

System-level impact and behavior of coordinated vehicle fleets in transportation networks

*Submitted in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy
in
the Department of Civil and Environmental Engineering*

Matthew Battifarano

B.A., Mathematics, University of Chicago
M.S., Advanced Infrastructure Systems, Carnegie Mellon University
M.S., Machine Learning, Carnegie Mellon University

Carnegie Mellon University
Pittsburgh, PA 15213
December 2022

To my wife.

Acknowledgments

I would like to thank my advisor Sean Qian for our frequent and extensive discussions, for allowing me to pursue interesting ideas, and for helping me craft ideas into papers. His questions and advice helped me to find the story in each one of my projects and made me a better researcher along the way. I would like to thank members of the Mobility Data Analytics Center past and present, particularly Wei Ma, Xidong Pi, Arnav Choudhry, Weiran Yao, Pablo Gaurda, and Rick Grahm, for the many insightful discussions over the years and for listening to me present these ideas many times over and still asking good questions. I would also like to thank my family for supporting me the entire way and always asking how my research or coursework was progressing, my mom Kathleen Robinson, my dad Ernest Battifarano, my sisters Oriana Battifarano and Jessica Robinson-Papp and her family, Alex Papp and Sophia and James. I would also like to thank my in-laws, Dave and Tammy Bogema and Stephanie Bogema, for welcoming me into their family and always expressing interest in my projects. Last, but not at all the least, I would like to thank my wonderful wife, Meghan Bogema, who not only offered me constant support and encouragement throughout my PhD, but had the distinct privilege of frequently hearing me bemoan my lack of progress. She had the resilience to remind me, each and every time, that I would see it through. I would not be writing this if not for all of you.

This research was funded in part by National Science Foundation Awards CMMI-1751448 and CMMI-1931827, the Pennsylvania Department of Community and Economic Development (DCED), and Mobility 21, a national University Transportation Center on mobility funded by U.S. Department of Transportation. I would like to thank Gridwise Inc. for providing consultative resources for this research. I would also like to extend thanks to Professors Glenn Shafer, Ruodu Wang, and Aaditya Ramdas for their course which introduced me to the statistical ideas leveraged in the final chapter of this thesis and to Ojash Neopane for several helpful conversations.

Finally I would like to thank my thesis committee, Sean Qian¹², Chair, Matteo Pozzi¹, Corey Harper¹² and Peter Zhang².

¹Department of Civil and Environmental Engineering, Carnegie Mellon University

²Heinz School of Information Systems and Public Policy, Carnegie Mellon University

Abstract

Over the past decade, travelers have increasingly utilized technology to aid their use of transportation systems, from real-time navigation apps like Google Maps, to ride-sourcing apps like Uber and Lyft, to bike and scooter-share apps like Lime and Spin. These applications of technology to mobility systems have fundamentally changed the way in which people travel by offering them the ability to navigate the physical world by first navigating a virtual representation of it. The layers of virtual and physical transportation infrastructure complicates transportation planning efforts and underscores the need for methodologies to understand and leverage the relationship between the virtual and physical transportation infrastructures.

This dissertation offers concrete methodologies to understand the behavior and impact of technologically-enabled mobility services on transportation infrastructure along three directions: 1) to *illuminate* the coordinating signals in virtual mobility services, 2) to *anticipate* benefits and harms of coordinating signals in general physical transportation networks, and 3) to *recover* coordinating signals from observed behavior on the physical network. These methodologies together expand the toolbox of the transportation planner to not only accommodate but leverage technological change in realizing the transportation system of the future.

In the first part of this thesis, we demonstrate that spatial-temporal imbalances in supply and demand on ride-sourcing platforms can be predicted from real-time urban data. This work highlights surge pricing as an imperfect market mechanism and also demonstrates that ride-sourcing supply and demand behavior in the aggregate can be recovered from data. In the second part of the thesis, we study how coordination within a virtual transportation infrastructure impacts travelers on the physical road network. In the best case, this coordination induces the system optimal network utilization, but in other cases may reduce network efficiency. Finally, we study how we might infer the coordinating signal within a virtual transportation network from observed behavior of its users on the physical network, even as the environment is changing. To this end we apply a novel statistical test to driver behavior during the COVID-19 lockdown in Pittsburgh. This research may also be applied to assess the validity of network equilibrium models based on a particular notion of travel cost to describe the long-run behavior of a network from data collected on that network over time. Taken together this thesis outlines a methodology for transportation planners to recover coordinating signals from data and leverage them in long-term transportation plans to improve transportation systems.

