality due to heuristic pruning of vehicle-trip matching options, or suboptimality due to optimizer runtime limits. Furthermore, as demand shifts, the fleet may benefit from relocating to areas of upcoming demand even if there have not been many requests there yet. We apply the method described in Alonso-Mora et al. (2017) to rebalance vehicles to unserved areas. Adapted from Alonso-Mora to use total travel costs, Equation 3.4a (in Table 3.6) minimizes the total travel costs of rebalancing. Equation 3.4b requires that the total vehicle rebalanced equal the number of unserved trips, or the number of unused vehicles if that number is smaller. While Equation 3.4c constrains rebalancing assignments to zero or one, the constraint matrix for this formulation is totally unimodular, implying that it can be solved very quickly as a linear program and still yield an integer optimal solution. This optimization considers vehicle-rider matches that may have been pruned away for tractability by the primary optimization.

This rebalancing step is a defensible approximation of how fleets react to demand in real-time, but the model’s myopic nature is not well-suited for a more comprehensive analysis of different, realistic rebalancing strategies (Uber and Lyft have teams that develop sophisticated demand forecasting algorithms based on years of operational data).

As a sensitivity test of the importance of forecasting, we consider (as a second stage of rebalancing) two methods of repositioning vehicles based on upcoming demand. To model an optimistic version of effective forecasting, we assume the the fleet operator has perfect information regarding total volume of upcoming ride requests starting in each of Chicago’s 77 designated community areas, but that individual ride request locations are unknown until each current batch of ride requests is optimize over. For a sense of scale, each community

Table 3.4: Rebalancing for anticipated demand (efficiency-centric version)

minimize	$\sum_{v \in \mathcal{V}^U} \sum_{z \in \mathcal{Z}^U} c_{v,z} a_{v,z}$	Minimize routing costs	(3.2a)
subject to			
	$\gamma = \gamma^*$	Reward must be maximized	(3.2b)
	$\gamma = \sum_{v \in \mathcal{V}^U} \sum_{z \in \mathcal{Z}^U} \gamma_z$	Reward is defined as the sum across zones	(3.2c)
	$\gamma_z \leq \sum_{v \in \mathcal{V}^U} a_{v,z} \quad \forall z \in \mathcal{Z}^U$	Reward is earned by assigned cars	(3.2d)
	$\gamma_z \leq d_z \quad \forall z \in \mathcal{Z}^U$	No reward for cars above demand	(3.2e)
	$a_{v,z} \in \{0, 1\} \quad \forall v \in \mathcal{V}^U, z \in \mathcal{Z}^U$	Assignments are binary	(3.2f)

Table 3.5: Rebalancing for anticipated demand (equity-centric version)

minimize	$\sum_{v \in \mathcal{V}^U} \sum_{z \in \mathcal{Z}^U} c_{v,z} a_{v,z}$	Minimize routing costs	(3.3a)
subject to			
	$\gamma = \gamma^*$	Reward must be maximized	(3.3b)
	$\gamma \leq \gamma_z \quad \forall z \in \mathcal{Z}^U$	Reward is defined by the lowest-reward zone	(3.3c)
	$\gamma_z \leq \sum_{v \in \mathcal{V}^U} \frac{a_{v,z}}{d_z} \quad \forall z \in \mathcal{Z}^U$	Reward is earned by assigned cars	(3.3d)
	$\gamma_z \leq 1 \quad \forall z \in \mathcal{Z}^U$	No reward for cars above demand	(3.3e)
	$a_{v,z} \in \{0, 1\} \quad \forall v \in \mathcal{V}^U, z \in \mathcal{Z}^U$	Assignments are binary	(3.3f)

area on average consists of around 12 census tracts. As a basis for rebalancing decisions, the fleet uses each community area’s highest-volume 30-minute period within the upcoming 90 minutes. This step’s goal is not to perfectly match supply of vehicles to demand, but rather to ensure the fleet errs on the side of over-provisioning vehicles so that ridesplitting decisions are made for optimal-cost reasons rather than supply shortages.

In one version, shown in Table 3.4, we prioritize efficiency by assigning as many vehicles as possible to neighborhoods with unmet upcoming demand d_z (where d_z is demand above current supply of vehicles rather than absolute demand). As a secondary objective, the problem minimizes costs of routing specific vehicles to those locations. This formulation is solved as a hierarchical objective problem (to minimize costs subject to a maximized reward term) using Gurobi. We also consider a variant of this problem that prioritizes equity, shown in Table 3.5. In this variant, reward is defined by the lowest-reward neighborhood instead of a sum across neighborhoods, where each neighborhood’s reward is the percent of anticipated demand that will be matched with supply. This instead favors matching supply to demand similarly well across zones. These results are discussed in Section 3.4.4.

Table 3.6: Unmet trip rebalancing

$$\begin{array}{ll} \underset{a_{v,r}}{\text{minimize}} & \sum_{v \in \mathcal{V}^U} \sum_{r \in \mathcal{R}^U} \gamma_{v,r} a_{v,r} \quad \text{Minimize travel costs} \end{array} \quad (3.4a)$$

$$\begin{array}{ll} \text{subject to} & \sum_{v \in \mathcal{V}^U} \sum_{r \in \mathcal{R}^U} a_{v,r} = \min(|\mathcal{V}^U|, |\mathcal{R}^U|) \quad \text{Assign the lesser of unused vehicles or unmet trips} \end{array} \quad (3.4b)$$

$$a_{v,r} \in \{0, 1\} \quad \forall v \in \mathcal{V}^U, r \in \mathcal{R}^U \quad \text{Assignments are binary} \quad (3.4c)$$

3.3 Data for Chicago case study

Our model takes as input exogenous historical TNC trip demand for Chicago, Illinois, optimizing dispatch to minimize costs. Those trip and cost data sources are detailed here.

3.3.1 Cost data

Each test case is run twice to minimize two cost functions: (1) the private costs of operating the fleet and (2) those private costs plus a Pigovian tax on externalities. Table 3.7 lists the

Table 3.7: Cost components included in each test case.

Cost	Units	Included in case	
		No tax	Tax
Fuel	\$/vehicle-mile	✓	✓
Maintenance	\$/vehicle-mile	✓	✓
Value of wait and delay times	\$/passenger-minute	✓	✓
Value of privacy (lost in shared ride)	\$/shared pick-up	✓	✓
Pollutant and greenhouse gas emissions	\$/vehicle-mile		✓
Congestion	\$/vehicle-mile		✓
Collisions	\$/vehicle-mile		✓
Noise	\$/vehicle-mile		✓

cost components included in each case, which are explained in detail here. All values are expressed in 2019 US dollars.

Routing private costs: We include fuel and maintenance as a flat per-mile routing private cost. For fuel price per gallon, we take 2019 annual averages for the Chicago area (\$3.00) from Energy Information Administration data.(US Energy Information Administration, 2018b). For fuel efficiency, we take GREET’s default assumptions (24.9 miles/gallon) for a combustion vehicle of model year 2015, where 2015 is the use-weighted average age in 2019 of vehicles found in Chicago’s dataset of ridesourcing vehicles per our own analysis.(City of Chicago, 2021e) For per-mile maintenance costs, we reuse estimates for combustion vehicles (8.5¢/mile) taken from a recent comparison study of different vehicle powertrains; we exclude per-mile insurance costs due to little evidence of the costs increasing per marginal mile for ridesourcing drivers.(Compostella et al., 2020)

Wait and delay time private costs: We include as a per-passenger-minute private cost the costs of waiting (before pickup) and delays (after pickup, due to en-route detours to pick up or drop off other riders). Using US Department of Transportation’s recommended methodology, this cost is the total wait and delay time multiplied by 60% of the Chicago area’s mean wage, with adjustments made using 2017 National Household Travel Survey data to reflect the proportion of urban ride-hailing passenger-miles traveled that are personal versus business.(U.S. Department of Transportation, 2016; US Bureau of Labor Statistics,

2020)

Lost privacy private costs: We include in private costs the loss of an individual’s welfare resulting from sharing a ride with a stranger. Estimates are uncertain in the literature, so we use the median estimate across three studies (\$1.24) on a per-passenger basis.(Yan et al., 2019; Kang et al., 2021; Lavieri and Bhat, 2019) For example, if one passenger is already in a car and a second passenger enters, a cost of \$1.24 will be added twice to the trip (once per each passenger inconvenienced).

Emissions externalities: Emissions factors for tailpipe combustion and fuel refining are taken from Argonne National Laboratory’s GREET model.(Argonne National Laboratory, 2019) For emissions from vehicle manufacturing, disposal, and recycling, we take GREET’s default assumptions for a combustion vehicle of model year 2015 (the current average age of vehicles found in Chicago’s dataset of ridesourcing vehicles per our own analysis).(City of Chicago, 2021e) For tailpipe damage factors (\$/mile), we use tract-level data from the InMAP reduced complexity model (RCM) in the base case.(Tessum et al., 2017) Refining damage factors use InMAP county-level data and will be sited as the weighted average of operating refining capacities in the broad vicinity around Illinois (per current Energy Information Administration datasets). For greenhouse gas emissions, we adopt the \$50/tonne social cost of carbon estimated by the Interagency Working Group on the Social Cost of Carbon (United States Government, 2016).

Traffic externalities: A recent study of ridesourcing impacts conducted a thorough review of literature on traffic and other externalities resulting from ridesourcing.(Ward et al., 2021a) We reuse that study’s central estimates for collisions (9.0¢/mile), congestion (6.0¢/mile), and noise (0.02¢/mile) as flat per-mile externalities.

3.3.2 Travel data

Road network: For the Chicago area road network, we use a downloadable version of OpenStreetMap data for Illinois, in node and edge list formats.(Boeing, 2020) To reduce the size of the network, we treat all roads as two-way and only keep road segments that fall

inside the convex hull of origin and destination latitude-longitudes from rides included in the model. For tractability reasons, we exclude segments of the road network that are marked as residential, and keep all others (including highways, primary, secondary, and tertiary road designations).

Ridesourcing demand: Chicago publishes TNC trip data from November 2018-present and taxi trip data from 2013-present.(City of Chicago, 2021d,c) Each trip includes origin and destination location and time (anonymized to the census tract or community area and the nearest 15 minutes), as well as the trip’s recorded exact duration and travel distance.

Ride request times are not provided, so we treat ride start times as request times (and convert rounded times to specific times by sampling uniformly from $+/-7.5$ minutes within the rounded ride start time). To convert rounded locations to specific pickup points, we randomly sample an intersection from the road segments within that tract (or community area).

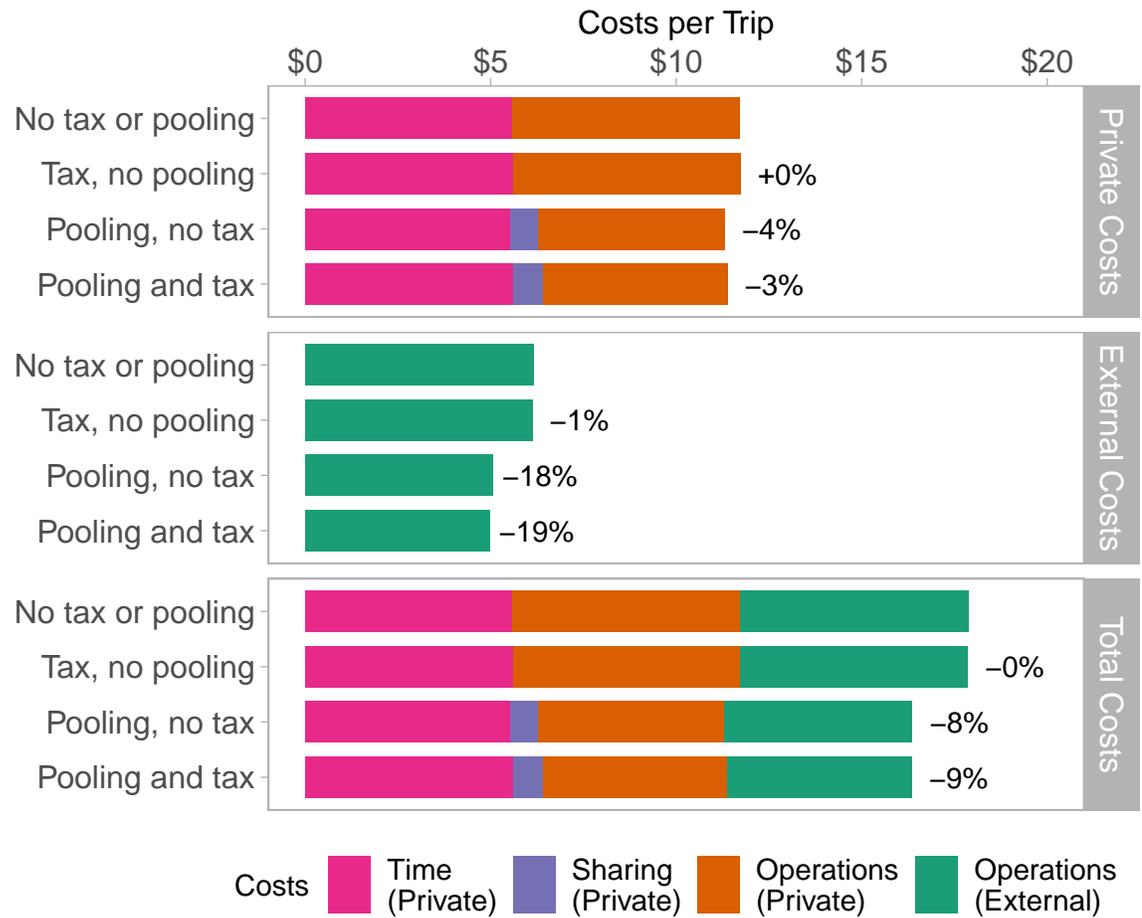
This dataset includes a flag for whether the user authorized ridesplitting (e.g., requested UberPool rather than a private ride) and whether the ride was ultimately pooled with another ridesplitting ride. However, we do not use this variable, which is treated as endogenous to our model. Instead, we treat all riders as open to ridesplitting and optimize those decisions for modeled costs and constraints.

For tractability, we select a representative week of ridesourcing demand to model. To do so, we use a two-step unsupervised learning approach that seeks to make sure not only that overall travel volumes are typical, but also that the underlying composition of travel patterns are typical. First, we use partitioning around medoids on location’s total ride request volume per hour of the week (using dynamic time warping distance, appropriate for time series data clustering) to determine four representative hourly profiles of Chicago ridesourcing demand. We repeat this using request origins and then request destinations as the basis for volume counts, yielding eight total profiles (four using origins and four using destinations). We use those eight profiles to create eight metrics for each week in 2019, corresponding to the total level of demand to or from locations fitting each of the eight profiles. Next, we use k -medoid clustering, with a $k = 1$, along those eight dimensions to

determine the week most representative of ridesourcing demand in Chicago in 2019. The chosen week is from September 9 to September 15.

3.4 Results and Discussion

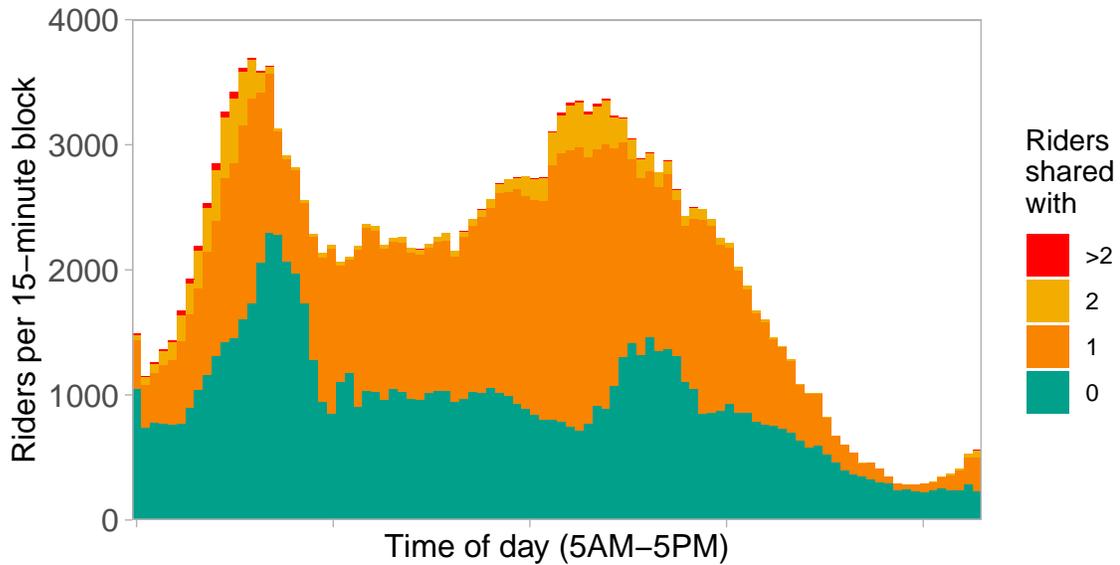
Figure 3.1: Private, external, and total costs associated with ridesourcing, with and without a Pigovian tax on external costs. Dollar values are summed across all phases of ridesourcing operations (including “deadheading” travel of empty vehicles) then divided by the total number of trips served.



3.4.1 “Status quo” use of ridesplitting

First, we consider how much a ridesourcing fleet may use the option of ridesplitting to minimize private costs of to the fleet (operations costs) and to travelers (travel time costs and sharing costs). As our model is currently calibrated, 53% of rides use ridesplitting (i.e.,

Figure 3.2: Time series view of riders per 15-minute block who have a private ride (“Shared” = 0) or who overlap for some duration with one or more other rides in the same car (“Shared” > 0). “Change” represents each ride’s amount of sharing under a Pigovian tax on external costs minus their amount of sharing in the case with no Pigovian tax.



the rider overlaps with another rider in the same car for at least some portion of their ride) in the private cost minimization case.

The option of ridesplitting itself leads to cost reductions, even without any tax on external costs. Figure 3.1 summarizes on a per-trip basis the private (top panel), external (middle panel), and total (bottom panel) costs associated with ridesourcing, with different combinations of the ridesplitting option and a Pigovian tax. In the status quo, no tax, case, external costs from ridesourcing (congestion, collisions, life cycle emissions, and noise) represent 33% of total costs. Another 33% arises from private costs of operation (gasoline, maintenance, and depreciation). 28% arises from travel waits or delays (difference between actual and best-case drop-off times), and the remaining 5% comes from the disutility associated with sharing a ride with a stranger.

Versus a baseline in which ridesplitting is strictly forbidden, adding the option of ridesplitting reduces private costs 4%, composed of a reduction in operations costs (lower vehicle-distance traveled) that is only partially offset by additional time delay costs and sharing costs. Due to lower vehicle-distance traveled, private cost-optimal ridesplitting itself also reduces external costs of traffic and air emissions by 18%. This suggests that

ridesplitting services do have a niche to fill in the broader picture of ridesourcing, regardless of the question of a Pigovian tax.

3.4.2 Changes induced by a Pigovian tax on externalities

Figure 3.1 summarizes on a per-trip basis the change in private, external, and total cost induced by a Pigovian tax on the external (unpriced) portion of those costs. The net effect of the Pigovian tax is small. On net, external costs are reduced by only 1%, due to more efficient routing and pooling. This increase is mostly offset by a slightly smaller increase in private costs.

In terms of citywide averages, the changes resulting from a Pigovian tax are minor. The proportion of rides that overlap at some point with at least one other ride in the same car rises slightly, from 42% to 44%. Due in part to this increase in ride pooling, total citywide vehicle-miles traveled (VMT) decline by 0.6%. Also due to increased pooling, the average ride’s wait plus delay time (defined as actual dropoff time minus best-case dropoff time, whether there was a wait prior to pickup or detours en route) increased very slightly, by 15 seconds (from a status quo delay time of 10 minutes and 30 seconds).

3.4.3 Variation in results across times and areas

With or without a tax, the use of ridesplitting varies across times and areas. For the first day in our representative week, Figure 3.2 illustrates how the use of ridesplitting varies with demand fluctuations. When demand is very low (seen in the early morning hours at the right-hand side of the plot), ridesplitting use is very low. Ridesplitting tends to be more common when demand is higher, but it is not strictly used for peak demand. Instead, during peak periods, both solo rides and ridesplitting increase. Further investigation is needed, but these trends may be due to nonlinearities in the viability of pooling: a certain critical mass of riders is needed along a corridor to enable any pooling. As a share of total demand, pooling “ramps up” to a peak during the late afternoon, while more private rides are used for the post-commute evening hours. It is conceivable this could reflect late afternoon commutes largely being concentrated from the central business district, but nightlife destinations being more geographically scattered (and thus harder to pool).

While the net change in outcomes due to a Pigovian tax is relatively small, larger changes exist in specific times and areas. Empirical analyses of Chicago have shown that ride pooling is as common or more common in outlying neighborhoods as it is downtown. Chi (2013) It appears that the use of ridesplitting, and the effects of a tax, may vary systematically across neighborhoods in this model’s results, but additional robustness checks are needed to determine which variations are robust across input assumptions and represent meaningful operational changes (which could have equity implications across neighborhoods), or are dependent on certain specific modeling parameters.

3.4.4 Sensitivity to rebalancing strategies

In the base case, we use the Alonso-Mora et al. (2017) method of rebalancing the fleet after each timestep, but minimizing our model’s objective (private costs with or without a tax on external costs). (Alonso-Mora et al., 2017a) To consider how much difference the use of a rebalancing strategy (or the specific priorities of rebalancing) may make, we consider two possible formulations of a rebalancing optimization that prioritize efficiency (Table 3.4) or equity (Table 3.5). Results to come.

3.4.5 Discussion

These results demonstrate that ridesplitting plays a valuable role within the set of services ridesourcing fleets offer. Even when minimizing private fleet and traveler costs, the option of ridesplitting lowers those private costs substantially relative to a 100% solo ride service.

However, the net effect of a Pigovian tax is limited, both in terms of operational changes (3% increase in pooling) and external cost reductions. This small net effect may be explained due to the relatively small number of opportunities to reduce externalities given fixed demand: each rider must be picked up, and the rides that can be pooled most efficiently (without leading to large detours and associated private and external costs) are already pooled to minimize private costs. Despite this small net effect, there is a large difference arising from who pays external costs. In the “no tax” case, traffic and emissions costs are borne by society at large (many of whom do not see use ridesourcing or see benefits from it) and do not factor into private decision-making. In the case of a tax, even though those costs

are still borne by society at large, the parties who opt into a ridesourcing transaction pay those costs (as a transfer to the government) and consider them in decision-making.

These results do not provide evidence of a substantial market failure regarding ridesplitting that would justify correction via a policy intervention on economic efficiency ground. However, there may be other sources of market failure (e.g., market concentration, imperfect information, or transaction frictions related to the complexity of ridesplitting) that lead to inefficient outcomes, or other grounds (e.g., equity) that would justify policy interventions to encourage more ridesplitting.

3.5 Data availability

Data and code supporting the findings of this study are available at

<https://github.com/mbbruch/>.

3.6 Acknowledgements

We thank Haojia Zuo for sharing a reimplementaion of the Alonso-Mora et al. (2017) model code, which served as a starting point for our model code development. This work was supported by the Carnegie Mellon Center for Climate and Energy Decision Making through a cooperative agreement (Awards SES-0949710 and SES-1463492) between the National Science Foundation and Carnegie Mellon University. This work was also developed as part of the Center for Clean Air Climate Solution (CACES), which was supported under Assistance Agreement No. R835873 awarded by the U.S. Environmental Protection Agency. It has not been formally reviewed by EPA. The views expressed in this document are solely those of authors and do not necessarily reflect those of the Agency. EPA does not endorse any products or commercial services mentioned in this publication. This project is funded in part by Carnegie Mellon University's Mobility21 National University Transportation Center, which is sponsored by the US Department of Transportation. This material is based upon work supported by Carnegie Mellon University's College of Engineering Moonshot Award for Autonomous Technologies for Livability and Sustainability (ATLAS). It was also supported by the Department of Engineering & Public Policy at Carnegie Mellon University. This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which

is supported by National Science Foundation grant number ACI-1548562. Specifically, it used the Bridges-2 system, which is supported by NSF award number ACI-1928147, at the Pittsburgh Supercomputing Center (PSC). (Towns et al., 2014; Nystrom et al., 2015)

CHAPTER IV

Effects of Chicago's congestion price on solo and shared ridesourcing use

This chapter was developed with coauthors Connor Forsythe, Jeremy Michalek, Kate Whitefoot, and Charlotte Andreasen.

My first two papers consider ridesourcing fleets that behave optimally and estimate the potential effect of changing incentives: whether internalizing unpriced externalities affects ridesourcing operations and whether changing the use of specific levers (electrification and ridesplitting) may improve policy outcomes. Instead of modeling optimal behavior under economically efficient price signals, this paper instead conducts an empirical analysis of a recently implemented change to ridesourcing incentives.

We consider the case of a per-ride ridesourcing fee change instituted in Chicago in January 2020 and its effects ridesourcing use and operations. The change (described in detail in Section 4.0.2) increased the fee associated with all rides and simultaneously discounted fees for pooled rides. Its stated goals were to discourage private TNC rides in favor of shared vehicles downtown, and in doing so, to reduce peak hour congestion where it is worst. (City of Chicago Business Affairs and Consumer Protection, 2020) The City of Chicago's own study found that annual TNC use increased 271% from 2015 to 2018, TNC cars occupy 26 miles of road space downtown during evening commute periods, and that higher proportions of TNC riders request private rides downtown relative to other parts of the city. (Chi, 2013) The effects of Chicago's ride fees warrant investigation on their own, but are also of interest because New York and San Francisco have instituted similar policies, (New York City Taxi & Limousine Commission; City and County of San Francisco Treasurer and tax Collector)

and understanding the efficacy of these policies may inform policymakers in all urban areas on the value of such fees. Furthermore, these TNC congestion charges may serve as a proof of concept for broader road use fees for other forms of private vehicle use. New York plans to implement a charge on all private vehicle use within certain neighborhoods in 2023, and Chicago has stated it plans to study a more comprehensive charge in the future.(Gold, 2021; Chi, 2013)

We will use a difference-in-differences causal inference framework considering how areas treated by the policy differed in outcomes from the rest of the city. The question we hope to address is: *Did Chicago's ridesourcing fee change accomplish the stated goals of discouraging private rides and encouraging shared rides downtown?*

4.0.1 Prior work

The impact of TNC use on congestion is the subject of much ongoing discussion and research. Qian et al. analyze vehicle trajectory data in New York City, and find that TNCs may be responsible for substantial increased congestion and reduced road speeds.(Qian et al., 2020) However, within the gray literature, a 2019 Fehr & Peers report analyzed private TNC use data alongside public road use data for for Chicago and six other cities, and found that TNC use in Cook County (where Chicago is located) makes up only 3% of total vehicle-distance traveled (VDT).(Balding et al., 2019) One study gathered data by serving as a driver for ridesourcing apps, estimating that at the individual trip level, ridesourcing use increases the traveler's resulting VMT by 83.5%, but they do not estimate the overall citywide impacts of TNC use.(Henao and Marshall, 2019) An across-city econometric study used macro-level data to reach conclusions on certain TNC impacts, but do not find significant results for VDT.(Ward et al., 2019) Li et al. (2021a) suggest that one source of difficulty in assessing TNC impacts on congestion is the correlation across time and space between use of TNCs and other private vehicles and argue that estimates of congestion using average VDT may not be correlated with more granular estimates.(Li et al., 2021b)

A separate branch of study considers the impact of pricing on TNC rider travel choices. Mehr et al. consider the inefficiencies arising for various possible network equilibria, and find that differential pricing at the road link level—taking into account dependencies of road

capacity on autonomous vs. human drivers—can lead to minimal overall delays.(Mehr and Horowitz, 2019) Li et al. (2021b) formulate an optimal pricing problem considering network equilibrium under the scenarios of a downtown zone entry fee, a fee upon entry and exit, and a trip-based fee, finding that an entry fee most effectively reduces congestion but that a trip-based fee most effectively raises revenue. Li et al. (2021a) apply a market equilibrium model to the case study of San Francisco, finding that a time-based charge (that considers occupied and idle time) leads to higher vehicle occupancy and tax revenue relative to a trip-based charge. Simoni et al. use agent-based modeling to estimate that while various road pricing schemes can effectively reduce congestion, including from shared vehicles, their social welfare impacts differ (with more complex schemes often performing better).(Simoni et al., 2019)

There is suggestive evidence in the gray literature of the impact of congestion fees. As one recent example, a consultancy commissioned by the New York City Metropolitan Transportation Authority recently projected that the city’s planned congestion pricing on all vehicle travel may reduce VDT by 7-9%.(Colon, 2021)

This work differs from these studies in that it investigates the effects that resulted from a ridesourcing congestion charge, using empirical analysis of historical data within a city, rather than deriving results from a more theoretical model. In doing so, it considers an implemented congestion charge and its realized impacts, rather than considering hypothetical or optimal charges and their impacts on optimal or assumed traveler decision-making processes.

4.0.2 Policy defining treatment

Chicago’s modified TNC ride fees were enacted on January 6, 2020. The fees assessed before and after this change are shown in Table 4.1. The change increased per-ride fees citywide and at all times, but also created an offsetting discount for riders who opt in to shared ridesplitting services. For rides beginning or terminating within a defined downtown zone between the “peak hours” of 6AM-10PM, the per-ride fee is larger, but the discount for opting into shared ridesplitting is also larger.

The per-ride fees shown in Table 4.1 can alternately be expressed as a set of stacking charges and discounts:

Table 4.1: Summary of Chicago’s January 6, 2020 TNC ride fees changes, which raised fees, made them time-varying, and introduced a discount for opting into shared ride pooling.

Location (start or end)	Time	Old fee	New fee	
			Solo ride	Ridesplitting
Special locations	Peak	\$5.72	\$8.00	\$6.25
	Off-peak	\$5.72	\$6.25	\$5.65
Downtown, excluding special locations	Peak	\$0.72	\$3.00	\$1.25
	Off-peak	\$0.72	\$1.25	\$0.65
All other locations	Peak	\$0.72	\$1.25	\$0.65
	Off-peak	\$0.72	\$1.25	\$0.65

1. **Base fee (time- and location-invariant):** the fee for all trips begins at \$0.72.
2. **Special location charge (time-invariant):** Trips to/from airports, the convention center, or the Navy Pier pay \$5.00 more.
3. **New fee schedule (location-invariant):** All trips post-1/6/2020 on pay \$0.53 more than before 1/6 and are discounted \$0.60 for opting in to shared ride pooling (i.e., rider selects UberPool, Lyft Line, or Via).
4. **Downtown zone surcharge (treatment):** Peak hour trips to/from downtown/special locations, on weekdays from 6AM-10PM, post-1/6/2020, pay \$1.75 more and are discounted \$1.15 for opting in to shared ride pooling.

Only surcharge #4, the downtown zone surcharge, defines a time period in which certain travel corridors are charged differently. For this reason, the difference-in-differences model described below only identifies the effect of that change—that is, the post-1/6/2020 differential between downtown or special location travel corridors and all other travel corridors. We consider two different treatment groups—ride origin-destination pairs (ODs) whose locations fall under the downtown zone surcharge (item #4 above), and those whose locations do not—in order to identify the effect of the downtown zone surcharge.

4.1 Materials and methods

4.1.1 Modeling strategy

For OD travel corridors i and dates t , we represent the downtown zone surcharge (post-treatment fee differential between treated and un-treated ODs) as a binary treatment indicator d_{it} (equivalent to the interaction of a binary indicator for treatment group and a binary indicator for dates on or after January 6, 2020) and each outcome (described below) as y_{it} :

$$y_{it} = \pi d_{it} + \delta_i + \tau_t + \epsilon_{it} \quad (4.1)$$

Fixed effects are included for i (δ_i) and t (τ_t) to capture unobserved factors that influence outcomes for each travel corridor, and for each date. Date fixed effects also capture the effect of all non-location-specific changes the policy induces, leaving π to capture the location-specific, time-specific treatment effect of the downtown zone surcharge. This formulation is augmented to consider unobserved factors that lead to autocorrelated outcomes between each corridor i and its neighboring corridors j :

$$y_{it} = \pi d_{it} + \rho \sum_{j=1}^N w_{ij} y_{jt} + \delta_i + \tau_t + \epsilon_{it} \quad (4.2)$$

This formulation follows the method used in other spatial autoregressive difference-in-differences studies. (Dubé et al., 2014; Chagas et al., 2016; Basu and Ferreira, 2021; Qiu and Tong, 2021) Neighbors' autocorrelative relationships are assumed to be time-invariant, defined by entries w_{ij} in a precomputed spatial weights matrix, and a single ρ is estimated to capture the change in outcomes arising from spatial autocorrelations. The details of our formulation of spatially autocorrelated outcomes are in Section 4.1.1.

If the downtown zone surcharge's treatment effect π is statistically significant, it identifies how much change treatment induced on outcomes of interest. To be precise, the question we can answer is: *How were outcomes associated with weekday, daytime, downtown OD travel corridors from January 6-March 3, 2020 affected on average by the downtown zone surcharge?*

We use this specification for three outcomes, each using ridesourcing ride data aggregated

at the level of OD o and date t :

- **Outcome 1: total (solo plus ridesplitting) ride volumes.** Daily ridership counts measure the number of private plus shared ridesourcing trips (i.e., the total number of transactions served by Uber and Lyft). This assesses whether the downtown zone surcharge affected demand in total along treated corridors. We model this outcome in log space, so that we can estimate the surcharge’s effect in percentage terms.
- **Outcome 2: percent of total rides that allowed ridesplitting.** To assess the downtown zone surcharge’s efficacy at encouraging ridesplitting, we will consider the proportion of total rides for which the user accepted a discount to allow other riders to potentially share the vehicle. We consider this outcome in level space (from 0% to 100% and estimate the effect as a percentage point change.
- **Outcome 3: percent of total rides that were matched with another ridesplitting ride:** Even if the downtown zone surcharge encouraged riders to opt in for ride sharing, there is no guarantee more rides were actually matched together. We will consider the proportion of rides that were shared with another rider. We consider this outcome in level space (from 0% to 100% and estimate the effect as a percentage point change.

We hypothesize that a downtown zone surcharge reduces total ride volumes (Outcome 1) increases the share of ridesplitting rides (Outcome 2), and increases the share of rides that are successfully matched with another ridesplitting ride (Outcome 3). It is not clear *a priori* whether we might expect the magnitude of shift for Outcome 2 or 3 to be larger. It is possible that network effects of additional shared ride opt-ins could have a better-than-linear impact on match rates. However, it is also possible that if new opt-ins are not successfully matched as often, the increase in match rates may be positive but still smaller than opt-in rates.

Figure 4.1 shows pre- and post-treatment trends for each of these three outcome variables. For each variable, trends prior to January 6, 2020 appeared relatively similar for travel corridors that were later subject to / not subject to the downtown zone surcharge. In the

case of daily TNC rides, there is substantial noise when looking at pre-treatment trends; even though some holidays are excluded (as described in Section 4.1.2) the largest outliers for daily rides occurs during the week of Thanksgiving, Christmas, or July 4th. Another notable outlier is at the end of January 2019, when a record-setting cold snap occurred in Chicago. In the case of rides for which pooling was allowed, and for which pooling matches occurred, there was a pronounced downward trend in ride sharing over time, but the trend appears similar for the two groups.

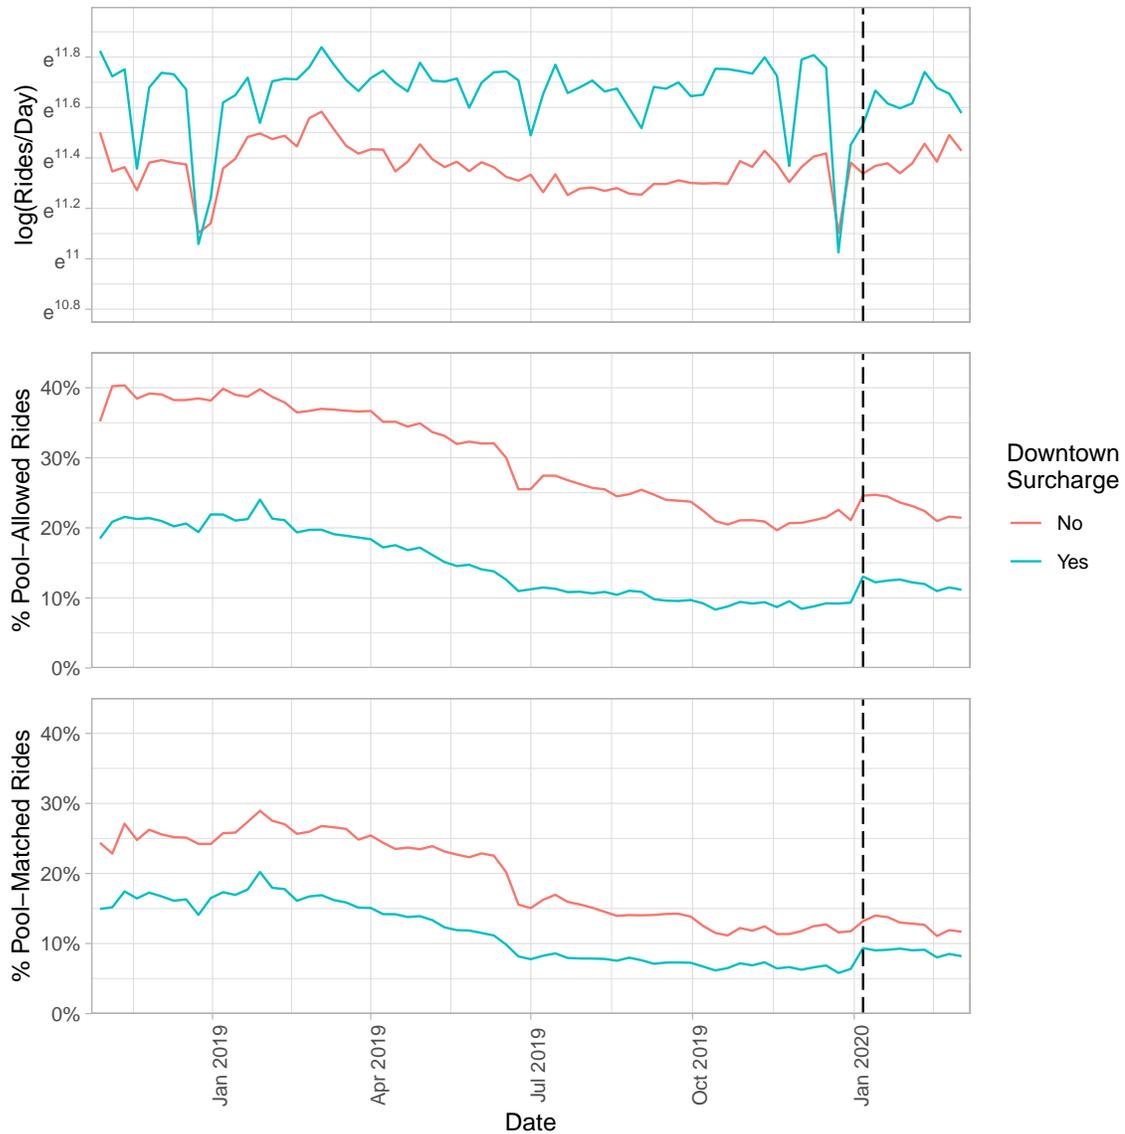
Handling of spatial autocorrelation

Our motivation for estimating a spatial autocorrelation effect, ρ , is twofold. First, if significant autocorrelations exist but they are not explicitly modeled, model coefficients estimated using ordinary least squares may be biased.(LeSage and P., 2008) Second, explicitly modeling ρ allows us to estimate both the direct effects of treatment on the treated travel corridors, as well as the indirect “spillover” effects of treatment on the un-treated travel corridors.(Dubé et al., 2014; Chagas et al., 2016)

Spatial weights matrices are often used to model those autocorrelations, including in a panel data context.(Elhorst, 2014; Salima et al., 2018; Baltagi et al., 2015) However, Lee & Yu (2010) demonstrate that when fixed effects and an autocorrelation term are both included, OLS estimates are inconsistent and a correction method must be used.(fei Lee and Yu, 2010) We model spatial autocorrelation using the R package ‘splm’, which implements the correction recommended by Lee & Yu.(Millo and Piras, 2012; fei Lee and Yu, 2010)

In order to define spatial weights for each travel corridors’ neighboring corridors, we use the measure of merge distance, which measures the relative length of the “supertrajectory” combining two origin-destination pairs into one shared ride. Aside from its property of directly measuring the ease of joining two travel corridors in a shared ride scheme—of direct relevance to spatial autocorrelations in ridesplitting opt-in and matching—it also captures endpoint similarity. For example, a reduction in TNC rides to the airport may reduce the number of free TNC vehicles at the airport, leading (through increased wait times and prices) to fewer rides starting at the airport. Because the first travel corridor terminates at the next one’s starting point—the airport—they would have a merge distance of zero, reflecting their

Figure 4.1: Pre- and post-treatment trends for three outcome variables: daily TNC rides, the percentage of TNC rides for which the rider opted into a shared ride, and the percentage of TNC rides for which the car was shared by two or more pooled ride requests for some period of time. Outcomes are shown as weekly averages, and the treatment date (January 6, 2020) is marked with a dashed line.



strong spatial relationship. Merge distance is formally described in Su et al. (2020). (Su et al., 2020)

4.1.2 Data

The City of Chicago publishes TNC trip data from November 2018-present and taxi trip data from 2013-present through its Open Data Portal. (City of Chicago, 2021d) Each trip includes origin and destination location and time (anonymized to the census tract or community area and the nearest 15 minutes), as well as the trip’s recorded exact duration and travel distance. Each ride’s price is broken out into fare, tip, and additional charges. An additional binary flag indicates whether a private or pooled ride was requested and, if pooling occurred, how many total passengers were pooled in the chain of rides (measured as total passengers from when the vehicle was empty until it was once again empty, not as the simultaneous passengers in the car at any given moment). To define the outcome of whether a ride was matched in a shared ride pool, we convert this integer-valued number of passengers into a zero (no ride pool matched) or one (rider matched with at least one overlapping ride).

Handling of censored origins and destinations

To protect rider privacy, Chicago does not provide true origin-destination (OD) locations, but instead rounds each origin and destination to the centroid of its census tract (henceforth “tract”, of which there are roughly 800 in Chicago). For tracts with fewer than three trips in a 15-minute time window, the tract is censored and location (for those time windows only) is rounded to the Community Area (henceforth “area”, of which there are 77 in Chicago). (City of Chicago Open Data Portal Team, 2019) This means that travel along any given OD may be reported at the tract or area level, depending on how high-volume the OD is at that time.

To create consistent observational units across dates, we must determine whether to define each OD as area-area, area-tract, tract-area, or tract-tract. Using areas eliminates the problem of missing (censored) data at the tract level and thus maximizes the number of trips we can assign to an observational unit. However, using areas also reduces our ability to ascertain which observational units were subject to the downtown zone surcharge treatment because many community areas span the downtown and non-downtown zones. On the other

hand, using tracts allows us to ascertain treatment group (except for 6 census tracts that partially overlap the downtown zone boundary), but increases the number of trips that cannot be assigned an observational unit due to missing geographical labels. In order to include a trip in our model, we must be able to both assign it to an observational unit and also be able to assign its observational unit to a treatment group.

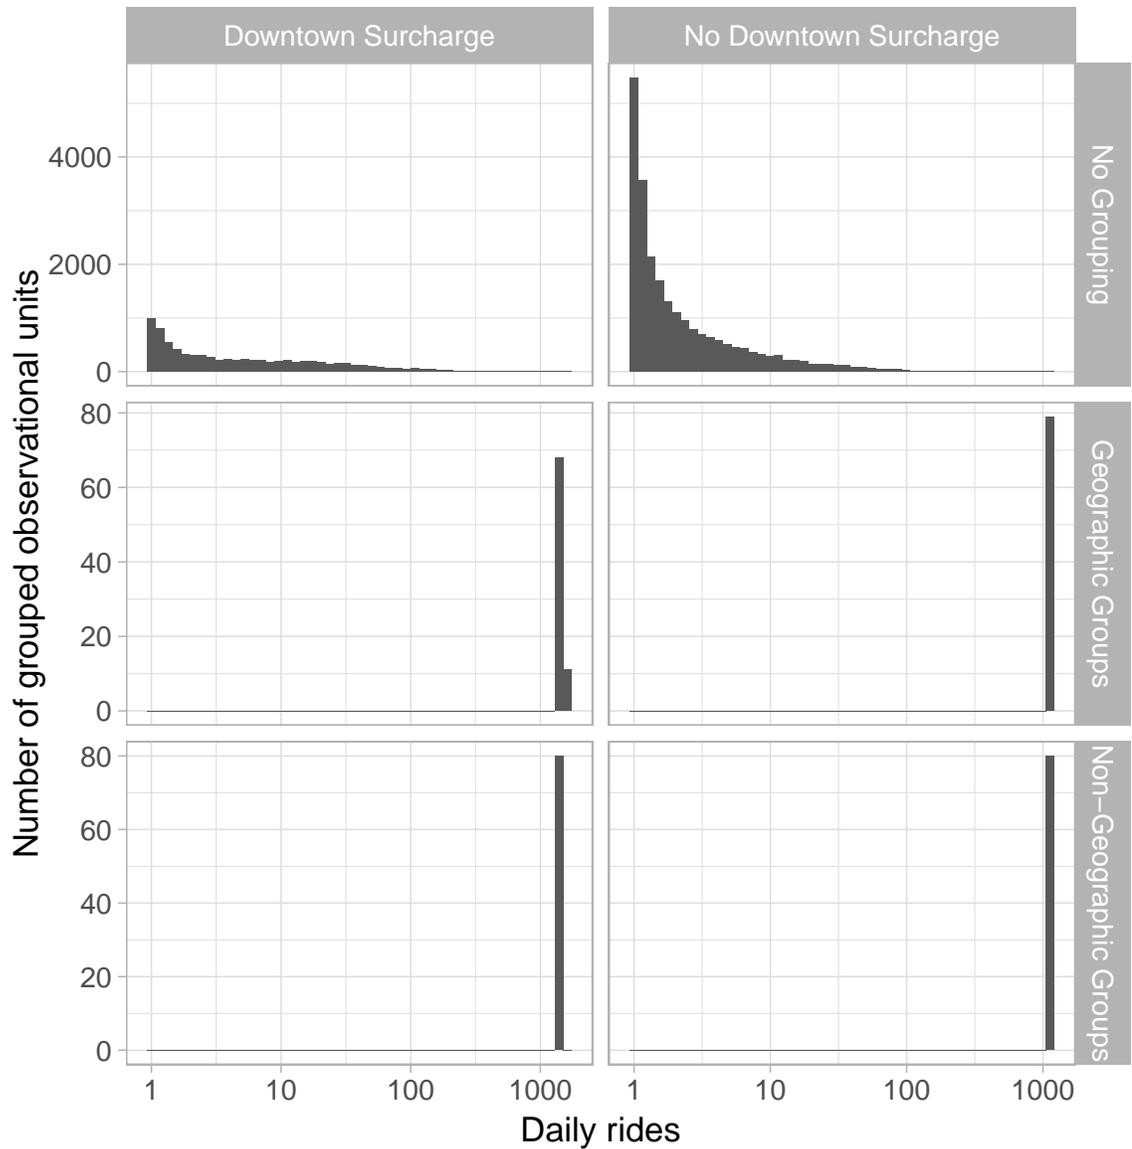
To maximize the number of trips we can include in the model, we define observational units using origin and destination area by default, but split origin and/or destination areas into tracts everywhere that doing so increases the number of trips we can both assign an observational unit (i.e., location is not missing) and assign treatment (i.e., observational unit does not overlap the downtown zone boundary). This represents a globally optimal number of trips we can assign to a treatment group, still yields some trips we cannot definitively label. We exclude those trips (2.7% of total).

Creation of similarly-sized observational units

Because ridesourcing demand is very unevenly distributed across the city, the distribution of trips is weighted very heavily towards a small number of ODs (particularly between different parts of downtown) with very high volumes. Those high-volume ODs are much more impactful in terms of absolute travel volume and policy outcomes than others, so the treatment effect as estimated across observational units may not be equivalent to the true average treatment effect per unit of travel downtown.

To make our average treatment effect estimate more equivalent to the average effect per unit of travel volume downtown, we combine low-volume entities into larger, similar-sized entities that preserve geographic similarity. The results of this grouping process are shown in Figure 4.2. We do this using a two-step clustering method. First, we use the R package ‘WeightedCluster’ to determine a set of k origin-destination pairs, weighted on travel volumes, that constitute the optimal k medoids of same-sized observational units. (Reynolds et al., 2006; Studer, 2013) Having done this, we adapt an approach that ranks ODs by gap between best and worst cluster assignment (distance to nearest medoid minus distance to farthest medoid) and iteratively assigns observational units to a medoid grouping to maximize size similarity of the k groups. (Schubert et al., 2020)

Figure 4.2: Distribution (log scale) of ride volumes for each travel corridor. With no grouping step, ride volumes vary drastically across observational units, which may lead to misleading effect estimates. Grouping by geographic proximity (main model specification) nearly matches ride volumes across all units. With geography-agnostic grouping (alternative specification), sizes are equal.



This does not create perfectly identically-sized groupings, but it drastically narrows the distribution of group sizes and makes it roughly normally distributed, rather than skewed.

Exclusion of certain dates

Outside of the hours during which the downtown zone surcharge applies, there is no location-varying treatment to identify the effect of. For that reason, we exclude trips that started and ended outside of 6AM and 10PM and exclude weekends. We exclude data beginning on March 7, 2020, a conservative estimate of when community spread of COVID-19 began in Chicago. This leaves us with a pre-treatment period from November 2, 2018 to January 3, 2020 and a post-treatment period from January 6 to March 6, 2020. Due to atypical demand patterns, we also exclude the following holidays when they fall on weekdays: winter holidays (December 23–January 2), Thanksgiving (+/−1 day), Memorial Day weekend (Friday–Tuesday), July 4 (+/−1 day), Valentine’s Day, and St. Patrick’s Day.

4.2 Results

Table 4.2 summarizes regression findings from our base specification. Relative to travel corridors not subject to the downtown zone surcharge, the surcharge reduced ridesourcing ridership, increased shared ride pooling opt-in rates, and increased the share of riders matched into a shared ride pool. For each outcome, we also find substantial effects of spatial autocorrelation, implying either that policy effects spill over into travel corridors that were not subject to the downtown zone surcharge or simply that co-occurring factors influence proximal travel corridors.

For each outcome, both the direct treatment and spatial autocorrelation effects are significant (with un-clustered standard errors). Neither root mean squared error (RMSE) nor R^2 are of primary importance, since the goal of this model is not to predict y_{it} accurately, but they do provide suggestive evidence of model fit.

4.2.1 Outcome 1: total TNC ride volumes

Our base specification estimates that the surcharge’s direct impact on travel corridors it covered was to reduce total (private plus shared) ridesourcing ridership by a mean [95%

confidence interval] of 8.2% [7.6%, 8.7%]. In absolute terms, this direct effect equates to a reduction of around 9,500 TNC rides to or from treated zones per day.

The model estimates substantial spatial autocorrelation, ρ , of 0.75 [0.73, 0.76]. This coefficient is open to interpretation. If the main underlying dynamic are spatial causal dynamics (e.g., increased ride volumes along one corridor leads to additional supply, lower costs, and additional ride volume for spatially related corridors), it may imply that on for the average travel corridor, a 10% reduction in neighboring corridors' ride volumes itself reduces that corridor's volumes 7.5% (separate from any direct effect of treatment). However, it may also simply capture unobserved "common cause" factors leading to similar outcomes in geographically related areas.

4.2.2 Outcome 2: percent of riders opting in to pooled ride services

The estimated direct effect of the surcharge on travel corridors it covered was to increase the rate of opt-ins to pooled ride services by an estimated 4.2 [4.0, 4.3] percentage points. If we were to take this result in tandem with the estimated drop in total rides, it would provide suggestive evidence of around 4,100 new ride pooling opt-ins per day to or from treated zones. Substantial spatial correlation, $\rho = 0.93$ [0.93, 0.94], exists for pooling opt-in rates, suggesting strong spillover effects or unobserved common factors.

4.2.3 Outcome 3: percent of riders matched into a shared ride pool

The estimated direct effect of the surcharge on travel corridors it covered was to increase the rate of pooled ride matches by an estimated 3.4 [3.3,3.5] percentage points. If we were to take this result in tandem with the estimated drop in total rides, it would provide suggestive evidence of around 3,800 new pooled rides per day to or from treated zones. This increase suggests that some additional rides were actually pooled as a result of increased opt-ins to pooled ride services. Again, spatial autocorrelation also has a significant and large coefficient, $\rho = 0.95$ [0.95,0.95].

4.2.4 Sensitivity to handling of geographic relationships

Our main specification groups observational units by geographic proximity and includes a spatial autocorrelation matrix, which relies upon design choices for the spatial weight

Table 4.2: Results for models that group origin-destinations by geographic proximity and include a spatial autocorrelation term using merge distance (main specification).

Dependent Variables: Model:	log(Rides) (1)	% Pooling Opt-Ins (2)	% Pool-Matched (3)
<i>Variables</i>			
Direct treatment effect	-0.0818*** (0.0027)	0.0416** (0.0006)	0.0342*** (0.0005)
Spatial autocorrelation	0.7476*** (0.0088)	0.9339*** (0.0019)	0.9485*** (0.0015)
<i>Number of fixed effects</i>			
Date	321	321	321
Corridor	297	297	297
<i>Fit statistics</i>			
Observations	95,337	95,337	95,337
Adjusted R ²	0.188	0.912	0.850
RMSE	0.138	0.029	0.027

Un-clustered standard errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

matrix. To test the sensitivity of our findings to these design choices, we compare results to an alternative model specification that does not consider geographic relationships within, or between, observational units. This classical difference-in-differences formulation (Equation 4.1 does not include a spatial autocorrelation term, and its approach to same-sizing uses a location-agnostic number partitioning method (Table 4.3). One additional difference in this model is that it uses two-way clustered standard errors, a more conservative approach, and uses an ordinary least squares estimator.

Table 4.3: Multi-way partitioning problem for same-volume clusters

$$\begin{array}{ll} \underset{a_{i,k}, s_k}{\text{minimize}} & s_K - s_1 & \text{Minimize range of cluster sizes} & (4.3a) \end{array}$$

subject to

$$\sum_{k=1}^K a_{i,k} = 1 \quad \forall i \in 1 \dots I \quad \text{Each O-D assigned to exactly one cluster} \quad (4.3b)$$

$$\sum_{i=1}^I v_i a_{i,k} = s_k \quad \forall k \in 1 \dots K \quad \text{Cluster size is sum of member O-D travel volumes} \quad (4.3c)$$

$$s_k \leq s_{k+1} \quad \forall k \in 1 \dots K - 1 \quad \text{Symmetry breaking constraint} \quad (4.3d)$$

$$a_{i,k} \in \{0, 1\} \quad \forall i \in 1 \dots I, k \in 1 \dots K \quad \text{Assignments are binary} \quad (4.3e)$$

These results are shown in Table 4.4. Neglecting spatial relationships does not drastically alter findings. Across travel corridor subject to the downtown zone surcharge, the surcharge is estimated to reduce ride volumes by 9.6% [7.1%,12.1%] relative to corridors with no downtown zone surcharge. This is similar in size to the direct treatment effect estimated by the spatial model, though the interpretation is different: here, it reflects the estimated total effect of treatment on treated units, rather than only the direct portion of those treatments.

Findings are similar for outcomes related to ride pooling. For the choice to allow pooling, this specification estimates treatment increases opt-in rates by 5.1 [4.6,5.6] percentage points and increases pooling matches by 4.4 [4.0,4.8] percentage points. As with ride volumes, these (slightly higher) effect estimates are not directly comparable to the spatial model, since they reflect total treatment effects rather than only the direct effects of treatment on treated entities.

As a robustness check, we also consider an alternative specification for the spatial weights matrix in which distances are defined using the simpler metric of origin-to-origin proximity; this may capture any “common cause” factors that might lead to co-occurring outcomes for travel corridors with similar starting locations. These results (Table 4.5) have a slight attenuation in treatment effect estimates, but do not drastically differ.

4.2.5 Discussion and future work

Across the main and alternative model specifications, we find that Chicago’s surcharge on private TNC rides downtown reduced total TNC ridership, increased ride pooling opt-in rates, and the rate of successful pooling matches. This suggests that a location-based surcharge can alter traveler behavior and meet policy goals.

These findings must be taken in context. We do not assess the impact of the full set of fee changes that were made in January 2020 (which included some increases across the board). Rather, we only consider the effects of the pricing differential for corridors levied the downtown zone surcharge, active from 6AM-10PM on weekdays. We also only consider impacts on TNC use on average along surcharge and no-surcharge travel corridors, and not how those impacts may vary from neighborhood to neighborhood or across the course of the day. This view of average impacts does not consider whether ridesourcing use is impacted

Table 4.4: Results for models that group origin-destinations with no consideration of geographic proximity and do not include a spatial autocorrelation term.

Dependent Variables: Model:	log(Rides) (1)	% Pooling Opt-Ins (2)	% Pool-Matched (3)
<i>Variables</i>			
Treatment	-0.0959*** (0.0128)	0.0508*** (0.0025)	0.0444*** (0.0020)
<i>Number of fixed effects</i>			
Date	321	321	321
Corridor	280	280	280
<i>Fit statistics</i>			
Observations	89,880	89,880	89,880
Adjusted R ²	0.925	0.951	0.925
Within R ²	0.032	0.121	0.124
RMSE	0.090	0.023	0.020
<i>Two-way clustered (Date & OD) standard errors in parentheses</i>			
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>			

Table 4.5: Results for models that group origin-destinations by geographic proximity and include an alternative spatial autocorrelation term (average origin-to-origin distance).

Dependent Variables: Model:	log(Rides) (1)	% Pooling Opt-Ins (2)	% Pool-Matched (3)
<i>Variables</i>			
Direct treatment effect	-0.07033*** (0.0025)	0.0411** (0.0006)	0.0343*** (0.0005)
Spatial autocorrelation	0.7476*** (0.0088)	0.8154*** (0.0037)	0.8296*** (0.0035)
<i>Number of fixed effects</i>			
Date	321	321	321
Corridor	297	297	297
<i>Fit statistics</i>			
Observations	95,337	95,337	95,337
Adjusted R ²	0.144	0.907	0.852
RMSE	0.240	0.028	0.026
<i>Un-clustered standard errors in parentheses</i>			
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>			

during the hours or on the streets where congestion is most severe.

The findings are also limited somewhat in external validity. They describe the fee's effects within the post-treatment period of roughly eight weeks, until community spread of COVID began. It is conceivable that if ride fees had been instituted at a different time of year (or, for that matter, during a pandemic), their first eight weeks' effect may have been different. For similar reasons, we cannot say whether these early impacts would increase, remain steady, or dissipate over time. They also are limited in scope to Chicago, and it is unclear whether a similar policy would have similar effects in different cities with different patterns of ridesourcing use and different traveler preferences, priorities, and budgets.

Available data only includes rides that were taken, which limits the findings of this analysis. For instance, we cannot say whether a reduction in ridership was caused by fewer users using ridesourcing apps post-treatment or by users opening an app, seeing higher prices, and choosing not to use the app). Importantly, we also cannot directly assess whether a relative reduction in riders downtown and a relative increase in shared rides led immediately to fewer TNC vehicles on the road, or whether the same vehicles operated at lower utilization (thus leading to many of the same impacts on the city). Further analysis could use the city's congestion data to assess how congestion was altered by the downtown zone surcharge.

While this study considers the efficacy of a surcharge on private TNC rides, it is worth bearing in mind that the congestion impacts of TNC riders are small in comparison to overall vehicle use. By one estimate, passenger-free "deadheading"—not directly penalized by Chicago's congestion surcharge—makes up 45% of total TNC travel, and the two combined only make up 3% of road use in Chicago. (Balding et al., 2019) The early stages of the downtown zone surcharge provide evidence that individual travel choices are affected by price signals, suggesting the question of taxing broader vehicle use is worth continued investigation and experimentation.

4.3 Data availability

Data and code supporting the findings of this study are available at

<https://github.com/mbbruch/>.

4.4 Acknowledgements

The authors thank Nick Muller, Inês Azevedo, Kara Kockelman, and Akshaya Jha for their input on model development. This work was developed as part of the Center for Clean Air Climate Solution (CACES), which was supported under Assistance Agreement No. R835873 awarded by the U.S. Environmental Protection Agency. It has not been formally reviewed by EPA. The views expressed in this document are solely those of authors and do not necessarily reflect those of the Agency. EPA does not endorse any products or commercial services mentioned in this publication. This project is funded in part by Carnegie Mellon University's Mobility21 National University Transportation Center, which is sponsored by the US Department of Transportation. This material is based upon work supported by Carnegie Mellon University's College of Engineering Moonshot Award for Autonomous Technologies for Livability and Sustainability (ATLAS). It was also supported by the Department of Engineering & Public Policy at Carnegie Mellon University. This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number ACI-1548562. Specifically, it used the Bridges-2 system, which is supported by NSF award number ACI-1928147, at the Pittsburgh Supercomputing Center (PSC). (Towns et al., 2014; Nystrom et al., 2015)

CHAPTER V

Conclusions and future work

Across three studies, this dissertation explores external costs of ridesourcing fleets, policy “levers” to address the associated market failures, and pathways (technological, operational, and behavioral) to improve outcomes. The synthesized findings of these studies, and potential research paths forward, are discussed in turn.

5.1 Summary of findings

The growth of electrification and shared mobility seem likely to continue. It seems conceivable that their impacts on society will be transformative, but the size and direction of those impacts remain subject to debate. This uncertainty is due in part to the ever-changing nature of urban travel patterns themselves and also a result of the lack of observational travel data in nearly all cities. All three studies seek to address these knowledge shortcomings. The first and second studies use optimization to compare optimally-chosen ridesourcing strategies with and without a Pigovian tax on the service’s external costs in order to identify the size of market failure. The size of this market failure suggests the size of socially suboptimal, distorted outcomes that may justify some form of policy intervention. The third focuses on Chicago, (the city with the most comprehensive public collection of observational travel data) and contributes new insights on the effectiveness of price-based regulations to affect travel behavior.

5.1.1 First study’s findings

The first study focuses on electrification and air emissions, finding that externalities induce a substantial (and in some cases drastic) shift towards vehicle electrification and a similarly

large reduction in air emission externalities. The size of this shift, and its robustness across analysis regions and sensitivity cases, implies that the market failure of unpriced emission externalities is large enough to justify policy intervention on economic efficiency grounds, and that electrification is a viable lever to reduce those externalities. However, the shift does not result in a homogeneous, 100% battery electric vehicle fleet. In each case, the socially optimal fleet is a mixture of technologies that depends on factors that vary with location. This implies cost-centric or goal-centric policies may be more appropriate than technology mandates, since EVs make more sense in some contexts than others.

This study holds demand fixed and minimizes costs of meeting that demand. However, if ridesourcing emissions costs were taxed, some reduction in ridesourcing demand would likely result in addition to, or instead of, vehicle electrification. This modeling assumption's effect on top-level findings is unclear. It is conceivable, for example, that more price-sensitive trips occur in the times or locations that tend to be more amenable to electric vehicles, and that an emissions tax would shrink the niche that electric vehicles are best equipped to fill. However, it is also conceivable that more trips an optimal fleet serves with conventional vehicles (e.g., trips during peak demand times) would be affected by pricing changes, thus further increasing the proportion of trips best served by electric vehicles. In either of these situations, the implications would depend on what those riders do instead of a ridesourcing trip (e.g., forego travel, walk, take transit, or use a personal vehicle).

The first study also demonstrates that economically efficient ridesourcing electrification alone will not yield the “deep decarbonization” some believe is needed to meet stated climate goals. In other words, efficient pricing is a low-regret policy but its outcomes may be, in some policymakers' eyes, far from sufficient. (Alternately, a technology mandate or a higher-than-economically-efficient nudge towards electrification may be warranted if policymakers prioritize robustness to the tail risk of catastrophic climate change over economic efficiency, or for pragmatic reasons of political viability versus other decarbonization options.)

Furthermore, vehicle electrification does not address the increase in total vehicles distance traveled and the resulting emissions and congestion impacts that may be caused by ridesourcing fleets. This helps to motivate the second study.

5.1.2 Second study's findings

An additional pathway to reduce negative impacts of ridesourcing fleets' emissions and traffic is to reduce vehicle distance traveled per passenger distance traveled, which is the focus of the second study. This study contributes to the literature by sizing the market failure of unpriced emissions and traffic externalities (and, correspondingly, sizing how much internalizing externalities changes optimal strategies). It finds the market failure's distortion to be comparatively small, around a 2 percentage point shift in ridesplitting opt-in rates. It finds the tax's effect to be similarly small, a 1% reduction in externalities and a slightly smaller net change in total social (private plus external) costs.

As in the first study, this study assumes travel demand is fixed (though each trip may be served as a solo or shared ride). If a tax on ridesourcing traffic and emissions costs were created, it seems conceivable that some trips would not happen, others would shift to other travel modes, and others would shift to times where the traffic and emissions impacts are less severe. Not all trips are equal in impacts; depending on proximity to other trips and available vehicles, some trips may result in almost no marginal vehicle distance traveled, while others may require long detours by empty vehicles. This study does not consider any of these shifts, and by doing so, considers only the market failure as it pertains to ridesplitting.

To some policymakers and stakeholders, the small size of this market failure may be taken as good news: it provides evidence that private costs already incentivize an economically efficient level of ridesplitting. This could imply that even though a market failure exists, market forces (i.e., the pressure to reduce vehicle distance traveled in order to reduce operations costs) may already provide incentives to implement a level of ridesplitting that is nearly socially optimal.

However, this result must be interpreted in the context of real-world fleet operations. This study shows a private cost-optimal fleet uses ridesplitting for just over half of all trips, around twice the proportion observed in historical Chicago trip data. This may be because this model assumes a stronger ability to plan ahead when matching rides than ridesourcing services can do in practice, and because it assumes all riders can be matched into a shared ride (at some cost of lost privacy and time). Another plausible explanation for this gap is that

incentives may be misaligned. For example, perhaps the parties most affected by operations costs are drivers, and those driver costs are not perfectly reflected in the costs considered by riders (who make the decision to opt in to ridesplitting) and the fleets' algorithms (which make the decision to combine ride requests into cars). If so, perhaps policymakers or fleet operators may have a role to play in fixing incentive structures, rather than in correcting for unpriced externalities.

A secondary finding pertains to the potential of ridesplitting as a pathway to reduced ridesourcing impacts. Similar to a finding regarding electrification from the first study, this study finds that the option of ridesplitting, if used on a private cost-optimal set of rides, can meaningfully reduce external costs even without a Pigovian tax. However, also like the first study, this study provides little direct evidence that the pathway of economically efficient ridesourcing operations holds much potential by itself to induce a massive change in societal outcomes pertaining to vehicle-distance traveled. For policymakers who believe drastic reductions in vehicular impacts are warranted—for example, to improve air quality or for Vision Zero campaigns to fully eliminate traffic fatalities—an “all of the above” approach may be needed, including additional pathways such as infrastructure changes and shifts from passenger automobiles to other transportation modes.

5.1.3 Third study's findings

The third study contributes to the literature a historical lookback of Chicago's congestion zone pricing, and concludes that the pricing policy induced a meaningful shift in outcomes. Relative to the rest of the city, areas within the downtown congestion surcharge zone saw an 8.2% [7.6%,8.7%] reduction in total rides, a 4.2 [4.0,4.3] percentage point increase in ridesplitting opt-in rates, and a 3.4 [3.3,3.5] percentage point increase in rides that are ultimately matched with another ridesplitting ride.

This suggests that Chicago's policy may have shifted behavior and ridesplitting operations in the intended direction. It is for policymakers to say whether the size of the estimated shift met their goals, and further research is needed to determine whether these induced shifts led to reductions in congestion externalities.

5.1.4 Comparison across studies

While each study considers a different research question using a different analysis population, it is worth briefly comparing results across studies.

The first and second studies both consider a Pigovian tax's impact on operations (electrification or ridesplitting) and resulting outcomes. Comparing findings, the difference in effects of a Pigovian tax may be explained by underlying mechanics. The first study shows that private cost tradeoffs for conventional versus electric powertrains depend on many factors including each vehicle's use level, and that air emissions external costs can tip the scales (whether towards hybrid or battery electric vehicles) for a substantial fraction of the fleet. In the case of the second study, the external cost incentives for ridesplitting appear to be largely aligned with private cost incentives: many rides are inefficient to pool together (due to increased vehicle-distance traveled or increased travel delays) regardless of external costs, many that are efficient to pool together regardless of external costs (because they are geographically convenient), and relatively few rides where a tax would tip the scales.

The second and third studies both consider behavioral or operational responses to price signals to use ridesplitting. The two studies are not directly comparable due to different model assumptions and different price mechanisms being considered. The Pigovian tax considered in the second study internalizes total external costs for every unit of vehicle travel in order to correct incentives and provide a socially optimal level of pooling, while Chicago's fee differential is explicitly intended to discourage solo rides and thus is a premium paid for downtown solo rides. Despite not being directly comparable, it is worth noting that relative to the second study's price signals, those considered in the third study yielded a different net shift in ridesplitting use. Chicago's solo downtown ride fee differential (\$1.75) is smaller in magnitude than the second study's average Pigovian tax on total externalities (around \$4 per trip on average taking into account empty "deadhead" travel), but that Pigovian tax affects every solo and shared ride to varying degrees. The second study's optimal fleet responds to a Pigovian tax by increasing ridesplitting by just under 3 percentage points, while Chicago's actual policy increased realized ridesplitting matches by 3.4 percentage points [95% CI: 3.3,3.5].

This difference is small and there are several mutually compatible explanations for it. One is that the second study’s myopic optimization, including heuristics to rebalance fleet vehicles, could create a suboptimal price response. Another is that the analyzed population differs: the second study considers citywide changes averaged across all times and locations, while the third study specifically estimates effects on high-surcharge areas and only during weekdays from 6AM-10PM. A third possibility is that the costs modeled in the second study (using flat values based on USDOT guidance) do not perfectly match true valuations of weekday travelers in downtown Chicago (which may swing drastically depending on weather, travel budget, time pressures, and other factors). This suggests a version of the second study calibrated to better match actual ridesplitting levels (e.g., by using different travel time valuations) is warranted. If findings pertaining to market failure size, and shift induced by a tax, do not change across traveler preferences, then the second study’s findings are likely robust. If the market failure appears larger when “status quo” ridesplitting use is lower, then a case may still exist for policy intervention on economic efficiency grounds.

5.2 Policy implications

Taken as a group, these three studies characterize the degree to which unpriced external costs lead to inefficient or inequitable outcomes, and the extent to which specific pathways (electrification and ridesplitting) may improve those outcomes. They also provide estimates—both simulation-derived and data analysis-derived—of the impacts of ridesourcing with and without policy interventions. In both the case of electrification and ridesplitting, a market failure exists; however, the size of market failure (and practical effect of a Pigovian tax) is much larger in the case of electrification. Prioritizing electrification versus ridesplitting may not be a true dilemma faced by policymakers, but this relative sizing may help inform policymakers choosing a pathway to prioritizing incentivize or regulate. Awareness that these market failures exist, and that the status quo operations are distorted in an economically inefficient manner, can also improve the public’s (and private firms’) understanding of the value of price-based policy responses.

This dissertation also considers tradeoffs between economically efficient price-based policies and more blunt policy instruments. In the first study, for example, we discuss

scenarios where the fleet is single-technology, which could represent a technology-forcing mandate rather than a Pigovian tax (discussion in the full manuscript provided in Appendix B and in that paper’s supporting information). An optimally mixed set of powertrains, as induced by the Pigovian tax, lowers social costs the most and is thus most economically efficient. However, an all-BEV scenario results in air emissions external costs that are lower than other single-technology fleets and, in two of three cities considered, lower than those of a socially optimal fleet. (In the third city, an all-BEV fleet has lower emissions costs than a private-optimal mixed fleet with no Pigovian tax, but slightly higher emissions costs when a tax is included). This implies a 100% technology standard could, in some circumstances, reduce emissions costs more substantially than a market-based price signal, but at the expense of lost efficiency (higher total social costs versus the optimal fleet). Despite this inefficiency, technology standards could conceivably be more politically viable and simpler to implement; they also may or may not lead to more equitable outcomes than a perfectly efficient pricing scheme. A similar tradeoff could exist with ridesplitting: an economically efficient Pigovian tax may best minimize total social costs, but standards-based approaches (e.g., a vehicle occupancy level or “deadheading” ratio mandate) or alternative price signals may be more feasible to implement than a tax or may lead to more substantial external cost reductions at the expense of efficiency (depending on how relaxed or aggressive those alternative designs are).

Within the set of possible pricing schemes, the third study does not consider the economic efficiency of Chicago’s chosen fee level, but it is clear that any flat rate is inherently less than perfectly efficient, since external costs of traffic vary in real-time as roads become more or less busy. However, it is still valuable to know that a pricing scheme that has been enacted—i.e., that was not algorithmically complex, and that was at a price level policymakers could champion and the public could tolerate—can induce meaningful behavioral shifts.

It is for policymakers and stakeholders to consider tradeoffs related to implementability, economic efficiency, and efficacy at meeting goals in choosing policies for ridesourcing. However, these studies make clear that the operations and impacts of ridesourcing are not predestined, but rather they depend on incentive structures. To the extent that those incentive structures do not reflect the full costs of travel, or do not match the priorities of

stakeholders, public policy can and should improve outcomes.

5.3 Methodological findings

5.3.1 Optimization methods

A well-known result in optimization, the “no free lunch” theorem of Wolpert and Macready (1997), implies that improving the performance of a “one-size-fits-all” optimization solver on any given problem requires customizing the solution steps in a problem-specific manner. For the first and second studies, substantial performance improvements were made by customizing solution methods to fit the problem, and customizing the problem formulation to fit the solver.

For the first study, two customized heuristic steps were used to make a challenging optimization model solvable and to help narrow the gap between the best integral solution found (primarily via the Gurobi solver’s heuristics) and the tightest linear relaxation solution found (primarily via Gurobi’s presolve steps, cutting planes, and, in fewer cases, the branch and bound search tree). One, adapted from Bertsimas et al. (2017), discovered and pruned likely sub-optimal pairings of vehicles to trips by running many simplified versions of the problem with inputs perturbed. Another, created from scratch, simplified the optimization formulation’s most complicating constraints (battery charge scheduling, which includes some integrality constraints) in a stepwise manner by optimizing one vehicle’s battery at full fidelity and the remainder of the fleets’ batteries in the aggregate. By iteratively removing solved vehicles from the solution and “shrinking” the size of the aggregate battery being solved, each iteration can solve with all integrality constraints met and is an increasingly close approximation to vehicle-level charging constraints. A more problem-agnostic stepwise heuristic would have been to model every individual vehicle’s charging first with integrality constraints relaxed, then iteratively make one additional vehicle observe integrality. This approach did not perform as well, because the linear relaxation itself (which included all decision variables for all vehicles) was a very large problem with a long solution time and a larger memory footprint.

For the second study, the primary bottleneck in solving the optimization was not in the solution itself, but in enumerating massive numbers of combinations of multiple rides

into ridesplitting trips. As an illustration of problem size, a typical five-minute time period could have a batch of around 1,000 new ride requests, and determining the costs of every three-way combination of requests would imply nearly a billion evaluations (multiplied by the number of available cars, perhaps 25,000)—too many options to fit in memory or to evaluate in reasonable time, and certainly too many to include in an integer programming model. Solving this bottleneck followed the approach from Alonso-Mora et al. (2017b), and used some additional “shortcuts”. We pre-compute travel times and costs for all possible routes in the city and store them in a vector for fast lookup, and we test each combination of rides for feasibility against an imaginary, perfectly located vehicle before testing it against any individual vehicles operating in the fleet. We also parallelize complicated steps as much as possible.

Both studies tested multiple optimization formulations with Gurobi (as well as CPLEX, which was not used for any final model runs) and customized the problem formulation to the solution method. In particular, both studies are solved more quickly by using totally unimodular constraint matrices when possible. A fundamental integer programming result (from Cramer’s Rule) is that when a constraint matrix is totally unimodular and its right-hand side is integral, basic solutions (including those found by the simplex method) are integral. This implies relatively fast and predictable solution times, since time-consuming branch and bound search is not needed. Network matrices are unimodular, and the first study’s pruning of likely sub-optimal options relies on one such matrix in a minimum-cost network flow model that solves quickly enough to run very many times with inputs perturbed (but with acquisition costs simplified and battery charge constraints excluded). Bipartite graph matching matrices are also unimodular, and the first stage of the second study’s vehicle rebalancing can solve very quickly as a result.

5.3.2 Machine learning to enable representative policy findings

In each study, machine learning was used to make input data, and resulting policy findings, representative in a data-driven manner.

The first study uses supervised machine learning to fill in gaps in ridesourcing data. Specifically, the available data indicated vehicles’ average travel speed during each passenger

ride, but provides no direct information about between-ride traffic conditions, which is necessary to estimate which rides are feasible to join into linked into chains of (solo, non-overlapping) rides. A k-nearest neighbors regression was used to estimate traffic conditions along a given corridor at a given time from similar nearby co-occurring trips. This is a simple approach that is nonparametric (i.e., it makes no assumptions about underlying data distributions), but out-of-sample validation showed it performed well except in the times and neighborhoods with very sparse trip data.

The second study uses unsupervised learning to choose a subset of trip data that is most representative of typical travel demand in Chicago. Most readily available computing environments cannot facilitate solving for a full year (tens or hundreds of millions of trips); the Pittsburgh Supercomputing Center, for example, has a 48-hour runtime limit. To choose a representative week, this study first uses partitioning around medoids (using dynamic time warping distance, appropriate for time-series data) to characterize the subgroups of hourly travel demand patterns that underly a typical week’s aggregate travel, then uses k -medoids clustering (where $k = 1$) to find the week with the most typical volumes from each underlying demand pattern.

The third study also uses unsupervised learning, but with the goal of creating observational units of similar travel volumes that preserve geographic relationships. As described in Chapter IV, similar-sized observational units are required for effect estimates to carry meaning for total downtown effects. To generate these similar-sized units, the study first uses weighted k-means to determine each observational unit’s “central” origin-destination pair, then iteratively assign origin-destinations to a nearby grouping to maximize travel volume similarity of the k groups. This approach results in nearly identically-sized observational units, but preserves enough geographic meaning that spatial autocorrelation still has explanatory power in the econometric model.

These data-driven estimation approaches may be more challenging to explain than a rule-based heuristic approach but may be more defensible. Taking the question of ridesplitting as an example, instead of using unsupervised methods to determine a representative subset of demand (as in the second study), one could simply choose to stratify a trip sample by season and day of week (as done in the first study). This would be simple to explain and defensible.

However, such a sample may not capture meaningful sources of variation—it is possible, for example, that travel demand patterns, and the viability of ridesplitting, differ very little between fall and spring but differ widely between holiday and non-holiday periods, or rainy and sunny periods. A data-driven approach can automatically characterize these variations (agnostic to their underlying source) and determine representativeness across longitudinal data.

These approaches are also easily customizable. For instance, the number of “typical” groups can be easily adjusted as needed to meet computational requirements. If atypical periods of demand are necessary to test model robustness, one can rank periods of time by difference from “typical” (using the same definitions of “distance” and “typical” as the initial unsupervised learning step). And if certain specific variables are of special importance to the model (e.g., concentration of trips downtown vs. non-downtown), those variables can be included in the learning step and weighted as desired.

5.4 Modeling lessons learned

In order to solve their research questions, each study involved a mix of computationally intensive solution methods (studies one and two) and large dataset processing (all studies, but particularly the second and third). This section summarizes lessons that were learned while building, testing, and solving each model.

5.4.1 Importance of benchmarking on realistic models

Conducting experiments for the first two studies illustrated the importance of testing models with realistic problem sizes, model parameters, and correct objective functions. For each study, early “toy models” used for exploration of the problem suggested certain solution steps that could be bottlenecks worth customized coding or reformulation. However, once these models were “scaled up” to meaningful problem sizes or certain parameters were changed (e.g., allowable travel delays in the second study), entirely different solution steps became bottlenecks. This is likely because the algorithmic complexity of certain steps is constant or linear with problem size, while other steps may increase in complexity faster. Using realistic problem sizes early can make time spent on code optimization more impactful.

A similar trend was observed for objective functions: in early versions which excluded certain cost components, the optimization solver was able to quickly find optimal solutions, while the final (more comprehensive) objective function solved more slowly. This suggests that when model parameters (e.g., steps that prune likely sub-optimal choices) are chosen to meet some specific solution time goal, those choices should be made using the final objective function.

5.4.2 Value of high performance computing resources

All three studies involved datasets or optimization models that would not have easily run on most personal computers. All three used the Pittsburgh Supercomputing Center’s Bridges and/or Bridges-2 systems for the majority of work and the Engineering & Public Policy (EPP) departmental computing server for additional model building and results analysis. Working extensively with these systems has provided lessons regarding areas in which supercomputing resources help solve problems, and areas in which they may provide less value.

One area where these computing resources added value for all three studies is their random access memory (RAM) capacity. Depending on the configuration used, the EPP departmental servers offer 128 gigabytes (Windows) or 200 gigabytes (Linux) of shared RAM, while Bridges-2 reserves for private use up to 256 or 512 gigabytes (2 or 4 gigabytes per computing core requested by the user, where each core costs additional computing credits). This is useful because some computational steps from these studies could not fit into RAM easily. Examples include the second study’s graph enumeration steps (because of the large number of feasible ridesplitting matches) and results processing (because there are many vehicles and their turn-by-turn routes are used), creating analysis datasets for the final two studies (because trip datasets include hundreds of millions of rows), and the final two study’s unsupervised clustering methods (because of large distance matrices being built). Other examples include large linear programs, especially when being solved using the barrier (interior point) method. While the final versions of each study’s optimization steps would fit on many personal computers, the early versions (or alternative versions, such as if vehicle routing had been optimized with a “set packing” formulation) may not have.

Allowing multiple modeling scenarios to run at once in parallel is another area where high performance computing added value. This was possible on Bridges and Bridges-2, which consist of multiple servers that can be reserved for private use (unlike the EPP departmental server, which is fully shared). The second study's model took over 24 hours to run, so being able to run multiple instances at once without sacrificing performance was helpful.

5.4.3 Limitations of parallelization

Within each individual model solving process, the value of parallelization was more limited. The EPP departmental servers have 20 computing cores each, while Bridges-2 allows reserving up to 128 computing cores (where computing credits are paid in units of core-hours used). In general, parallelization did not make any unsolvable optimizations solvable. For integer program solutions, using multiple parallel processing threads only offered marginal runtime improvements and added little additional value above around 10-15 threads. Parallelization did provide substantial speedups for linear program relaxation solves for the first study (using the barrier method, which benefits from parallelization), but did not do so for the second study's optimization steps (for which the simplex method, which runs serially, was as fast as the barrier method).

For the second study, the ability to run graph enumeration steps in parallel rather than as serial 'for' loops improved runtimes substantially, but required extensive customized code using the Open Multi-Processing (OpenMP) framework. Larger performance improvements were made by carefully choosing appropriate data structures and by running a code profiler tool to identify bottlenecks. (Experimentation was done with Intel's oneAPI Thread Building Block concurrent data structures, but computational overhead exceeded benefits relative to the C++ Standard Template Library.)

5.5 Future work

These studies help to address certain knowledge gaps in our understanding of ridesourcing and its associated market failures, but their findings suggest several additional research needs.

5.5.1 Direct extensions to this work

The findings described in this dissertation lead directly to several additional areas that warrant further research. One such area is the distribution of outcomes and the underlying equity implications. Each study was primarily focused on total or average outcomes. Those outcomes may provide sufficient justification for policy intervention, but other justifications may exist if different groups may be affected differently by “status quo” ridesourcing use or by any given policy mechanism. A natural extension to each study could consider how different subpopulations such as neighborhoods of the city, racial or other demographic groups, or segments of travel demand (e.g., commuters versus recreational travelers) are affected disproportionately by ridesourcing use or whether the analyzed policy mechanism redistributes external costs in an equitable manner. If not, then a role for policy exists not only to correct inefficient outcomes, but also to incentivize or enforce more equitable outcomes.

In particular, the third study’s findings suggest value in pursuing additional areas using breadth of data available for the Chicago region. Through the Chicago Data Portal that publishes the ridesourcing trip datasets used for the second and third studies, the city also publishes a range of other datasets. For example, their congestion datasets would enable an immediate follow-on to study three considering congestion as an outcome directly affected by congestion pricing (as measured using road speed of buses, which report speed data at regular intervals). City of Chicago (2021a,b) Similarly, the same portal’s datasets for individual and/or aggregated trips on public transit, yellow taxicabs, bikeshare, and electric scooters allow each mode’s use to be treated as a dependent variable. City of Chicago (2021b) Separately, crowdsourced air quality sensor data, publicly available on through the Purple Air sensor network’s webpage, would enable investigating relationships between ridesourcing use, ridesourcing regulation or taxation, and urban air quality. PurpleAir (2021) Another potential direction of study would use regression discontinuity on the repeated “sharp” switching off and on of the congestion pricing plan (at 6AM and 10PM each day) to consider price responses assuming nothing else changes immediately—for example, to see whether fee increases are fully passed on to the rider or whether fares adjust downward.

All three studies suggest that policies can nudge ridesourcing use towards electrification and ridesplitting, but they do not identify a singular fix for the negative impacts of ridesourcing. Rather, each one estimates small to moderately sized improvements in outcomes. This leaves open for further research a general question: *do other pathways have greater potential to reduce external impacts of ridesourcing?* In addition to electrification and ridesplitting, other pathways could be investigated in a similar manner—considering changes induced by a Pigovian tax—such as switching between ridesourcing and other modes (private vehicle, active transportation, transit) or managing demand (nudging users to combine multiple errands into one trip or to change departure time). Policymaking in this area could be improved by understanding the full space of pathways to improved societal outcomes, the relative efficacy of each option, and interaction effects of pursuing multiple pathways at the same time.

Another possible family of research questions consider how broadly applicable these findings may be: *which findings pertaining to ridesourcing carry over to broader personal vehicle travel?* Because ridesourcing trip datasets are more fine-grained than most available data on private vehicle use (and, in the small number of cities where they are published, they represent a full population rather than a surveyed sample), we might hope to be able to leverage that data to learn more about personal vehicle travel. It is conceivable, for example, that the cost-efficacy of ridesourcing electrification with and without a Pigovian tax could teach us something about market failures in personal travel air emissions; it is also conceivable the use cases are too different. Similarly, within the realm of congestion pricing, ridesourcing may serve as a test bed for schemes that could later apply to all vehicle travel within an urban area.

5.5.2 Adjacent research areas

A different area for future work pertains to the current status quo role of ridesourcing fleets and how they interact with cities. While this dissertation specifically considers market failures existing in the status quo, they skirt a more fundamental question: *on net, do ridesourcing fleets currently have more positive or more negative impacts on cities?* Prior studies have considered this question using across-city data, using various dimensions and

definitions for positive or negative impacts. Ward et al. (2021b, 2019); Hall et al. (2018); Babar and Burtch (2020) However, most existing causal studies have largely considered ridesourcing services' entry effects, rather than the effects of a marginal unit increase in ridesourcing use; the latter may more closely relate to future outcomes. Also, relatively few have used more granular trip-level data to consider, within a given city, how outcomes have evolved so far over time or how they differ from neighborhood to neighborhood. This could substantially change our understanding of how these services may affect cities going forward, or what their impacts are on transportation equity.

Given the lack of a clear panacea for reduced air emissions or vehicle-distance traveled within the ridesourcing sector, another research question left open is broader: *is there a set of policies or other developments that might drastically decarbonize passenger transportation?* In some respects, we might expect passenger transportation to serve as a comparatively “lower hanging fruit” for decarbonization than other segments of the transportation sector, because of passenger vehicles being smaller and easier to electrify than some other classes of vehicle (e.g., agricultural, industrial, long-haul freight). Furthermore, compared to other passenger vehicles, ridesourcing vehicles see higher utilization levels and faster fleet turnover than private vehicles, and one might surmise they could be easier to decarbonize quickly. If there is no economically efficient path forward to drastically and rapidly reduce carbon impacts from ridesourcing, the picture is even less optimistic for other forms of transportation.

5.5.3 Data improvements

These studies leverage publicly available datasets for ridesourcing use, and they would not have been possible without those datasets being made public. However, they also demonstrate the limitations of existing ridesourcing data and highlight the need for additional data.

As valuable as the available analysis datasets have been for the research community, there are substantial missing pieces that could aid future work. Specifically for ridesourcing data, available datasets indicate little about how much vehicle travel (or parking space occupancy) occurs in between trips. This “deadhead” travel may play an outside role in determining a service's overall impacts on a city, and providing it in summary form or as vehicle trajectory data would reduce the risk of unintentionally biased or misleading research

findings. Separately, analyses of demand flexibility and price signal efficacy are impeded due to the lack of data on ridesourcing on fare offers, acceptances, and rejections. At least one city (Chicago) collects this data but does not currently publish it; this data would strengthen the ability to model the full effects of price signals and aid in optimal price design.

More generally, without additional municipalities and regions making ridesourcing data public, it is not practical to generate policy findings that are robust across regions (or to check for robustness). We cannot easily estimate how specific to Chicago our findings for ridesplitting optimization may be; the underlying demand patterns may be unique in meaningful ways, and other external factors (e.g., weather and urban sprawl) that are regionally variable may have significant impacts on findings. There is value in regionally-specific analysis, and if more municipal and regional governments partner with ridesourcing services (and other transportation providing services) to make more data available, it can only benefit the policymaking process in those areas.

APPENDIX A

Chapter 3 Algorithms

Algorithm 1: Vehicle-to-request enumeration

Input: Vehicles \mathcal{V} , requests \mathcal{R}

Output: Request-vehicle graph G_{RV} of feasible connections from vehicle nodes \mathcal{V} to request nodes \mathcal{R} with edges denoting costs

Initialize: $G_{RV} \leftarrow \emptyset$

for $v \in \mathcal{V}$ **do**

for $r \in \mathcal{R}$ **do**

if $\delta(v, r) \leq \delta^{MAX}$ **then**

$G_{RV} \leftarrow G_{RV} \cup \text{edge}(v, r, \gamma(v, r))$

end

end

end

Algorithm 2: Request-to-trip enumeration**Input:** Request-vehicle graph G_{RV} connecting vehicle nodes \mathcal{V} to request nodes \mathcal{R} **Output:** \mathcal{T} , trips an ideal car could feasibly serve; \mathcal{D} , dependency map of trips to “exclusion rules” (trip t ’s entry in \mathcal{D} lists trip indices whose feasibility depends on t ’s)**Initialize:** $\mathcal{T} \leftarrow \mathcal{R}$ (all requests are added as trips of size $k = 1$)

```

for  $k = 2$  to  $k^{MAX}$  do
  // Build  $\mathcal{T}^*$ , a list of size  $k$  trip candidates
  if  $k = 2$  then
    if  $\exists v \in G_{RV} : \text{degree}(v) = |\mathcal{R}|$  then
      for  $(r_i, r_j) \in \mathcal{R}$  do
        | add trip( $r_i, r_j$ ) to  $\mathcal{T}^*$ ;
        | add  $r_i$  and  $r_j$  to  $D^{REVERSE}(\text{trip}(r_i, r_j))$ 
      end
    else
      for  $v \in \mathcal{V}$  do
        | for  $(r_i, r_j) \in G_{RV} : r_i, r_j$  adjacent to  $v$  do
        | | add trip( $r_i, r_j$ ) to  $\mathcal{T}^*$ ;
        | | add  $r_i$  and  $r_j$  to  $D^{REVERSE}(\text{trip}(r_i, r_j))$ 
        | end
      end
    end
  else
    for  $t \in \mathcal{T} : |t| = k - 2$  do
      | for  $(t_i^{DEP}, t_j^{DEP}) \in \mathcal{D}(t)$  do
      | | add trip( $\text{reqs}(t_i^{DEP}) \cup \text{reqs}(t_j^{DEP})$ ) to  $\mathcal{T}^*$ ;
      | | add  $t_i^{DEP}$  and  $t_j^{DEP}$  to  $D^{REVERSE}(\text{trip}(r_i, r_j))$ 
      | end
    end
  end
  // Add  $\mathcal{T}^*$ ’s candidates that are complete cliques and feasible to  $\mathcal{T}$ 
  for  $t^* \in \mathcal{T}^*$  that was added to  $\mathcal{T}$   $k * (k - 1)/2$  times do
    | for  $r \in t^*$  do
    | | place empty vehicle  $v^*$  at starting point of  $r$ ;
    | | if  $\delta(v^*, t^*) \leq \delta^{MAX}$  then
    | | | add  $t^*$  to  $\mathcal{T}$ ;
    | | | exit inner loop;
    | | end
    | end
    | if  $t^*$  was added to  $\mathcal{T}$  then
    | | for  $t \in D^{REVERSE}(t^*)$  do
    | | | add  $t^*$  to  $\mathcal{D}(t)$ ;
    | | end
    | end
  end
end

```

Algorithm 3: Trip-to-vehicle enumeration

Input: Vehicles \mathcal{V} , requests \mathcal{R} , Request-vehicle graph G_{RV} , trips an ideal vehicle could serve \mathcal{T} , map of trips to “exclusion rules” \mathcal{D}

Output: Request-trip-vehicle graph G_{RTV} of feasible connections from vehicle nodes \mathcal{V} to trip nodes \mathcal{T} with edges denoting costs

Initialize: $G_{RTV} \leftarrow G_{RV}$, with edges linking feasible vehicles and trips of size $k = 1$

```

for  $v \in \mathcal{V}$  do
  for  $k = 2$  to  $k^{\text{MAX}}$  do
    for  $t \in \mathcal{T} : |t| = k$  that is not ruled out for  $v$  do
      if  $\delta(v, t) \leq \delta^{\text{MAX}}$  then
         $G_{RTV} \leftarrow G_{RTV} \cup \text{edge}(v, t, \gamma(v, t))$  ;
      else
        for  $t^{\text{DEP}} \in \mathcal{D}(t)$  do rule out  $t^{\text{DEP}}$  for  $v$  ;
      end
    end
  end
end

```

APPENDIX B

Chapter 2 Manuscript

Effects of Air Emission Externalities on Optimal Ridesourcing Fleet Electrification and Operations

Matthew B. Bruchon, Jeremy J. Michalek,* and Inês L. Azevedo

Cite This: *Environ. Sci. Technol.* 2021, 55, 3188–3200

Read Online

ACCESS |

Metrics & More

Article Recommendations

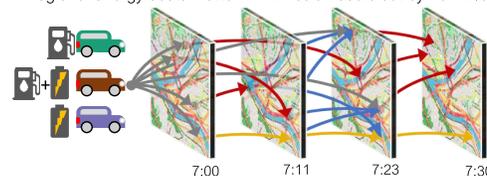
Supporting Information

ABSTRACT: Ridesourcing services from transportation network companies, like Uber and Lyft, serve the fastest growing share of U.S. passenger travel demand.¹ Ridesourcing vehicles' high use intensity is economically attractive for electric vehicles, which typically have lower operating costs and higher capital costs than conventional vehicles. We optimize fleet composition (mix of conventional vehicles (CVs), hybrid electric vehicles (HEVs), and battery electric vehicles (BEVs)) and operations to satisfy demand at minimum cost and compare findings across a wide range of present-day and future scenarios for three cities. In nearly all cases, the optimal fleet includes a mix of technologies, HEVs and BEVs make up the majority of distance traveled, and CVs are used primarily for periods of peak demand (if at all). When life cycle air pollution and greenhouse gas emission externalities are internalized via a Pigovian tax, fleet electrification increases and externalities decrease, suggesting a role for policy. Externality reductions vary from 10% in New York (where externality costs for both gasoline and electricity consumption are relatively high and a Pigovian tax induces a partial shift to BEVs), to 22% in Los Angeles (where high gasoline and low electric grid externalities lead a Pigovian tax to induce a near-complete shift to BEVs).

We optimize fleet purchases and operations for two objectives:

1. Private Costs 2. Private Costs + Emission Costs

- Fleet uses mix of powertrains
- Regional energy costs matter
- Electrification increases
- Emission costs cut by 10–22%



1. BACKGROUND

Passenger cars produce the largest share of greenhouse gas (GHG) emissions from U.S. transportation, which recently surpassed electric power as the country's highest-emitting economic sector.² Passenger cars also emit substantial conventional air pollution, and premature mortality from U.S. air pollution (28% of which results from transportation) is comparable to automobile accident fatalities, with an annual social cost of \$886 billion.³

Ridesourcing services are rapidly and dramatically changing the passenger car landscape: from 2009 to 2017, for-hire vehicles in the United States more than doubled their share of trips and their daily per capita usage, due primarily to the rapid growth of ridesourcing services,¹ and by 2016, 15% of intraurban trips in San Francisco were served by Uber and Lyft.⁴

Vehicle electrification has the potential to drastically reduce ridesourcing emissions while perhaps also lowering operating costs. Electricity is often cleaner and cheaper than gasoline per vehicle distance traveled (VDT), and for intensively used vehicles lower fuel costs and operation emissions might offset their higher upfront costs and manufacturing emissions. The Intergovernmental Panel on Climate Change recently stated that electric modes of transportation would “need to displace fossil-fuel powered passenger vehicles by 2035–2050 to remain in line” with pathways to hold global warming to 1.5 °C.⁵ Recognizing the potential of transportation network company (TNC) fleet electrification to reduce transportation emissions, the California Public Utilities Commission in 2018 released an

initial overview of regulatory approaches that are worth further research as a means to encourage TNC electrification, including technology mandates, distance-based fees on combustion engine usage, and financial incentives.⁶ Also in 2018, Uber announced a goal of an all-electric vehicle (EV) fleet within the city of London by 2025. This plan's stated motivation is to reduce pollution, and a per mile “clean air fee” will fund driver financing programs.⁷ Advances in vehicle electrification and automation may transform the way ridesourcing services operate.⁸

However, the premise that full fleet electrification is a viable or desirable policy goal warrants further investigation. At the current cost of lithium-ion batteries, battery electric vehicles (BEVs, which plug in to charge and rely entirely on electricity stored in large battery packs) have a much higher upfront cost than conventional vehicles (CVs); battery manufacturing emissions are nontrivial;^{9,10} and, depending on region, timing, and vehicle design, electric vehicles do not always reduce air pollutant emissions or greenhouse gas emission externalities compared to CVs (with lower-income census block groups more

Received: July 31, 2020

Revised: January 22, 2021

Accepted: January 25, 2021

Published: February 19, 2021



likely to face increased emission externalities from BEVs).^{11–14} Furthermore, the operations of BEVs suffer from logistical constraints of limited range and slower refueling (charging). BEVs cannot service demand while charging, so a larger fleet is required to satisfy a given level of demand. BEVs also must detour to recharge, increasing VDT. In contrast, gasoline hybrid electric vehicles (HEVs, which draw all net energy from gasoline but use a battery and electric motor to improve efficiency) have no additional range or refueling constraints, but they do burn gasoline and emit pollution from the tailpipe.

In general, it may be that the lowest-cost or lowest-emission fleet does not use a single homogeneous technology but, rather, a mixture of technologies, with different duty cycles (e.g., peak versus off-peak) being served by different technologies.

We investigate the optimal technology mix and operations of a ridesourcing fleet whose operator has perfect foresight of exogenous (and inflexible) passenger trip requests and total control over fleet acquisition and routing. Centralized control of fleet vehicle choices may represent ridesourcing companies that have owned or leased vehicles in some locales,^{13,16} a future with autonomous vehicle fleets,^{17,18} or vehicles that are purpose-built for ridesourcing fleets.^{19,20} Centralized vehicle routing may become widespread as autonomous vehicle technology advances, whereas today's ridesourcing services only approximate centralized routing via human drivers responding to ride requests and price signals. Also, regulations and incentives that operate at the level of the fleet, rather than the individual driver (e.g., California's under-development Clean Miles Standard, which will regulate fleet-wide annual CO₂ emissions per passenger-mile²¹), increase the role of centralized fleet-wide planning, coordination, and control.

We assess the policy opportunity of electrification by comparing costs and emissions of pure CV, HEV, and BEV fleets with mixed fleets across a range of scenarios. By comparing cases that include or exclude emission externality costs in fleet optimization, we assess the degree to which unpriced emission externalities bias fleet outcomes away from socially optimal solutions and consider whether policy intervention may be therefore justified on economic efficiency grounds.

1.1. Literature. A body of literature considers operations and outcomes of electrified vehicle fleets, but the question of electrification's role within a ridesourcing fleet's optimal technology mixture and its impact on resulting emissions is relatively unexplored.

Some studies use agent-based modeling (ABM) to explore the operational impacts of homogeneous all-electric fleets. Bauer et al. estimate that such a fleet operating in Manhattan would reduce private costs and emissions relative to a homogeneous fleet composed of either CVs or HEVs, and Bauer et al. find that fleet-wide coordination of charging would allow BEVs serving demand New York City or San Francisco to meet the same level of service as CVs at a lower cost even if charge networks are relatively sparse.^{22,23} Our fleet differs in its optimization of the fleet mix under different objectives and its consideration of multiple cities. Chen et al. find that electrification can meet ridesourcing demand while barely increasing empty VDT, but only if the fleet size is increased.²⁴

Other studies use ABMs in combination with a second model. The scenarios simulated in Chen et al.²⁴ were used as a case study and defined the inputs for a life cycle assessment framework in Gawron et al., which found that a fleet of electric autonomous taxis could reduce cumulative greenhouse gas emissions by 60% in the period from 2020 to 2050 in the base

case and up to 87% in additional scenarios.²⁵ Sheppard et al. use an ABM to generate simplified operational parameters for a national-scale optimal sizing of vehicles and infrastructure for an all-electric fleet, estimating that 12.5 million vehicles could replace the fleet of 276 million personally owned vehicles.²⁶ Chen and Kockelman incorporate a logit choice model into an ABM to estimate that a shared, autonomous, all-electric vehicle fleet could capture 14–39% of all passenger trips within the Austin, Texas region, depending on pricing.²⁷

Studies employing ABMs use simplifying assumptions or heuristics to model agents' behavior. These heuristics' ability to achieve representative behavior or near-optimal behavior cannot easily be evaluated for each test case, so comparisons across scenarios can conflate effects of the scenarios with effects of the heuristics. Specifically, it is difficult to determine the degree to which differences in results across scenarios are due to differences in the scenarios themselves or due to differences in the performance of the heuristics across scenarios. Bertsimas et al. find that for vehicle routing problems, optimization coupled with well-designed heuristics increases fleet revenue results by as much as 9% relative to a heuristic alone and that heuristics perform unevenly across problem instances;²⁸ it is conceivable that this 9% gap widens when a fleet's technology mix is jointly optimized with its routing. Heuristics are necessary to address city-scale problems at manageable computational cost, but they introduce challenges for comparing across cases—such as comparing solutions with and without internalized air emission externality costs. To address this limitation, we pair heuristics with mathematical optimization to understand heuristic quality, to gain intuition on their biases, and to compare fairly across cases.

There is also a separate stream of methodologically focused research applying optimization to the routing of range-limited electric vehicles. These are typically conducted at a very small scale (exact solutions for 100–200 trips or heuristic solutions for several hundred more), they do not consider external costs, and they rarely jointly optimize purchases and routing even in cases when a fixed mixture of powertrains is assumed.^{29,30} The **Supporting Information (SI)** describes some of these studies in greater detail. Optimizing fleet size and mix at any scale requires careful model formulation and development of problem-specific heuristics, which our study contributes to its problem (applied to an instance of 5000 trips).

In the somewhat-related context of round-trip car-sharing fleets, in which the user pays for short-term rental of a car and drives it themselves, Zoepf finds that BEVs do have a niche to fill, reducing private costs when 20–40% of a gasoline fleet is electrified but increasing private costs beyond that threshold.³¹ For TNC fleets, in the grey literature, a 2019 International Council on Clean Transportation report examined powertrain choice from the perspective of TNC driver costs of vehicle ownership.³² It found that hybrids may be financially favorable and that battery vehicles may become favorable around 2023–2028, using assumptions for factors such as the total distance traveled per year that, in practice, vary across vehicles in the fleet. A later analysis by the same group found that a per-trip fee, indexed to tailpipe emissions, between \$0.58 and 1.12 would suffice to make BEVs economically superior to HEVs.³³

We contribute to the prior literature by (1) constructing a mixed-integer optimization model with heuristics that make meaningfully sized problems tractable and provide near-optimal solutions for fair comparisons across scenarios and (2) applying the model to characterize how the optimal technology mix,

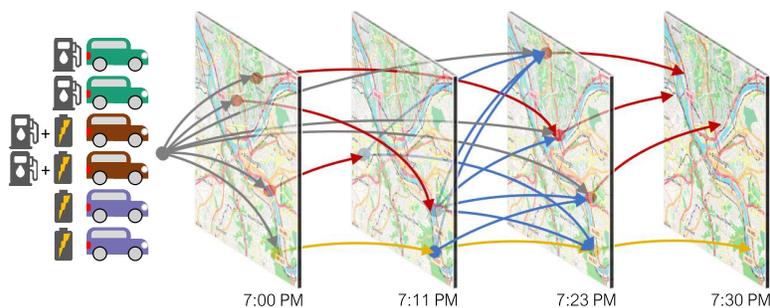


Figure 1. Illustration of the time–space graph showing passenger trips (red), charging arcs (yellow), dispatch arcs (gray), and relocation arcs (blue). Some arcs are omitted for simplicity.

Table 1. Formulation of the FullMILP Optimization Problem

$\begin{aligned} & \text{minimize} \sum_{k \in \mathcal{K}} \kappa_k + \sum_{(i,j) \in \mathcal{A}} c_{k,i,j} a_{k,i,j} + \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}_+} c_t \Delta q_{k,t}^{\text{CHG}} \\ & + \tau \left[\sum_{k \in \mathcal{K}} \delta_k + \sum_{(i,j) \in \mathcal{A}} d_{k,i,j} a_{k,i,j} + \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}_+} d_t \Delta q_{k,t}^{\text{CHG}} \right] \end{aligned}$	<p>(1a) minimize private and external car, fuel, and electricity costs</p>
<p>subject to</p> $\sum_{i \in \mathcal{V}} a_{k,i,j} = \sum_{i \in \mathcal{V}} a_{k,i,i}, \quad \forall k \in \mathcal{K}, j \in \mathcal{V} \setminus \{r, s\}$	<p>(1b) flow is preserved across nodes except source and sink</p>
$\sum_{k \in \mathcal{K}} a_{k,i,j} = n_{i,j}, \quad \forall (i, j) \in \{\mathcal{A}: n_{i,j} > 0\}$	<p>(1c) demand is satisfied</p>
$\sum_{j \in \mathcal{V} \setminus r} a_{k,r,j} = n_k, \quad \forall k \in \mathcal{K}$	<p>(1d) dispatched cars are purchased</p>
$\kappa_k \geq \alpha_{\omega,k}^{\text{COSTS}} n_k + \beta_{\omega,k}^{\text{COSTS}} \sum_{(i,j) \in \mathcal{A}} m_{i,j} a_{k,i,j}, \quad \forall k \in \mathcal{K}, \omega \in \Omega_k$	<p>(1e) per-vehicle private capital costs vary with usage</p>
$\delta_k \geq \alpha_{\omega,k}^{\text{X-COSTS}} n_k + \beta_{\omega,k}^{\text{X-COSTS}} \sum_{(i,j) \in \mathcal{A}} m_{i,j} a_{k,i,j}, \quad \forall k \in \mathcal{K}, \omega \in \Omega_k$	<p>(1f) per-vehicle manufacturing external costs vary with usage</p>
$q_{k,t+1} \leq q_{k,t} + \sum_{(i,j) \in \{\mathcal{A}: t_i=t\}} a_{k,i,j} \Delta q_{k,i,j}^{\text{MAX}}, \quad \forall k \in \mathcal{K}_{\text{B}}, t \in \mathcal{T}_{\text{Q}}$	<p>(1g) BEV charge level is tracked (times with a charger starting timeslot)</p>
$q_{k,t+1} = q_{k,t} + \sum_{(i,j) \in \{\mathcal{A}: t_i=t\}} a_{k,i,j} \Delta q_{k,i,j}^{\text{MAX}}, \quad \forall k \in \mathcal{K}_{\text{B}}, t \in \mathcal{T} \setminus \mathcal{T}_{\text{Q}}$	<p>(1h) BEV charge level is tracked (times with no charger starting timeslot)</p>
$q_{k,t}^{\text{END}} = q_{k,t}^{\text{START}}, \quad \forall k \in \mathcal{K}_{\text{B}}$	<p>(1i) final charge level equals initial charge level</p>
$0 \leq q_{k,t} \leq q_k^{\text{MAX}}, \quad \forall k \in \mathcal{K}_{\text{B}}, t \in \mathcal{T}$	<p>(1j) charge level does not exceed battery capacity</p>
$a_{k,i,j} \in \{0, 1\}, n_k \in \{0, 1\}, \quad \forall k \in \mathcal{K}_{\text{B}}, (i, j) \in \mathcal{A}$	<p>(1k) BEV routing and purchase decisions are binary</p>
$a_{k,i,j} \in \mathbb{Z}_+, n_k \in \mathbb{Z}_+, \quad \forall k \in \mathcal{K} \setminus \mathcal{K}_{\text{B}}, (i, j) \in \mathcal{A}$	<p>(1l) CV HEV routing and purchase decisions are integral</p>
$q_{k,t} \in \mathbb{R}_+, \quad \forall k \in \mathcal{K}_{\text{B}}, t \in \mathcal{T}$	<p>(1m) BEV charge level is always nonnegative</p>
$\kappa_k \in \mathbb{R}_+, \quad \forall k \in \mathcal{K}$	<p>(1n) per-vehicle private capital costs are nonnegative</p>
$\delta_k \in \mathbb{R}_+, \quad \forall k \in \mathcal{K}$	<p>(1o) per-vehicle manufacturing external costs are nonnegative</p>
<p>where</p> $\Delta q_{k,t}^{\text{CHG}} = q_{k,t+1} - q_{k,t} + \sum_{(i,j) \in \{\mathcal{A}: t_i=t, \Delta q_{k,i,j}^{\text{MAX}} < 0\}} a_{k,i,j} \Delta q_{k,i,j}^{\text{MAX}},$	<p>(1p) charging of each BEV at each timestep equals charge change minus routing losses</p>
$\forall k \in \mathcal{K}_{\text{B}}, t \in \mathcal{T}_{\text{Q}}$	

operations, and life cycle air emission externalities of a TNC fleet change across scenarios representing geographic and temporal variation, uncertainty, and the internalization of air emission externalities (as a Pigovian tax passed through to the fleet operator). Our model is also unique in its treatment of vehicle costs, incorporating into the optimization the effect of vehicle

usage on the period of use, resale value at end of use, and the resulting discounted future cash flow.

We include air emission externalities across the vehicle life cycle from greenhouse gas emissions (including carbon dioxide, methane, and nitrous oxide) and from criteria air pollutants (particulate matter, nitrogen and sulfur oxides, and secondary particulate matter from emissions of volatile organic com-

Table 2. Sets, Decision Variables, and Input Parameters

label	type	description
\mathcal{X}	set	all decision variables
\mathcal{V}	set	vertices representing points in space–time
\mathcal{A}	set	arcs connecting feasible pairs of vertices in \mathcal{V}
\mathcal{K}	set	vehicles or vehicle types (BEVs are represented individually, whereas CVs and HEVs are each tracked as a group)
\mathcal{K}_B	set	battery electric vehicles (subset of \mathcal{K} , indexed individually)
\mathcal{T}	set	all unique arc start and end times
\mathcal{T}_Q	set	all unique charging arc start times (subset of \mathcal{T})
Ω_k	set	linear constraints that make up the piecewise linear convex cost floor for capital cost κ_k for vehicle type k
n_k	variable	number of vehicle k purchased (BEVs are tracked individually, whereas CVs and HEVs are tracked as a group)
$a_{k,i,j}$	variable	assignment of vehicle k to arc (i, j)
$q_{k,t}$	variable	charge level of vehicle k at time t
$\Delta q_{k,t}^{\text{CHG}}$	variable	energy charged to vehicle k from the grid at time t
κ_k	variable	private acquisition cost for vehicle k
δ_k	variable	externality costs of manufacturing, disposal, and recycling emissions k
τ	parameter	flag controlling whether air emission externalities are included as a tax
r	parameter	source vertex from which all routes originate
s	parameter	sink vertex at which all routes terminate
t_i	parameter	time of vertex i
t^{START}	parameter	earliest time in \mathcal{T}
t^{END}	parameter	latest time in \mathcal{T}
$n_{i,j}$	parameter	number of trips requested along arc (i, j)
$m_{i,j}$	parameter	travel distance along arc (i, j) (annualized)
m^{MAX}	parameter	maximum lifetime travel distance of a vehicle
q_k^{MAX}	parameter	energy capacity of vehicle k (∞ for CVs and HEVs)
$c_{k,i,j}$	parameter	private cost for vehicle k to traverse arc (i, j)
$d_{k,i,j}$	parameter	external cost from vehicle k traversing arc (i, j)
c_t	parameter	private cost per kWh of electricity from the grid at time t
d_t	parameter	external cost per kWh of electricity from the grid at time t
$\Delta q_{k,i,j}^{\text{MAX}}$	parameter	maximum energy change for car k induced by travel on arc (i, j) (positive for charging arcs, negative for all others)
$\alpha_{\omega,k}^{\text{COSTS}}$	parameter	intercept term for line Ω , representing a portion of the convex piecewise linear VDT-dependent capital costs
$\beta_{\omega,k}^{\text{COSTS}}$	parameter	slope term for line Ω , representing a portion of the convex piecewise linear VDT-dependent capital costs
$\alpha_{\omega,k}^{\text{X-COSTS}}$	parameter	intercept term for line Ω , representing a portion of the convex piecewise linear VDT-dependent manufacturing external costs
$\beta_{\omega,k}^{\text{X-COSTS}}$	parameter	slope term for line Ω , representing a portion of the convex piecewise linear VDT-dependent manufacturing external costs

pounds) using reduced complexity models that estimate health costs caused by emissions of air pollutants. We use TNC trip data from Austin, Texas to represent TNC demand, but to consider how findings vary from city to city, we also model Los Angeles and New York by changing parameters related to energy prices, health costs impacts, and marginal emissions from the electric grid to represent each location. In each scenario, we find the fleet size, technology composition, and vehicle routing combination that satisfies TNC trip demand (matching origin–destination location and time) at minimum cost.

2. MATERIALS AND METHODS

We construct an optimization model to choose fleet composition (mix of CVs, HEVs, and BEVs) and operations (vehicle routing and BEV charging) to minimize the cost of satisfying exogeneous demand (origin and destination location and time) under a range of scenarios. We first describe our model and the customized methods we develop to solve it at scale, and then we describe the data that we use to instantiate the model.

2.1. Optimization Model. Figure 1 illustrates our modeling framework with an example. Vehicle purchase choices determine the vehicles available to dispatch (left). Routing options, jointly optimized with purchases, are represented using a graph, where each vertex (dot) represents a specific place and time, and the arcs connecting them include available options for:

- Trip arcs: Passenger trip requests that must be served.
- Charging arcs: Spending time parked (divided into 15 min charging increments) at a charging location while recharging a battery or waiting for the next trip.
- Dispatch arcs: Deadheading from a vehicle's home base to the first passenger trip request.
- Return arcs: Deadheading from a vehicle's final passenger trip back to its home base.
- Relocation arcs: Deadheading from the end of one passenger trip to the beginning of a next passenger trip or between passenger trips and recharge locations.

In describing our model, we first define the full mixed-integer linear programming (MILP) model FullMILP used to represent this problem, then describe a set of heuristics that we use to improve scalability.

2.2. MILP Formulation. Our FullMILP formulation, shown in Table 1, finds the cost-minimizing fleet technology mix and assignment of vehicles to trip arcs where the set of decision variables \mathcal{X} includes the number of vehicles n_k of each powertrain type k purchased, assignments $a_{k,i,j}$ of vehicles k to arcs (i, j) , charge level $q_{k,t}$ and energy charged from the grid $\Delta q_{k,t}^{\text{CHG}}$ for each vehicle k at each discrete time point t , and total annualized capital cost κ_k for each vehicle type (determined by vehicle utilization levels) for all vehicle types $k \in \mathcal{K}$, arcs

$(i, j) \in \mathcal{A}$, and times $t \in \mathcal{T}$. The full set of notation is shown in Table 2.

In all test cases, the objective function, eq 5, sums the relevant vehicle purchase costs κ_b , gasoline and per mile maintenance costs $c_{k,i,j}$ and time-varying battery charging costs c_t . In cases where air emission externalities are internalized, $\tau = 1$, so the fleet also considers a Pigovian tax on externalities from manufacturing, disposal, and recycling emissions δ_b , tailpipe and fuel refining emissions $d_{k,i,j}$ and grid emissions d_t .

At the core of FullMILP are equations that are standard for many vehicle routing problems. Constraint 6 ensures preservation of flow for each vehicle through the network (forcing vehicles to return to the depot after serving trips), Constraint 7 requires that all passenger trips be satisfied, and Constraint 8 requires that a vehicle must be purchased to be dispatched. The remainder of the formulation is customized for our case.

In our model, annualized mileage determines in what future year each vehicle is sold (either due to age or high mileage), its resale value, and the resulting discounted resale cash flow. Constraints 9–10 model capital costs and manufacturing external. For all vehicles, Constraint 9 uses a set of linear constraints Ω to define a convex piecewise linear cost floor representing the sum of annualized vehicle costs (including salvage value, which is a function of vehicle assignment) and, in relevant cases, internalized externality costs. We discuss this aspect of our formulation in more detail in the SI.

Constraints 12–14 manage the BEV charge level. Constraint 11 applies to timesteps at which regular (15 min) interval charging timesteps begin, defines charger usage, and tracks charge level changes. Constraint 12 applies to all other timesteps, at which there is no charging option, so that the charge level is fully determined by traversed arcs' energy requirements. Constraint 14 enforces bounds of BEV charge levels. The implied amount of electricity purchased from the grid is quantified for the objective function in the “where” statement as the change in charge unexplained by travel.

The set of vehicle types $k \in \mathcal{K}$ indexes individual vehicles for BEVs (each with binary purchase and routing decisions) but groups vehicles into types for CVs and HEVs (with integer purchase and routing decisions) for computational efficiency. This grouping means that FullMILP assumes refueling time and routing of CVs and HEVs is negligible, such that individually tracking fuel level is unnecessary and FullMILP need not separately index each car. Aside from these refueling implications, CV and HEV dispatch are otherwise representative of a fleet of discrete vehicles.

2.3. Heuristics. Solving the FullMILP problem with a standard commercial solver is impractical for city-scale problems with thousands of trips, particularly due to BEV charge constraints. To improve scalability, we introduce a set of customized heuristics that reduce problem size and tend to discover solutions quickly, allowing us to find near-optimal solutions to a sample of 5000 trips. This is a larger instance than commercial tools can solve for many vehicle routing problem variants in reasonable time and larger than the optimization state of the art for exact solutions (200 trips) described in the SI. We solve FullMILP first via a sequence of optimizations and heuristics:

1. A novel MCF_VaryingFleetSize heuristic reduces problem size by taking all feasible relocations from each trip to potential next trips and eliminating relocations that are

likely to be higher cost and therefore unused in optimal routing solutions. It adapts prior work²⁸ and uses MCF_CarLimit, a customization of the widely known minimum-cost network flow problem.³⁴

2. A novel ShrinkingBattery heuristic builds an initial feasible solution from an aggregated simplification of the electric subset of the vehicle fleet, iteratively making the aggregation more realistic.
3. A customized variant of a widely used RuinAndRecreate heuristic randomly selects pieces of the solution to reoptimize, improving the ShrinkingBattery solution.
4. The FullMILP formulation is executed, taking the best solution found by steps 1–3 as a starting point and upper bound on cost. It measures solution quality relative to a lower bound on cost defined by FullMILP's linear relaxation, which is iteratively tightened. In many of the cases we test, this step simply verifies that the upper bound found by steps 1–3 is within a tolerance of the solution, but in some cases, this step also improves the solution.

The MCF_VaryingFleetSize, ShrinkingBattery, and RuinAndRecreate heuristics constitute a substantial portion of this study's contribution—and this research question would be unanswerable at a meaningful scale without them—but because they are all tools to help solve the FullMILP formulation, we present their underlying intuition and algorithmic steps in the SI.

2.4. Passenger Trip Data and Driver Relocations. We instantiate the model using a dataset of 1.5 million passenger trips from June 6, 2016 to April 13, 2017, released in 2017 by RideAustin, a nonprofit ridesourcing service in Austin, Texas. We use the same set of trips from Austin to also model Los Angeles and New York City (varying private and external costs by region but not travel demand). We extract passenger trip origin and destination, starting and ending timestamps, and distance traveled to define trip arcs $n_{i,j}$. All demand must be satisfied, and passenger pickup times are inflexible. We sample down to 5000 trips using the weekday-season categories shown in Figure S17, plus a separate category for the high-demand days of the South By Southwest Festival. This sample size equates to a fleet size ranging from 37 to 39 vehicles in the base case (37–44 vehicles across all sensitivity cases), depending on the optimal technology mix. The number of trips sampled from each category is proportional to average daily demand (which increased season to season as RideAustin became more popular), and costs and distance values (which affect capital costs of each vehicle) are scaled up to annual quantities based on the number of days per year represented by each category. For tractability, we use k -means clustering to group locations into 25 clusters and round times to the nearest 5 min. Because efficiency varies with driving conditions,¹² we estimate each trip's efficiency for each powertrain type using average speed, computed from the known distance and duration, and interpolating efficiency (gallons or kWh per mile) between standard test city and highway drive cycles (EPA drive cycles with average speeds of 21.2 and 48.3 miles/h, respectively).³⁵

RideAustin data does not include travel between passenger trips. For every potential relocation from each trip to each subsequent trip (or charging node), we estimate the required distance traveled and duration using k -nearest neighbors

regression³⁶ on the RideAustin trips. This method and its implications are described further in the SI.

Relocations from the prior trip to the next trip were disallowed if the actual time gap (from the first trip's end to the second trip's start) was shorter than the estimated relocation duration or longer than 30 min. For tractability, vehicles may chain trips more than 30 min apart but must park at the central charge station depot between those trips. When the estimated duration is shorter than the time gap between trips, we assume that the vehicle travels at the estimated speed for the estimated trip duration, then idles for the remainder of the excess time (assuming the combustion engine, if applicable, is shut off using a start–stop system). For relocations between passenger trips and the charging station or the source/sink nodes, we instead assume that the vehicle parks immediately at the station and departs the station as late as possible.

2.5. Vehicle and Charger Technology. We model a typical present-day ridesourcing vehicle with otherwise-identical CV, HEV, and BEV counterparts. For model year 2018 in the United States market, there are five light-duty passenger vehicles with BEV and CV variants. Of those, the Kia Soul is best suited for ride hailing due to sufficient backseat space, so we adopt it for this study. Figure S25 shows that its efficiency and range are representative of model year 2018 BEVs excluding Teslas (likely too expensive for mass-market TNCs), the Chevrolet Bolt, and the BYD e6.

We assume one charging station (also the depot from which all vehicles must begin and end trip chains) and place it at the centroid of all trip origin–destinations. In practical contexts, optimal sizing and siting of charge capacity is a challenging problem that requires planning and investment, and it would add a great deal of complexity to our optimization. However, because our results do not show a substantial increase in VDT from BEVs routing to and from charging stations—perhaps due to perfect demand information—we do not consider sensitivity cases with more charging stations or a different charge station location. There is no capacity constraint for charging or parking at this location. The charger is the fast-charger specification (CHAdeMO) that is compatible with the Soul, which can charge its 30 kWh battery to 90% in 46 min (linearized to a rate of 35.2 kWh/h for simplicity). The BEV has an MSRP of \$33 950, a city efficiency of 27.3 kWh/100 miles, and a highway efficiency of 36.1 kWh/100 miles. The middle-trim version of the Soul CV is used, with an MSRP of \$20 500, city efficiency of 26.1 miles/gallon (mpg), and highway efficiency of 30.9 mpg. The hypothetical hybrid version of the Soul's parameters is estimated using differences in cost and efficiency between the similarly sized Kia Optima sedan's gasoline and hybrid variants, resulting in a cost of \$25 000, a city efficiency of 40.7 mpg, and a highway efficiency of 39.7 mpg. CV and BEV variant efficiencies are taken from the fueleconomy.gov;³⁷ their MSRPs were accessed from the manufacturer's product websites.

In the base case of present-day Austin, energy prices come from EIA-estimated 2017 Austin Energy annual averages for transportation sector retail electricity prices (10.90¢/kWh) and gasoline prices for 87 octane gasoline (\$2.20/gal).^{38,39} These time-invariant energy prices are shown in Table S8.

To annualize vehicle purchase costs, the MSRP minus a discounted future cash flow from resale of the vehicle (whether due to high mileage or age) is multiplied by a capital recovery factor F , as shown in eq 2

$$f_{\text{CR}}(r, N) = \frac{r(1+r)^{N-1}}{(1+r)^N - 1} \quad (2)$$

where N is the age, in years, of the vehicle at which it ceases fleet operation and is sold in the used car market (N may be a noninteger) and r is the discount rate. Note that this capital recovery factor is for equivalent annual payments from years 0 to $N - 1$ (rather than years 1 to N).

We assume that vehicles are retired from the fleet and sold in the used market after N_{MAX} years or d_{MAX} miles, whichever happens first. Given a private firm discount rate r , a vehicle purchase price p , and vehicle resale value function $v(N, d)$ that depends on age N and annual distance traveled d , the private costs of each vehicle investment are

$$\kappa = \left(1 - \frac{v(N, d)}{(1+r)^N} \right) p \times f_{\text{CR}}(r, N) \quad (3)$$

where $N = \min\left(N_{\text{MAX}}, \frac{d_{\text{MAX}}}{d}\right)$ and d is defined for each vehicle k as $\sum_{(i,j) \in \mathcal{A}} m_{i,j} a_{k,i,j}$. Here, we use the symbol κ for capital cost loosely because the MILP model treats κ as a decision variable bound below by a set of constraints that represent a piecewise linear convex function approximating eq 3. We describe this in more detail in the SI.

We assume a private firm real discount rate of 7% (with annual inflation of 2%), a maximum vehicle age of $N_{\text{MAX}} = 12$ years, and a maximum VDT of $d_{\text{MAX}} = 170\,000$ miles, based on Argonne National Laboratory's Greenhouse gases, Regulated Emissions, and Energy use in Transportation Model (GREET).⁴⁰ For each powertrain type, a separate regression (described in the SI) estimated the relationship between age, miles driven, and resale value using resale values queried from Kelley Blue Book.

2.6. Air Emission Externality Costs. In scenarios where external costs of emissions are considered, air emission externalities from the manufacturing stage are added as a Pigovian tax on vehicle investments

$$T_{\text{MFG}} = \sum_{i \in \mathcal{P}} \gamma_i^{\text{MFG}} c_i^{\text{P}} \quad (4)$$

where \mathcal{P} is the set of pollutants considered, γ_i^{MFG} is the quantity of pollutant i produced during manufacturing, and c_i^{P} is the external cost per unit of pollutant i emitted. We consider greenhouse gas emissions from CO₂, methane, and N₂O; we consider health costs from PM_{2.5}, SO_x, NO_x, and VOC.

To compute external costs per unit of greenhouse gas emissions, we adopt the social cost of carbon \$50 per ton of CO₂ equivalent estimated by the Interagency Working Group on the Social Cost of Carbon.⁴¹ For conventional air pollutants, external costs depend on emission location, and we use the AP3 model⁴² to compute and monetize estimated health costs associated with these emissions. AP3 is one of several reduced complexity models that estimate the health impacts resulting from air pollution. In contrast to the estimates generated by complex chemical transport-based air pollution models, reduced complexity models generate estimates at an acceptable level of accuracy while enabling estimates to be found for large numbers of scenarios quickly.

We adopt estimates of emissions from manufacturing each vehicle technology from GREET,⁴⁰ adjusting inputs to the modeled vehicles' curb weight and battery weight, and we assume that manufacturing emissions from each production step occur in U.S. counties where similar economic activity occurs.

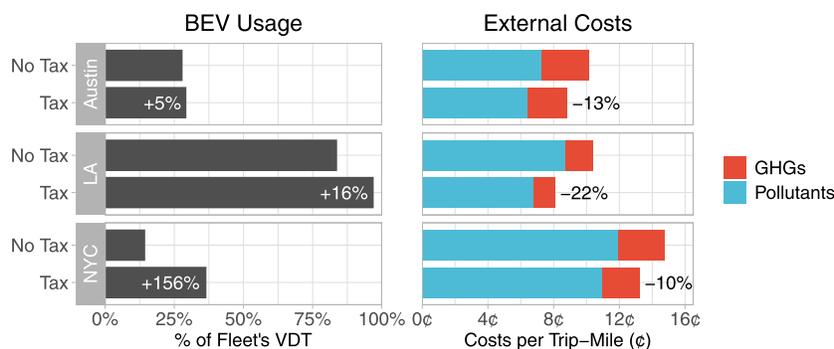


Figure 2. Summary of changes to the optimal ridesourcing fleet when air emission externalities are internalized, including share of fleet-wide vehicle-distance traveled (VDT) from BEVs (left) and total air emission externality costs per trip-mile (right) in three cities for the optimal fleet technology mix and routing to serve exogenous travel demand. Each measure's relative change induced by a Pigovian tax (expressed as a percentage of the "no-tax" case) is annotated. Assumed private and external costs of energy inputs vary by city, as described in Sections 2.5 and 2.6. All cases use a 7% real private firm discount rate, no labor costs, the vehicles described in Section 2.5 including the 2019 Kia Soul BEV, \$50/tonne CO₂ externality price, the AP3 external cost model, \$9.41 million (2018) value of statistical life, and the Pope et al.⁴⁷ concentration–response function. Results using alternative assumptions are summarized in Section 3.3.

When air emission externalities are included, p in eq 3 is the vehicle's MSRP + T_{MFG} . When externalities are excluded, p is simply the vehicle's MSRP. The SI includes further details and input values.

Air emission externalities associated with vehicle operations were estimated in a similar manner.

$$T_{\text{OP}} = \sum_{i \in \mathcal{P}} \gamma_i^{\text{OP}} c_i^{\text{P}} \quad (5)$$

As with manufacturing emissions, we use the social cost of carbon and AP3 to estimate external costs per unit of pollutant emitted from vehicle operations. We adopt GREET tailpipe and upstream estimates of emissions per gallon of gasoline consumed and compute emissions based on the fuel consumption rate of each vehicle technology on each route arc. We assign tailpipe emissions to each scenario's relevant county (Travis County, TX in the base case of Austin). For upstream emissions associated with BEV charging, there is a body of literature estimating the time-varying marginal grid emissions from end uses, including studies by Graff Zivin, Kotchen, and Mansur and Siler-Evans et al.^{43,44} that use regression approaches. We use the Siler-Evans et al. methodology, recalculated on recent data and averaged by season and hour of day in 2017 for each analysis city's eGrid subregion (ERCOT in the case of Austin).⁴⁵ For upstream emissions associated with feedstock production and transportation for making gasoline and fuel for power plants, we use emissions quantities from GREET and AP3 costs factors of refinery counties in the city's region. This approach is described further in the SI, and values used as inputs are given in Tables S2 (time-invariant gasoline externalities) and S3 (average values of time-variant grid externalities).

For emission external costs from all sources, we use the AP3 external cost model, a value of statistical life of \$9.41 million (2018), a carbon price of \$50/tonne, and the Pope et al. concentration–response relationship.⁴⁷

3. RESULTS AND DISCUSSION

Across a wide range of scenarios for three cities—Austin, Los Angeles, and New York City—we find the optimal fleet composition and operations for (1) minimizing private costs and (2) minimizing private costs plus air emission externality

costs, and we compare resulting outcomes of policy interest. The second case assumes the firm faces a Pigovian tax on direct emissions as well as other life cycle emissions passed through suppliers to the fleet operator without inducing other changes in the economy. Each test case has the same total trip miles, since demand is exogenous and must be met, and we present results per trip-mile with outcomes annualized and monetary values in 2018 USD. Costs labeled as "external" refer to life cycle air emission externalities from vehicle manufacture and use (computed with a social discount rate of 3%), and costs labeled as "social" refer to the sum of private and external costs.

In the base test case for each city, we assume a 7% real discount rate used by the fleet operator, no labor costs, a BEV price of \$33 950 (2019 Kia Soul), \$50/tonne CO₂ externality valuation, the AP3 model of conventional air pollution emission mortality effects, \$9.41 million value of statistical life, and the Pope et al.⁴⁷ air pollution concentration–response function. We use a trip dataset from Austin for all three cities, but private and external costs related to gasoline (at the tailpipe and refinery) and electricity vary across cities. These assumptions are discussed in Materials and Methods.

We first describe the impacts of a Pigovian tax on our results and assess the cost reductions possible through technology mixing. We then summarize the key results from an extensive sensitivity analysis. In the SI, we provide additional analysis of the base case results and a range of sensitivity cases.

3.1. Impact of a Pigovian Tax. Figure 2 summarizes key cost outcomes in each city when optimized with and without a Pigovian tax on air emission externalities. Across cities in our base case, with no Pigovian tax, private costs range from 45.5 to 49.0¢ and external costs range from 10.1 to 14.8¢ per trip-mile (that is, total annualized life cycle costs divided by the number of annualized miles of passenger trips served). Depending on the variability of regional costs, a tax leads the fleet to increase its usage of BEVs by 5–156% and dispatch these vehicles in a manner that reduces emission externalities per trip-mile by 10–22%. In absolute terms (x -axis of Figure 2), these reductions range from 1.3 to 2.3¢ per trip-mile. These values are broken down in greater detail in Figures S1–S4, SI.

External cost reductions are greatest in percentage and absolute terms in Los Angeles, where fuel emission externalities are high and electricity generation externalities are low relative to the other cities modeled (a larger relative difference for

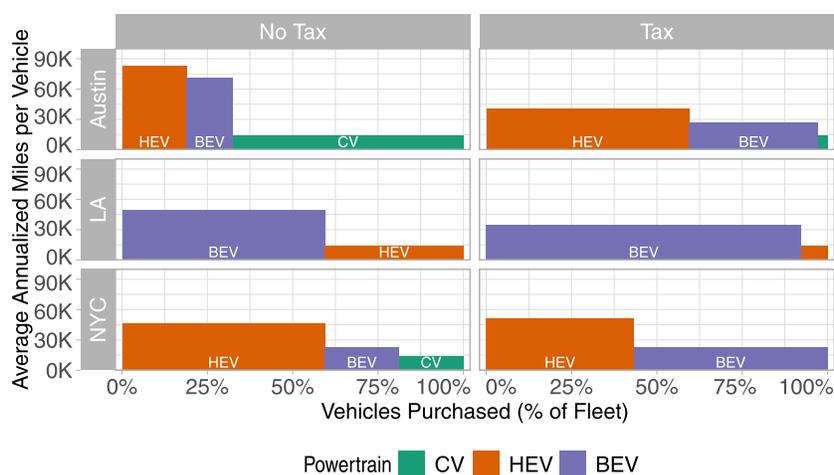


Figure 3. Vehicle purchases (*x*-axis) and average utilization (*y*-axis) by powertrain type for cost-minimizing fleets when excluding (left) and including (right) a Pigovian tax on air emissions. Private and external costs of energy inputs vary across cities, as described in Sections 2.5 and 2.6. All cases use a 7% real private firm discount rate, no labor costs, the vehicles described in Section 2.5 including the 2019 Kia Soul BEV, \$50/tonne CO₂ externality price, the AP3 air emission external cost model, \$9.41 million (2018) value of statistical life, and the Pope et al.⁴⁷ concentration–response function.

criteria pollutant external costs than for GHGs). In percentage terms, they are smallest in New York City, where the external costs of electricity generation are highest of the three cities. Austin sees the largest increase in BEV usage, partially because of lower gas prices that lead a private cost-minimizing fleet to use many CVs and few BEVs. However, due to lower health external costs per unit of tailpipe emissions and a less “clean” grid than Los Angeles, Austin’s external cost reductions fall between the other two cities’ in percentage terms and are smallest in absolute terms.

To put these per trip-mile results in context, a recent Fehr and Peers consulting report estimated that Uber and Lyft drive 104 million monthly trip miles in Los Angeles.⁴⁸ Multiplying those trip miles by the 2.3¢ per trip-mile decrease in externalities, we can roughly estimate external cost reductions of \$29 million per year in Los Angeles (\$24 million in reduced criteria pollutant emissions and the remainder in reduced GHG emissions).

As shown in Figure S2, these external cost reductions occur alongside private cost increases up to 1% in our base case (increases no higher than 0.4¢ per trip-mile). The net effect of these cost changes is a reduction in overall social costs (private costs plus external costs) ranging from 2 to 3% (0.9–2.0¢ per trip-mile). While this net effect is small in relative terms, the distributional impacts are significant since the tax shifts the fleet’s external costs away from the public, many of whom do not benefit from the fleet’s services and onto the fleet operator (and potentially its customers).

These effects are not uniform across life cycle stages. Figure S5 shows that across analysis regions, as more BEVs are used, per trip-mile manufacturing external costs increase by 6–11% (0.3–0.5¢), almost entirely due to criteria pollutants. Tailpipe and refining external costs drop by 17–80% (0.7–2.9¢) as internal combustion engines are used less. In New York City, where BEV usage is low without a Pigovian tax and the increase is largest in relative terms, grid external costs increase by a factor of four (0.9¢); in Austin and Los Angeles, where the shift is less drastic in relative terms, changes in charge scheduling offset emissions from increased grid energy usage.

These effects also vary by type of emissions. The share of per trip-mile external cost reductions attributed to reduced criteria pollutant emissions ranges from 63 to 85% (0.8–2.0¢), with the

remaining 15 to 37% coming from reduced GHG emissions (0.3–0.5¢). Criteria pollutant external cost reductions range from 8 to 11%, while greenhouse gas external cost reductions range from 16 to 19%.

These external cost reductions are accomplished in each city not only by shifting VDT away from gasoline usage (in CVs and HEVs) and toward electricity usage (in BEVs), but also by a corresponding change in vehicle purchases. Figure 3 illustrates for each city, with and without the Pigovian tax, the share of vehicle purchases for each powertrain (out of an optimal fleet size ranging from 37 to 39 vehicles in the base case) and the annual miles driven per car of each powertrain type. For all three cities, the Pigovian tax results in increased fleet electrification, both per vehicle and per mile, but the details of each city’s private-optimal and socially optimal fleets differ:

- In Austin, where gas prices are low relative to other modeled cities, a private-cost-minimizing fleet is composed of a majority of CVs, but those CVs are used infrequently, primarily during periods of high demand, while HEVs serve as “baseload” supply and are responsible for a plurality of total miles driven. When air emission externalities are internalized, the fleet uses HEVs to serve baseload and BEVs for nearly all remaining trips, almost eliminating CV usage.
- In Los Angeles, a private-cost-minimizing fleet uses no CVs due to higher gasoline prices. Instead, BEVs serve a majority of demand with HEVs used primarily in periods of high demand. Due to high gasoline externalities and low electricity externalities, a Pigovian tax results in a fleet that is almost entirely composed of BEVs.
- In New York City, where gasoline is more expensive than Austin but cheaper than Los Angeles, a private cost-minimizing fleet relies heavily on HEVs, using a mix of BEVs and CVs for high-demand periods. A Pigovian tax eliminates CVs from the fleet and makes the fleet majority BEV, but due in part to relatively high externalities of electricity generation, HEVs are still used as the baseload.

Across the three cities, the number of BEVs in the optimal fleet increases by 63–180% when a Pigovian tax is imposed on the fleet, and BEVs’ total vehicle-distance traveled increases by 5–

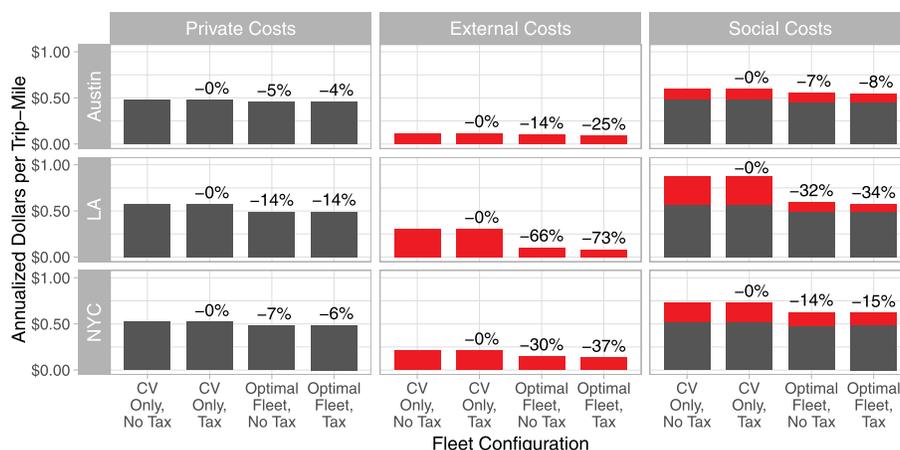


Figure 4. Average private costs, external costs from air emissions, and social costs (private + external) per trip-mile for TNC fleets in three cities, considering an all-conventional vehicle fleet (“CV Only”) and an optimally mixed fleet (“optimal fleet”) with and without a Pigovian tax on air emission externalities. The percentage cost reduction relative to the “CV only, no-tax” case is annotated. Assumed private and external costs of energy inputs vary across cities. All cases use a 7% real private firm discount rate, no labor costs, the vehicles described in Section 2.5 including the 2019 Kia Soul BEV, \$50/tonne CO₂ externality price, the AP3 external cost model, \$9.41 million (2018) value of statistical life, and the Pope et al.⁴⁷ concentration–response function.

156%. HEVs serve virtually all of the remaining demand in these three Pigovian tax cases, while CVs are at or near 0% of the fleet’s purchases and distance traveled.

3.2. Value of Optimally Mixing Technologies. Across cities, a fleet that optimally determines the mixture of powertrains to purchase and dispatch substantially reduces its private costs and the air emission externalities it produces. Figure 4 illustrates each cost component for four fleet configurations: (1) an all-CV fleet optimized for private costs with no Pigovian tax on emission externalities (the single-technology fleet option that is arguably closest to the business-as-usual case of present-day fleets); (2) the same fleet facing a Pigovian tax; (3) a mixed fleet optimized for private costs; and (4) the same fleet facing a Pigovian tax. For an all-CV fleet, internalizing emission externalities has virtually no ability to reduce them because routing decisions for CVs that minimize private fuel and capital costs also nearly minimize external costs (a very small reduction occurs because internalizing externalities alters tradeoffs between energy usage and distance-based net capital costs).

In all three cities, an all-CV fleet is suboptimal enough that many of the external cost reductions seen from a Pigovian tax are also achieved simply by having the fleet optimally choose powertrains (while still providing it with the full foresight to know what that optimal mix is). Relative to an all-CV fleet, a fleet optimized for private costs reduces private costs by 5–14% and, in doing so, also reduces emission externalities by 14–66%. In the SI, we also compare the optimal mixed fleet to optimal homogeneous fleets composed of either CVs HEVs or BEVs. Across the three cities, the best homogeneous fleet does not depend on a Pigovian tax: it is all-HEV in Austin and New York City and all-BEV in Los Angeles regardless of whether a tax is included. Relative to the best homogeneous fleet, the mixed fleet optimized without a Pigovian tax reduces private costs by 1–4% and the mixed fleet with a Pigovian tax reduces social costs by 1–4%.

Unlike all-CV and all-HEV fleets, internalizing an all-BEV fleet’s emission externalities can shift charging to lower-polluting times of the day to reduce externalities (assuming perfect day-ahead information regarding external costs of the grid’s marginal

generator). The magnitude of this reduction ranges from 4 to 6% depending on the scenario. All-BEV fleets are also slightly larger than other fleets, due to the need for some portion of the fleet to recharge during high-demand hours. These results are provided in the SI.

3.3. Sensitivity to Model Inputs. Here, we briefly summarize findings across additional test cases. In Figures S7–S16, we provide more exhaustive results from all sensitivity cases.

3.3.1. External Cost Model and Social Cost of Carbon. Our base case assumes a \$50/tonne CO₂-equivalent externality valuation and uses the AP3 reduced complexity model to estimate health costs from criteria air pollutants. Using a very high CO₂ externality valuation of \$300/tonne increases private-optimal externality estimates by a factor of roughly two to three across cities and leads the fleet to use nearly all BEVs under a Pigovian tax, reducing externalities as much as 39% (in Austin).

We also consider EASIUR and InMAP, two alternative reduced complexity models downloaded from a Center for Air, Climate, and Energy Solutions Database.^{46,49,50} Using either model rather than AP3 in Austin (where those models’ marginal grid external cost estimates are each roughly half of AP3’s) leads to substantially increased BEV uptake in Pigovian tax scenarios. In tax scenarios for Los Angeles, BEV usage is nearly maximized with all three models. In New York City, EASIUR results in similar outcomes as AP3, but InMAP leads to estimated externality reductions of 75% due to near-total electrification with a tax. This change is primarily driven by a large difference in county-level tailpipe external cost estimates for Manhattan (where InMAP’s are around 5 times greater than AP3’s and 3 times greater than EASIUR’s). These models’ structural differences, and regional variations in differences between their external cost estimates, are examined systematically in Gilmore et al.⁵¹

3.3.2. Discount Rate, Resale, and Labor. Our base case assumes that the fleet pays drivers no hourly wages (instead effectively assuming either a flat percentage of fare or driverless cars) and uses a 7% real discount rate for future operation costs and future resale value of its vehicles at the end of TNC use (resale value estimation is described in Section 2.5, eq 3).

Using a lower discount rate of 1%, the fleet places greater value on the future cash flow from reselling each car; this reduces the capital cost advantage of CVs and they are used negligibly even when minimizing private costs. With a higher discount rate of 13%, the capital cost advantage of CVs increases, but when a tax is introduced, they still serve no more than 22% of VDT and BEVs still serve 23–97% of VDT.

If we instead assume fleet vehicles have no resale value, outcomes shift slightly. This is because our resale value regression model estimates faster depreciation for BEVs than for HEVs and faster depreciation for HEVs than for CVs. This means when the resale value is removed, the effective purchase costs of CVs increase by more than HEVs or BEVs in percentage terms, but the gap is narrow in absolute dollar terms. When minimizing private costs in Austin, for example, CVs fall from 27% in the base case (7% discount rate) to 26% of total VDT (no resale value).

If we assume the fleet pays its drivers an hourly wage of \$12, including when BEVs must go out of service to recharge, BEV VDT decreases by 26 (in New York City) to 52 (in Los Angeles) percentage points in no-tax scenarios. This demonstrates that charging may not only increase planning complexity but also impose labor costs that change a fleet's optimal strategy.

3.3.3. Battery Capacity and Cost. Our base case uses the 2019 Kia Soul with a retail price of \$33 950 and a 30 kWh battery. Rather than model an explicit cost per kWh, we consider a sensitivity case in which the BEV's sticker price is reduced to \$28 950 (a cost reduction of \$167/kWh if all price reductions are attributed to lower battery costs). In that case, the private cost-minimizing fleet would be the majority BEV in each city (in terms of purchases and VDT) and a Pigovian tax would lead BEVs to serve 60–96% of VDT.

If we instead used a 2020 Chevrolet Bolt as the reference vehicle, with a price of \$36 620 and a 66 kWh battery, a tax increases electrification slightly in Austin, where the higher sticker price makes BEVs less competitive as baseload, and increases it slightly in New York City, where each BEV can serve additional trips in high-demand periods. In Los Angeles, where BEVs already served a nearly all VDT, they serve roughly the same share of VDT but require fewer purchases to do so.

3.3.4. Electricity Rates. Our base case uses average retail electricity rates from 2017. In the future, higher rates could result from large amounts of demand induced from economy-wide BEV charging, which may alter grid dispatch and require additional generating capacity. For sensitivity cases, we focused on Austin, which has the most evenly mixed fleet in the base case, and varied its electricity rates from 5.5¢ (half of the base case value) to 21.8¢ (double the base case value) per kilowatt hour. Across all price ranges, usage of CVs does not change much: they make up 68–73% of the fleet with no-tax (but primarily serve peak demand, around 30% of total VDT) and are nearly eliminated with a tax. However, the portion of the remaining demand served by BEVs varies by electricity price. When rates are doubled, they are used only in the tax scenario (for 17% of VDT versus 29% in the base case); when rates are halved, at least 68% of VDT are served by BEVs with or without a tax. Since this affects scenarios with and without a tax in similar ways, the range of reductions in external costs induced by a tax is relatively consistent, from 10% (with rates doubled) to 15% (with rates halved). Because Austin is the least favorable for BEVs in our base case, it loosely corresponds to a pessimistic lower bound on BEV usage under increased electricity rates.

3.3.5. Marginal External Costs from Electricity Generation. Our base case (described in Section 2.6) uses marginal generation estimates from 2017 for each eGrid subregion. It is unclear how external cost values from marginal generation may change in the future: it is possible that in some regions, demand induced from economywide BEV charging could shift less efficient coal plants to the margin, but it is also possible that long-term increases in renewable energy may result in some hours of the day having zero marginal emissions. Rather than explicitly model these possibilities, in each region, we consider three additional sensitivity cases with three time-invariant values for marginal external costs: the region's highest-external cost marginal emissions from 2017 (the most damaging plant is always on the margin), the region's lowest-external-cost marginal emissions from 2017 (the least damaging plant is always on the margin), and 50% of the region's lowest-external-cost marginal emissions from 2017 (on average, renewables and the least damaging plant are each on the margin half the time). For Austin's lowest-external-cost case, BEVs serve 79% of VDT with a tax; in the highest-external-cost case, they serve only 18%. External cost estimates shift with generation assumptions across tax and no-tax scenarios, and the external cost reductions induced by a tax vary in the range from 9 to 19% (0.8–1.8¢ per trip-mile). The trend is less pronounced in New York City, where BEV VDT in tax cases ranges from 31 to 42% and external cost reductions range from 7 to 15% (1.6–2.3¢). In Los Angeles, where BEVs serve a majority of VDT across all cases, external cost reductions shift even less, ranging from 17 to 22% (1.9–2.3¢).

Detailed results from sensitivity analyses, along with an expanded discussion, are available in the SI.

3.4. Discussion. Across a wide range of scenarios, our results consistently suggest that internalizing air emission externalities results in a greater degree of electrification (shift from CV to HEV and BEVs and shift from HEV to BEV) as well as operational changes that together reduce air emission externality costs (by 10–22% in the base case and 4–75% across sensitivity cases, depending on the city) and lower social costs (by 2–3% in the base case and 0–18% across sensitivity cases, depending on the city). This suggests a potential role for policy because when emission externalities are unpriced, firms have incentives to lower private cost in ways that increase air emissions, implement a lower degree of electrification, and charge BEVs when the grid is less clean than socially optimal. While the change in social cost is fairly small across most of the scenarios examined, the change in who bears the cost (private versus external costs) can be significant—as estimated above for Los Angeles, as high as \$29 million of annual environmental and health outcomes.

Pigovian taxes offer efficiency and flexibility, but in the absence of such an option, other policies that encourage similar outcomes, such as policies encouraging increased electrification, could potentially improve economic efficiency. However, any such policy should be designed with care. A blunt instrument favoring one technology over others may not be desirable because (1) the optimal fleet is generally a mixed fleet; (2) beyond fleet composition, it is important how intensively each vehicle type is used; and (3) factors that vary with location and over time, like energy prices, vehicle cost, population density, and grid emission factors, can dramatically change the degree of electrification that is optimal.

It is worth noting that unlike private vehicles, ridesourcing fleet vehicles spend a substantial portion of time deadheading with no passengers while they wait for their next ride request and

travel to its pickup location. Across test cases, even with our model's assumption of perfect information, around 47% of total distance traveled result from these empty miles (due in part to relatively low demand density in the RideAustin dataset); this implies that a similar share of external costs is due to deadheading—a largely unavoidable aspect of ridesourcing—regardless of powertrain decisions. Because demand is exogenous in our model, the Pigovian tax (and the resulting electrification) does little to reduce the degree of deadheading.

These results should be interpreted in context. Our model is relatively detailed in its treatment of supply-side investment and operation costs and constraints, but it considers a single ridesourcing firm with perfect information and full control of fleet acquisition and operation that must satisfy all demand with inflexible pickup times. In practice, current ridesourcing fleets in the U.S. are staffed by workers who choose their own vehicles, which often serve dual uses as personal vehicles and choose when to work in response to incentives. Where vehicles are purchased by each worker, it is unclear what options (e.g., a powertrain externality based driver incentive program) may be most viable for fleet owners to induce these shifts. Our model may approximate today's dispatch to the degree that accurate demand prediction is possible and to the degree that drivers respond to incentives about when to work, but we ignore the pricing mechanisms altogether, as well as the potential for dual-use vehicles. Our model may be a better approximation of a future fleet centrally owned and routed by the ridesourcing firm (e.g., a fleet of autonomous vehicles; a fleet owned and leased to drivers for TNC work) or perhaps one responding to policies requiring greater fleet-wide coordination and optimization.²¹

Our exclusion of vehicles' outside value for dual uses overestimates the extent to which CVs (which have lower capital costs) reduce costs for peak demand hours relative to BEVs and HEVs. However, this overestimation is partially mitigated by our formulation of endogenous capital costs, which lowers them for less heavily used cars. Accordingly, excluding dual uses affects the composition of vehicle purchases, particularly for low-use vehicles, but its effect on the vehicle distance traveled by each vehicle type is smaller by comparison.

Assuming perfect information and control may overestimate the fleet's ability to opportunistically schedule battery charging around gaps in demand and fluctuations in marginal grid emissions, overestimating the number of trips each BEV may be able to serve. Ridesourcing services also need not meet all demand at the exact start and end time they were served in the RideAustin data. If we allowed flexible time windows for passenger pickup and dropoff time (perhaps with a cost for additional waiting time), the fleet could improve operational efficiency and the optimal fleet composition could potentially change.⁵²

We provide results for three cities with varying private and external assumptions but use RideAustin data from 2016 to 2017 in all scenarios. Because this study does not have accurate and current inputs for each city's TNC travel demand, it misses differences resulting from the urban form that may alter the trip and relocation distances and speeds across regions (e.g., compact versus sprawled development, gridded versus irregular design, congested versus free-flowing travel, and low-speed neighborhood streets versus urban highways). This may change optimal fleets. For example, where trips are more stop-and-go, BEVs may be more optimal; where trips are longer, BEVs may face more difficult charging constraints.

We do not consider distributional impacts of ridesourcing fleet externalities, but rather optimize total social costs across disparate regions. Approaches considering equitable outcomes could, for example, require that no region may see external costs increase by more than some margin. By shifting emissions from the tailpipe to the grid, fleet electrification could increase health impacts of air emissions in areas outside the city even while reducing total air emission health impacts.¹³ Distributional impacts within a status-quo city also warrant further attention; ridesourcing fleets may be used by relatively affluent people but may impose disproportionately large health costs on populations less likely to use the services.

Despite these limitations, the ability to observe changes in optimal fleets under a variety of scenarios helps in developing intuition about fleet technology choices and operations as well as the implications of unpriced externalities in technology choice and operations.

We discuss a range of additional cases and considerations in greater depth in the SI.

■ ASSOCIATED CONTENT

SI Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.est.0c05141>.

Detailed results for test cases, summary results for additional sensitivity cases, discussion of model limitations, detailed methods for dataset preparation and estimation of uncertain inputs, and discussion of optimization methods and related literature (PDF)

■ AUTHOR INFORMATION

Corresponding Author

Jeremy J. Michalek – *Department of Engineering & Public Policy and Department of Mechanical Engineering, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, United States*; orcid.org/0000-0001-7678-8197;
Email: jmichalek@andrew.cmu.edu

Authors

Matthew B. Bruchon – *Department of Engineering & Public Policy, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, United States*; orcid.org/0000-0003-4321-1345

Inês L. Azevedo – *Department of Energy Resources Engineering, Woods Institute for the Environment, and Precourt Institute for Energy, Stanford University, Stanford, California 94305, United States*; orcid.org/0000-0002-4755-8656

Complete contact information is available at:
<https://pubs.acs.org/10.1021/acs.est.0c05141>

Notes

The authors declare no competing financial interest. Data and code that support the findings of this study are available at <https://github.com/mbruchon/FleetElectrification>.

■ ACKNOWLEDGMENTS

The authors thank Julian Marshall for providing guidance on appropriate usage and interpretation of reduced complexity air quality models. This work was supported by the Carnegie Mellon Center for Climate and Energy Decision Making through a cooperative agreement (Awards SES-0949710 and SES-1463492) between the National Science Foundation and Carnegie Mellon University. This work was also developed as

part of the Center for Clean Air Climate Solution (CACES), which was supported under Assistance Agreement no. R835873 awarded by the U.S. Environmental Protection Agency. It has not been formally reviewed by EPA. The views expressed in this document are solely those of authors and do not necessarily reflect those of the Agency. EPA does not endorse any products or commercial services mentioned in this publication. This work was also supported by the Department of Engineering & Public Policy at Carnegie Mellon University. This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number ACI-1548562. Specifically, it used the Bridges system, which is supported by NSF award number ACI-1445606, at the Pittsburgh Supercomputing Center (PSC).^{53,54}

REFERENCES

- (1) Conway, M.; Salon, D.; King, D. Trends in Taxi Use and the Advent of Ridehailing, 1995-2017: Evidence from the US National Household Travel Survey. *Urban Sci.* **2018**, *2*, No. 79.
- (2) United States Environmental Protection Agency. *Inventory of U.S. Greenhouse Gas Emissions and Sinks*; United States Environmental Protection Agency, Office of Air and Radiation, Climate Change Division, 2018; pp 1–505.
- (3) Goodkind, A. L.; Tessum, C. W.; Coggins, J. S.; Hill, J. D.; Marshall, J. D. Fine-scale damage estimates of particulate matter air pollution reveal opportunities for location-specific mitigation of emissions. *Proc. Natl. Acad. Sci. U.S.A.* **2019**, *116*, 8775–8780.
- (4) San Francisco County Transportation Authority. TNCs Today: A Profile of San Francisco Transportation Network Company Activity, 2017. <https://www.sfcta.org/emerging-mobility/tncs-today>.
- (5) *Strengthening and Implementing the Global Response*; IPCC, 2018; Chapter 4.
- (6) George, S. R.; Zafar, M. Electrifying the Ride-Sourcing Sector in California, 2018. <http://www.cpuc.ca.gov/General.aspx?id/6442457050>.
- (7) Uber. The Clean Air Plan: Helping London Go Electric, 2018. <https://www.uber.com/en-GB/blog/a-cleaner-future-for-our-cities/>.
- (8) Fulton, L.; Mason, J.; Meroux, D. *Three Revolutions in Urban Transportation*; ITDP, 2017; p 38.
- (9) Tessum, C. W.; Hill, J. D.; Marshall, J. D. Life cycle air quality impacts of conventional and alternative light-duty transportation in the United States. *Proc. Natl. Acad. Sci. U.S.A.* **2014**, *111*, 18490–18495.
- (10) Michalek, J. J.; Chester, M.; Jaramillo, P.; Samaras, C.; Shiao, C.-S. N.; Lave, L. B. Valuation of plug-in vehicle life-cycle air emissions and oil displacement benefits. *Proc. Natl. Acad. Sci. U.S.A.* **2011**, *108*, 16554–16558.
- (11) Weis, A.; Jaramillo, P.; Michalek, J. Consequential life cycle air emissions externalities for plug-in electric vehicles in the PJM interconnection. *Environ. Res. Lett.* **2016**, *11*, No. 024009.
- (12) Yuksel, T.; Tamayao, M.-A. M.; Hendrickson, C.; Azevedo, I. M. L.; Michalek, J. J. Effect of regional grid mix, driving patterns and climate on the comparative carbon footprint of gasoline and plug-in electric vehicles in the United States. *Environ. Res. Lett.* **2016**, *11*, No. 044007.
- (13) Holland, S. P.; Mansur, E. T.; Muller, N. Z.; Yates, A. J. Distributional Effects of Air Pollution from Electric Vehicle Adoption. *J. Assoc. Environ. Resour. Econ.* **2019**, *6*, S65–S94.
- (14) Holland, S. P.; Mansur, E. T.; Muller, N. Z.; Yates, A. J. Are there environmental benefits from driving electric vehicles? the importance of local factors. *Am. Econ. Rev.* **2016**, *106*, 3700–3729.
- (15) Corporation, T. M. Toyota develops “Total-care Service” Designed for Ridesourcing Companies, 2018. <https://global.toyota/en/newsroom/corporate/25737570.html>.
- (16) Lyft, Inc. Lyft Express Drive, 2020. <https://www.lyft.com/expressdrive>.
- (17) Lyft, Inc. Lyft Level 5 Self-Driving Division, 2020. <https://level5.lyft.com/>.
- (18) Uber Technologies, Inc., Uber Advanced Technologies Group, 2020. <http://www.uber.com/us/en/atg/>.
- (19) Pram, D. GM’s Cruise Shows Off Its Self-Driving Origin Shuttle with No Room for a Driver, 2020. <https://www.bloomberg.com/news/articles/2020-01-21/gm-s-cruise-to-show-first-car-built-to-operate-without-a-driver>.
- (20) Sun, Y.; Shirouzu, N. China’s Didi Sets Up Joint Ventures with BAIC Unit, Volkswagen, 2019. <https://www.reuters.com/article/us-china-didi-chuxing/chinas-didi-sets-up-joint-ventures-with-baic-unit-volkswagen-idUSKCN1PM079>.
- (21) California State Senate, Senate Bill #1014, 2018. https://leginfo.ca.gov/faces/billTextClient.xhtml?bill_id=201720180SB1014.
- (22) Bauer, G. S.; Greenblatt, J. B.; Gerke, B. F. Cost, Energy, and Environmental Impact of Automated Electric Taxi Fleets in Manhattan. *Environ. Sci. Technol.* **2018**, *52*, 4920–4928.
- (23) Bauer, G. S.; Phadke, A.; Greenblatt, J. B.; Rajagopal, D. Electrifying urban ridesourcing fleets at no added cost through efficient use of charging infrastructure. *Transp. Res. Part C: Emerging Technol.* **2019**, *105*, 385–404.
- (24) Chen, T. D.; Kockelman, K. M.; Hanna, J. P. Operations of a shared, autonomous, electric vehicle fleet: Implications of vehicle & charging infrastructure decisions. *Transp. Res. Part A: Policy Pract.* **2016**, *94*, 243–254.
- (25) Gawron, J. H.; Keoleian, G. A.; De Kleine, R. D.; Wallington, T. J.; Kim, H. C. Deep decarbonization from electrified autonomous taxi fleets: Life cycle assessment and case study in Austin, TX. *Transp. Res. Part D: Transp. Environ.* **2019**, *73*, 130–141.
- (26) Sheppard, C. J. R.; Bauer, G. S.; Gerke, B. F.; Greenblatt, J. B.; Jenn, A. T.; Gopal, A. R. Joint Optimization Scheme for the Planning and Operations of Shared Autonomous Electric Vehicle Fleets Serving Mobility on Demand. *Transp. Res. Rec.* **2019**, *2673*, 579–597.
- (27) Chen, T. D.; Kockelman, K. M. Management of a Shared Autonomous Electric Vehicle Fleet. *Transp. Res. Rec.* **2016**, *2572*, 37–46.
- (28) Bertsimas, D.; Jaillet, P.; Martin, S. Online Vehicle Routing: The Edge of Optimization in Large-Scale Applications. *Oper. Res.* **2017**, *1–48*.
- (29) Pourazarm, S.; Cassandras, C. G.; Wang, T. Optimal routing and charging of energy-limited vehicles in traffic networks. *Int. J. Robust Nonlinear Control* **2016**, *26*, 1325–1350.
- (30) Sassi, O.; Cherif, W. R.; Oulamara, A. Vehicle Routing Problem with Mixed fleet of conventional and heterogenous electric vehicles and time dependent charging costs. *Int. J. Math., Comput., Phys., Electr. Comput. Eng.* **2015**, *9*, 167–177.
- (31) Zoepf, S. M. Plug-in Vehicles and Carsharing: User Preferences, Energy Consumption and Potential for Growth. Ph.D. Thesis, Massachusetts Institute of Technology, 2015.
- (32) Pavlenko, N.; Slowik, P.; Lutsey, N. *When Does Electrifying Shared Mobility Make Economic Sense*; International Council on Clean Transportation, 2019; pp 1–15.
- (33) Slowik, P.; Wappelhorst, S.; Lutsey, N. *How Can Taxes and Fees on Ride-Hailing Fleets Steer Them to Electrify*; International Council on Clean Transportation, 2019; pp 1–15.
- (34) Klein, M. A. A Primal Method for Minimal Cost Flows with Applications to the Assignment and Transportation Problems. *Manage. Sci.* **1967**, *14*, 205–220.
- (35) U.S. Department of Energy Office of Energy Efficiency and U.S. Environmental Protection Agency. Detailed Test Information, 2018. https://www.fueleconomy.gov/feg/fe_test_schedules.shtml.
- (36) Wang, H.; Kuo, Y.-H.; Kifer, D.; Li, Z. In *A Simple Baseline for Travel Time Estimation Using Large-Scale Trip Data*, Proceedings of the 24th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems—GIS ’16. New York, New York, 2016; pp 1–4.
- (37) U.S. Environmental Protection Agency. Download Fuel Economy Data: 2018 Datafile, 2018. <https://www.fueleconomy.gov/feg/download.shtml>.

- (38) U.S. Energy Information Administration. Retail Gasoline and Diesel Prices: Regular—Conventional Areas, Annual Period, 2018. https://www.eia.gov/dnav/pet/pet_pri_gnd_a_epmru_pte_dpgal_a.htm.
- (39) U.S. Energy Information Administration. Average Retail Price of Electricity, Texas, Annual, 2018. <https://www.eia.gov/electricity/data/browser/>.
- (40) Argonne National Laboratory. Greenhouse Gases, Regulated Emissions, and Energy Use in Transportation Model, 2019. <https://greet.es.anl.gov/greet.models>.
- (41) *Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866*, Interagency Working Group on Social Cost of Greenhouse Gases; 2016, pp 1–21.
- (42) Muller, N. Z.; Mendelsohn, R. *The Air Pollution Emission Experiments and Policy Analysis Model (APEEP): Technical Appendix*; Yale University, 2006; Vol. 1.
- (43) Graff Zivin, J. S.; Kotchen, M. J.; Mansur, E. T. Spatial and temporal heterogeneity of marginal emissions: Implications for electric cars and other electricity-shifting policies. *J. Econ. Behav. Org.* **2014**, *107*, 248–268.
- (44) Siler-Evans, K.; Azevedo, I. L.; Morgan, M. G. Marginal Emissions Factors for the U.S. Electricity System. *Environ. Sci. Technol.* **2012**, *46*, 4742–4748.
- (45) Azevedo, I. L.; Donti, P. L.; Horner, N.; Siler-Evans, K.; Vaishnav, P. T. *Electricity Marginal Factors Estimates*; Center For Climate and Energy Decision Making, Carnegie Mellon University: Pittsburgh, 2021. <https://cedm.shinyapps.io/MarginalFactors/>.
- (46) Center for Air, Climate, and Energy Solutions, CACES Reduced Complexity Models Data Download. 2018, <http://airclimateenergy.org/data/>.
- (47) Pope, C. A.; Lefler, J. S.; Ezzati, M.; Higbee, J. D.; Marshall, J. D.; Kim, S. Y.; Bechle, M.; Gilliat, K. S.; Vernon, S. E.; Robinson, A. L.; Burnett, R. T. Mortality Risk and Fine Particulate Air Pollution in a Large, Representative Cohort of U.S. Adults. *Environ. Health Perspect.* **2019**, *127*, No. 077007.
- (48) Balding, M.; Whinery, T.; Leshner, E.; Womeldorff, E. *Estimated TNC Share of VMT in Six US Metropolitan Regions*, 2019.
- (49) Heo, J.; Adams, P. J.; Gao, H. O. Reduced-form modeling of public health impacts of inorganic PM_{2.5} and precursor emissions. *Atmos. Environ.* **2016**, *137*, 80–89.
- (50) Tessum, C. W.; Hill, J. D.; Marshall, J. D. InMAP: A model for air pollution interventions. *PLoS One* **2017**, *12* (4), e0176131.
- (51) Gilmore, E. A.; Heo, J.; Muller, N. Z.; Tessum, C. W.; Hill, J. D.; Marshall, J. D.; Adams, P. J. An inter-comparison of the social costs of air quality from reduced-complexity models. *Environ. Res. Lett.* **2019**, *14*, No. 074016.
- (52) Alonso-Mora, J.; Samaranayake, S.; Wallar, A.; Frazzoli, E.; Rus, D. On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment. *Proc. Natl. Acad. Sci. U.S.A.* **2017**, *114*, 462–467.
- (53) Towns, J.; Cockerill, T.; Dahan, M.; Foster, I.; Gauthier, K.; Grimshaw, A.; Hazelwood, V.; Lathrop, S.; Lifka, D.; Peterson, G. D.; Roskies, R.; Scott, J. R.; Wilkens-Diehr, N. XSEDE: Accelerating scientific discovery. *Comput. Sci. Eng.* **2014**, *16*, 62.
- (54) Nystrom, N. A.; Levine, M. J.; Roskies, R. Z.; Scott, J. In *Bridges: A Uniquely Flexible HPC Resource for New Communities and Data Analytics*, Bridges: A Uniquely Flexible HPC Resource for New Communities and Data Analytics, 2015.

BIBLIOGRAPHY

2013. [Transportation network providers and congestion in the city of chicago](#). Technical report.
- Javier Alonso-Mora, Samitha Samaranyake, Alex Wallar, Emilio Frazzoli, and Daniela Rus. 2017a. [On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment](#). *Proceedings of the National Academy of Sciences of the United States of America*, 114(3):462–467.
- Javier Alonso-Mora, Alex Wallar, and Daniela Rus. 2017b. [Predictive routing for autonomous mobility-on-demand systems with ride-sharing](#). In *IEEE International Conference on Intelligent Robots and Systems*, volume 2017-Septe, pages 3583–3590. Institute of Electrical and Electronics Engineers Inc.
- Argonne National Laboratory. 2019. [Greenhouse gases, Regulated Emissions, and Energy use in Transportation Model](#).
- Inês L. Azevedo, Priya L Donti, Nathaniel Horner, Kyle Siler-evans, and Parth T Vaishnav. [Electricity Marginal Factors Estimates](#). *Center for Climate and Energy Decision Making*. Pittsburgh: Carnegie Mellon University. [online].
- Yash Babar and Gordon Burtch. 2020. [Examining the Heterogeneous Impact of Ride-Hailing Services on Public Transit Use](#). *Information Systems Research*, 31(3):820–834.
- Melissa Balding, Teresa Whinery, Eleanor Leshner, and Eric Womeldorff. 2019. Estimated TNC Share of VMT in Six US Metropolitan Regions.
- Badi H. Baltagi, Lung-fei Lee, and Jihai Yu. 2015. [Spatial Panel Data Models](#). *The Oxford Handbook of Panel Data*.
- Rounaq Basu and Joseph Ferreira. 2021. [Planning car-lite neighborhoods: Does bikesharing reduce auto-dependence?](#) *Transportation Research Part D: Transport and Environment*, 92:102721.

- Gordon S. Bauer, Jeffery B. Greenblatt, and Brian F. Gerke. 2018. [Cost, Energy, and Environmental Impact of Automated Electric Taxi Fleets in Manhattan](#). *Environmental Science and Technology*, 52(8):4920–4928.
- Gordon S. Bauer, Amol Phadke, Jeffery B. Greenblatt, and Deepak Rajagopal. 2019. [Electrifying urban ridesourcing fleets at no added cost through efficient use of charging infrastructure](#). *Transportation Research Part C: Emerging Technologies*, 105:385–404.
- Dimitris Bertsimas, Patrick Jaillet, and Sebastien Martin. 2017. Online Vehicle Routing : The Edge of Optimization in Large-Scale Applications. *Operations Research*, pages 1–48.
- Geoff Boeing. 2020. [Global Urban Street Networks Node/Edge Lists](#).
- Matthew B. Bruchon, Jeremy J. Michalek, and Inês L. Azevedo. 2020. [Effects of Air Emission Externalities on Optimal Ridesourcing Fleet Electrification and Operations](#). *Environmental Science & Technology*, 55(5):3188–3200.
- California Air Resources Board. 2021. SB 1014 Clean Miles Standard 2018 Base-year Emissions Inventory Report. *Technical report*.
- California State Senate. 2018. [Senate bill #1014](#).
- André L.S. Chagas, Carlos R. Azzoni, and Alexandre N. Almeida. 2016. [A spatial difference-in-differences analysis of the impact of sugarcane production on respiratory diseases](#). *Regional Science and Urban Economics*, 59:24–36.
- T. Donna Chen and Kara M. Kockelman. 2016. [Management of a Shared Autonomous Electric Vehicle Fleet](#). *Transportation Research Record: Journal of the Transportation Research Board*, 2572:37–46.
- T. Donna Chen, Kara M. Kockelman, and Josiah P. Hanna. 2016. [Operations of a shared, autonomous, electric vehicle fleet: Implications of vehicle & charging infrastructure decisions](#). *Transportation Research Part A: Policy and Practice*, 94:243–254.
- City and County of San Francisco Treasurer and tax Collector. [Traffic congestion mitigation tax \(tcm\)](#).
- City of Chicago. 2021a. [Chicago Traffic Tracker - Historical Congestion Estimates by Segment - 2018-Current](#).

- City of Chicago. 2021b. [CTA - Ridership - Bus Routes - Daily Totals by Route](#).
- City of Chicago. 2021c. [Taxi Trips](#).
- City of Chicago. 2021d. [Transportation Network Providers - Trips](#).
- City of Chicago. 2021e. [Transportation Network Providers - Vehicles](#).
- City of Chicago Business Affairs and Consumer Protection. 2020. [City of Chicago Congestion Pricing](#).
- City of Chicago Open Data Portal Team. 2019. [How Chicago Protects Privacy in TNP and Taxi Open Data](#).
- Dave Colon. 2021. [Revealed: Early congestion pricing study shows congestion pricing works](#).
- Junia Compostella, Lewis M. Fulton, Robert De Kleine, Hyung Chul Kim, and Timothy J. Wallington. 2020. [Near- \(2020\) and long-term \(2030–2035\) costs of automated, electrified, and shared mobility in the United States](#). *Transport Policy*, 85:54–66.
- Matthew Conway, Deborah Salon, and David King. 2018. [Trends in Taxi Use and the Advent of Ridehailing, 1995–2017: Evidence from the US National Household Travel Survey](#). *Urban Science*, 2(3):79.
- Jean Dubé, Diègo Legros, Marius Thériault, and François Des Rosiers. 2014. [A spatial difference-in-differences estimator to evaluate the effect of change in public mass transit systems on house prices](#). *Transportation Research Part B: Methodological*, 64:24–40.
- J. Paul Elhorst. 2014. [Spatial Panel Data Models](#). pages 37–93.
- Lung fei Lee and Jihai Yu. 2010. [Estimation of spatial autoregressive panel data models with fixed effects](#). *Journal of Econometrics*, 154(2):165–185.
- James H Gawron, Gregory A Keoleian, Robert D De Kleine, Timothy J Wallington, and Hyung Chul Kim. 2019. [Deep decarbonization from electrified autonomous taxi fleets: Life cycle assessment and case study in Austin, TX](#). *Transportation Research Part D: Transport and Environment*, 73:130–141.
- Simi Rose George and Marzia Zafar. 2018. [Electrifying the Ride-Sourcing Sector in California](#).
- Michael Gold. 2021. [Congestion pricing is coming to new york. everyone has an opinion](#).

- Jonathan D. Hall, Craig Palsson, and Joseph Price. 2018. [Is Uber a substitute or complement for public transit?](#) *Journal of Urban Economics*, 108:36–50.
- Robert Cornelius Hampshire, Chris Simek, Tayo Fabusuyi, and Xi Chen. 2017. [Measuring the Impact of an Unanticipated Suspension of Ride-Sourcing in Austin, Texas.](#) *SSRN Electronic Journal*.
- Alejandro Henao and Wesley E. Marshall. 2019. [The impact of ride-hailing on vehicle miles traveled.](#) *Transportation*, 46(6):2173–2194.
- Jinhyok Heo, Peter J. Adams, and H. Oliver Gao. 2016. [Public Health Costs of Primary PM_{2.5} and Inorganic PM_{2.5} Precursor Emissions in the United States.](#) *Environmental Science and Technology*, 50(11):6061–6070.
- Stephen P. Holland, Erin T. Mansur, Nicholas Z. Muller, and Andrew J. Yates. 2016. [Are there environmental benefits from driving electric vehicles? the importance of local factors.](#) *American Economic Review*, 106(12):3700–3729.
- Stephen P. Holland, Erin T. Mansur, Nicholas Z. Muller, and Andrew J. Yates. 2019. [Distributional effects of air pollution from electric vehicle adoption.](#) *Journal of the Association of Environmental and Resource Economists*, 6(S1):S65–S94.
- IPCC. 2018. [Chapter 4: Strengthening and Implementing the Global Response.](#) Technical Report June 2018.
- Shuqing Kang, Aupal Mondal, Aarti C. Bhat, and Chandra R. Bhat. 2021. [Pooled versus private ride-hailing: A joint revealed and stated preference analysis recognizing psycho-social factors.](#) *Transportation Research Part C: Emerging Technologies*, 124:102906.
- George Karypis and Vipin Kumar. 1999. [A fast and high quality multilevel scheme for partitioning irregular graphs.](#) *Siam Journal on Scientific Computing*, 20:359–392.
- Morton Klein. 1967. [A Primal Method for Minimal Cost Flows with Applications to the Assignment and Transportation Problems.](#) *Management Science*, 14(3):205–220.
- Patrícia S. Lavieri and Chandra R. Bhat. 2019. [Modeling individuals’ willingness to share trips with strangers in an autonomous vehicle future.](#) *Transportation Research Part A: Policy and Practice*, 124:242–261.
- James P. LeSage and James P. 2008. [An Introduction to Spatial Econometrics.](#) <http://journals.openedition.org/rei>, 123(123):19–44.

- Sen Li, Kameshwar Poolla, and Pravin Varaiya. 2021a. [Impact of congestion charge and minimum wage on TNCs: A case study for San Francisco](#). *Transportation Research Part A: Policy and Practice*, 148:237–261.
- Sen Li, Hai Yang, Kameshwar Poolla, and Pravin Varaiya. 2021b. [Spatial pricing in ride-sourcing markets under a congestion charge](#). *Transportation Research Part B: Methodological*, 152:18–45.
- Yang Liu, Prateek Bansal, Ricardo Daziano, and Samitha Samaranyake. 2019. [A framework to integrate mode choice in the design of mobility-on-demand systems](#). *Transportation Research Part C: Emerging Technologies*, 105:648–665.
- Lyft, Inc. 2020a. [Lyft express drive](#).
- Lyft, Inc. 2020b. [Lyft Level 5 Self-Driving Division](#).
- Negar Mehr and Roberto Horowitz. 2019. [Pricing traffic networks with mixed vehicle autonomy](#). *Proceedings of the American Control Conference*, 2019-July:2676–2682.
- Jeremy J Michalek, Mikhail Chester, Paulina Jaramillo, Constantine Samaras, Ching-Shin Norman Shiau, and Lester B Lave. 2011. [Valuation of plug-in vehicle life-cycle air emissions and oil displacement benefits](#). *Proceedings of the National Academy of Sciences of the United States of America*, 108(40):16554–8.
- Giovanni Millo and Gianfranco Piras. 2012. [Journal of Statistical Software splm: Spatial Panel Data Models in R](#). *Journal of Statistical Software*, 47(1).
- Nicholas Z. Muller. 2014. [Boosting gdp growth by accounting for the environment](#). *Science*, 345(6199):873–874.
- National Highway Traffic Safety Administration. 2021. [Fatality Analysis Reporting System 2019 National Statistics](#).
- New York City Taxi & Limousine Commission. [New york state’s congestion surcharge](#).
- Nicholas A. Nystrom, Michael J. Levine, Ralph Z. Roskies, and J Scott. 2015. [Bridges: a uniquely flexible HPC resource for new communities and data analytics](#). *Proceedings of the 2015 XSEDE Conference: Scientific Advancements Enabled by Enhanced Cyberinfrastructure*, 1.
- Nikita Pavlenko, Peter Slowik, and Nic Lutsey. 2019. [When does electrifying shared mobility make economic sense?](#) Technical Report January, ICCT.

- C. Arden Pope, Jacob S. Lefler, Majid Ezzati, Joshua D. Higbee, Julian D. Marshall, Sun Young Kim, Matthew Bechle, Kurtis S. Gilliat, Spencer E. Vernon, Allen L. Robinson, and Richard T. Burnett. 2019. [Mortality Risk and Fine Particulate Air Pollution in a Large, Representative Cohort of U.S. Adults](#). *Environmental health perspectives*, 127(7):77007.
- Sepideh Pourazarm, Christos G. Cassandras, and Tao Wang. 2016. [Optimal routing and charging of energy-limited vehicles in traffic networks](#). *International Journal of Robust and Nonlinear Control*, 26(6):1325–1350.
- PurpleAir. 2021. [Purple Air Data Download](#).
- Xinwu Qian, Tian Lei, Jiawei Xue, Zengxiang Lei, and Satish V. Ukkusuri. 2020. [Impact of transportation network companies on urban congestion: Evidence from large-scale trajectory data](#). *Sustainable Cities and Society*, 55:102053.
- Feng Qiu and Qingmeng Tong. 2021. [A spatial difference-in-differences approach to evaluate the impact of light rail transit on property values](#). *Economic Modelling*, 99:105496.
- Lisa Rayle, Danielle Dai, Nelson Chan, Robert Cervero, and Susan Shaheen. 2016. [Just a better taxi? A survey-based comparison of taxis, transit, and ridesourcing services in San Francisco](#). *Transport Policy*, 45:168–178.
- A. P. Reynolds, G. Richards, B. De La Iglesia, and V. J. Rayward-Smith. 2006. [Clustering Rules: A Comparison of Partitioning and Hierarchical Clustering Algorithms](#). *Journal of Mathematical Modelling and Algorithms 2006 5:4*, 5(4):475–504.
- Bouayad Agha Salima, Le Gallo Julie, and Védrine Lionel. 2018. [Spatial econometrics on panel data](#). *Handbook of Spatial Analysis: Theory and Application with R*, (7):179–203.
- Paolo Santi, Giovanni Resta, Michael Szell, Stanislav Sobolevsky, Steven H Strogatz, and Carlo Ratti. 2014. [Quantifying the benefits of vehicle pooling with shareability networks](#). *Proceedings of the National Academy of Sciences of the United States of America*, 111(37):13290–4.
- Ons Sassi, Wahiba Ramdane Cherif, and Ammar Oulamara. 2015. [Vehicle Routing Problem with Mixed fleet of conventional and heterogenous electric vehicles and time dependent charging costs](#). *International Journal of Mathematical, Computational, Physical, Electrical and Computer Engineering*, 9(3):167–177.
- Bruce Schaller. 2021. [Can sharing a ride make for less traffic? Evidence from Uber and Lyft and implications for cities](#). *Transport Policy*, 102:1–10.

- David Schrank, Luke Albert, Bill Eisele, and Tim Lomax.
- Erich Schubert, Andreas Koos, Tobias Emrich, Andreas Züfle, Klaus Arthur Schmid, and Arthur Zimek. 2020. [Same-size k-means variation](#).
- Joseph P Schwieterman. 2019. [Uber Economics: Evaluating the Monetary and Travel Time Trade-Offs of Transportation Network Companies and Transit Service in Chicago, Illinois](#). *Transportation Research Record*, 2673(4):295–304.
- Colin J. R. Sheppard, Gordon S. Bauer, Brian F. Gerke, Jeffery B. Greenblatt, Alan T. Jenn, and Anand R. Gopal. 2019. [Joint Optimization Scheme for the Planning and Operations of Shared Autonomous Electric Vehicle Fleets Serving Mobility on Demand](#). *Transportation Research Record: Journal of the Transportation Research Board*, 2673(6):579–597.
- Kyle Siler-Evans, Inês Lima Azevedo, and M. Granger Morgan. 2012. [Marginal Emissions Factors for the U.S. Electricity System](#). *Environmental Science & Technology*, 46(9):4742–4748.
- Andrea Simonetto, Julien Monteil, and Claudio Gambella. 2019. [Real-time city-scale ridesharing via linear assignment problems](#). *Transportation Research Part C: Emerging Technologies*, 101:208–232.
- Michele D. Simoni, Kara M. Kockelman, Krishna M. Gurusurthy, and Joschka Bischoff. 2019. [Congestion pricing in a world of self-driving vehicles: An analysis of different strategies in alternative future scenarios](#). *Transportation Research Part C: Emerging Technologies*, 98:167–185.
- Peter Slowik, Sandra Wappelhorst, and Nic Lutsey. 2020. [How can taxes and fees on ride-hailing fleets steer them to electrify?](#) Technical Report January, ICCT.
- Matthias Studer. 2013. [WeightedCluster Library Manual: A practical guide to creating typologies of trajectories in the social sciences with R](#). Technical report, LIVES Working Papers 24.
- Han Su, Shuncheng Liu, Bolong Zheng, Xiaofang Zhou, and Kai Zheng. 2020. [A survey of trajectory distance measures and performance evaluation](#). *VLDB Journal*, 29(1):3–32.
- Yilei Sun and Norihiko Shirouzu. 2019. [China’s Didi sets up joint ventures with BAIC unit, Volkswagen](#).
- R. Tachet, O. Sagarra, P. Santi, G. Resta, M. Szell, S. H. Strogatz, and C. Ratti. 2017. [Scaling law of urban ride sharing](#). *Scientific Reports*, 7.

- Christopher W Tessum, Jason D Hill, and Julian D Marshall. 2014. [Life cycle air quality impacts of conventional and alternative light-duty transportation in the United States](#). *Proceedings of the National Academy of Sciences of the United States of America*, 111(52):18490–5.
- Christopher W. Tessum, Jason D. Hill, and Julian D. Marshall. 2017. [InMAP: A model for air pollution interventions](#). *PLoS ONE*, 12(4):e0176131.
- John Towns, Timothy Cockerill, Maytal Dahan, Ian Foster, Kelly Gaither, Andrew Grimshaw, Victor Hazlewood, Scott Lathrop, Dave Lifka, Gregory D. Peterson, Ralph Roskies, J. Ray Scott, and Nancy Wilkens-Diehr. 2014. [XSEDE: Accelerating scientific discovery](#). *Computing in Science and Engineering*, 1.
- Toyota Motor Corporation. 2018. [Toyota develops “total-care service” designed for ride-hailing companies](#).
- Uber. 2018. [The Clean Air Plan: Helping London go electric](#).
- Uber Technologies, Inc. 2020. [Uber Advanced Technologies Group](#).
- United States Environmental Protection Agency, Office of Air and Radiation, Climate Change Division. 2018. [Inventory of U.S. greenhouse gas emissions and sinks](#). Technical report.
- United States Government. 2016. [Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866](#). *Interagency Working Group on Social Cost of Greenhouse Gases*, (August):1–21.
- US Bureau of Labor Statistics. 2020. [Occupational Employment and Wages in Chicago-Naperville-Elgin – May 2019](#).
- US Department of Energy Office of Energy Efficiency and US Environmental Protection Agency. 2018. [Detailed Test Information](#).
- U.S. Department of Transportation. 2016. [Revised departmental guidance on valuation of travel time in economic analysis](#). Technical report.
- US Energy Information Administration. 2018a. [Average retail price of electricity, Texas, annual](#).
- US Energy Information Administration. 2018b. [Retail Gasoline and Diesel Prices: Regular - Conventional Areas, Annual Period](#).

- Alex Wallar, Javier Alonso-Mora, and Daniela Rus. 2019a. [Optimizing vehicle distributions and fleet sizes for shared mobility-on-demand](#). In *Proceedings - IEEE International Conference on Robotics and Automation*, volume 2019-May, pages 3853–3859. Institute of Electrical and Electronics Engineers Inc.
- Alex Wallar, Wilko Schwarting, Javier Alonso-Mora, and Daniela Rus. 2019b. [Optimizing Multi-class Fleet Compositions for Shared Mobility-as-a-Service](#). In *2019 IEEE Intelligent Transportation Systems Conference, ITSC 2019*, pages 2998–3005. IEEE.
- Hongjian Wang, Yu-Hsuan Kuo, Daniel Kifer, and Zhenhui Li. 2016. [A simple baseline for travel time estimation using large-scale trip data](#). In *Proceedings of the 24th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems - GIS '16*, pages 1–4, New York, New York, USA. ACM Press.
- Jacob W. Ward, Jeremy J. Michalek, Inês L. Azevedo, Constantine Samaras, and Pedro Ferreira. 2019. [Effects of on-demand ridesourcing on vehicle ownership, fuel consumption, vehicle miles traveled, and emissions per capita in U.S. States](#). *Transportation Research Part C: Emerging Technologies*, 108:289–301.
- Jacob W. Ward, Jeremy J. Michalek, and Constantine Samaras. 2021a. [Air Pollution, Greenhouse Gas, and Traffic Externality Benefits and Costs of Shifting Private Vehicle Travel to Ridesourcing Services](#). *Environmental Science and Technology*, 55(19):13174–13185.
- Jacob W. Ward, Jeremy J. Michalek, Constantine Samaras, Inês L. Azevedo, Alejandro Henao, Clement Rames, and Tom Wenzel. 2021b. [The impact of Uber and Lyft on vehicle ownership, fuel economy, and transit across U.S. cities](#). *iScience*, 24(1):101933.
- Allison Weis, Paulina Jaramillo, and Jeremy Michalek. 2016. [Consequential life cycle air emissions externalities for plug-in electric vehicles in the PJM interconnection](#). *Environmental Research Letters*, 11(2):024009.
- David Welch. 2020. [Gm’s cruise shows off its self-driving origin shuttle with no room for a driver](#).
- D.H. Wolpert and W.G. Macready. 1997. [No free lunch theorems for optimization](#). *IEEE Transactions on Evolutionary Computation*, 1(1):67–82.
- X. Yan, Jonathan Levine, and X. Zhao. 2019. [Integrating ridesourcing services with public transit: An evaluation of traveler responses combining revealed and stated preference data](#). *Transportation Research Part C: Emerging Technologies*, 105:683–696.

Tugce Yuksel, Mili-Ann M Tamayao, Chris Hendrickson, Inês M L Azevedo, and Jeremy J Michalek. 2016. [Effect of regional grid mix, driving patterns and climate on the comparative carbon footprint of gasoline and plug-in electric vehicles in the United States](#). *Environmental Research Letters*, 11(4):044007.

Najib Zgheib, Maya Abou-Zeid, and Isam Kaysi. 2020. [Modeling demand for ridesourcing as feeder for high capacity mass transit systems with an application to the planned Beirut BRT](#). *Transportation Research Part A: Policy and Practice*, 138:70–91.

Stephen M Zoepf. 2015. [Plug-in Vehicles and Carsharing: User Preferences, Energy Consumption and Potential for Growth](#). Technical report.

Haojia Zuo. 2013. Ridesharing. <https://github.com/MetaZuo/RideSharing>.