Contents

1	Introduction	1
1.1	Premise and motivation	2
1.1.1	Enabling technologies	2
1.1.2	Implications	4
1.2	Problem statement	5
1.3	Contributions	7
1.4	Thesis organization	10
2	Background and prior work	11
2.1	Transportation network companies	11
2.1.1	Overview	11
2.1.2	Surge pricing	11
2.1.3	Congestion	14
2.1.4	Research gaps	14
2.2	Transportation network equilibrium theory	15
2.2.1	Overview	15
2.2.2	Assumptions and implications	17
2.2.3	Stochastic user equilibrium	18
2.2.4	Equilibrium of mixed traffic	18
2.2.5	Research gaps	20
2.3	Predictive modeling in transportation systems	22
2.4	Choice modeling in transportation systems	24
2.4.1	Overview	24
2.4.2	Inverse reinforcement learning	25
2.4.3	Choice modeling and transportation network companies	26
2.4.4	Research gaps	26
2.5	Sequential hypothesis testing with non-negative super-martingales	27
2.5.1	Overview	27
2.5.2	Ville's inequality	29
2.5.3	E-processes	29

3	What drives a ride-sourcing fleet?: real-time prediction of supply-demand imbalances in ride-sourcing networks	31
3.1	Introduction	31
3.2	Descriptive analysis	35
3.2.1	Surge multipliers	35
3.2.2	Traffic speed measurements	37
3.2.3	Road closures	38
3.3	Methodology	38
3.3.1	Data collection	41
3.3.2	Data processing	41
3.3.3	Temporal segmentation	43
3.3.4	Clustering	44
3.3.5	A log-linear model with L_1 regularization	44
3.4	Results	46
3.4.1	Model performance on Uber surge multipliers	47
3.4.2	Model performance on Lyft surge multipliers	50
3.4.3	Selected features	52
3.4.4	Parameter clusters	57
3.5	Policy implications	60
3.6	Conclusion	62
4	On the impact of optimized fleets in transportation networks	65
4.1	Introduction	65
4.2	Preliminaries	69
4.3	Fleet optimality paradox	71
4.3.1	Total delay under mixed equilibrium with fleets	71
4.3.2	An example of the fleet optimality paradox	72
4.3.3	System optimality paradox	74
4.3.4	Concluding remarks	75
4.4	Critical fleet size	75
4.4.1	Critical fleet size on a parallel network	76
4.4.2	Critical fleet size as a mathematical program with equilibrium constraints	78
4.4.3	Critical fleet size heuristic solution algorithm	79
4.4.4	Exact critical fleet size solution algorithm via mixed integer programming	80
4.4.5	Critical fleet size bounds via linear programming	82
4.4.6	Solution algorithm via column generation	84
4.5	Experiments	84
4.5.1	Experimental setting and main results	86
4.5.2	Link flow patterns at CFS-SO	87

4.5.3	Fleet O-D demand patterns at CFS-SO	89
4.5.4	Characteristics of paths and O-D pairs with fleet flow at CFS-SO . .	90
4.5.5	Sensitivity of CFS-SO and CFS-UE to aggregate demand	91
4.6	Discussion	94
4.6.1	CFS-SO as optimization-by-proxy	95
4.6.2	Applications of CFS-SO and CFS-UE	97
4.7	Regularized critical fleet size	97
4.7.1	Demand-regularized critical fleet size	98
4.7.2	Maximum likelihood critical fleet size	98
4.8	Demographic analysis of critical fleet size	100
4.8.1	Data and methods	101
4.8.2	Results	102
4.9	Conclusions and future work	104
5	Learning from learners: behavioral inference from non-stationary poli-	
	cies	107
5.1	Introduction	107
5.2	Methods	109
5.2.1	Inferring service-related stops from GPS trajectories	109
5.2.2	Non-stationary binary sequence testing	110
5.3	Data	113
5.3.1	Ride-sourcing driver GPS trajectories	113
5.3.2	OpenStreetMaps feature tags	114
5.4	TNC driver GPS analysis	115
5.4.1	OpenStreetMap analysis	116
5.5	Experiments	119
5.5.1	Simulated data	119
5.5.2	Ride-hail driver GPS trajectories	121
5.6	Results	122
5.6.1	Inferred service stops	122
5.6.2	Statistical hypothesis test results	122
5.7	Discussion	125
5.8	Conclusion	134
5.9	Testing network equilibrium	135
6	Conclusion	139
6.1	Future work	140
A	On the impact of optimized fleets Appendices	143
A.1	Proof of Equation (4.3)	143
A.2	System optimality paradox	144

A.3	Proofs of critical fleet size on a parallel network145
A.3.1	Proof of Proposition 1145
A.3.2	Proof of Proposition 2145
A.3.3	Proof of Proposition 3146
A.3.4	Proof of Proposition 4146
A.4	Critical fleet size heuristic solution algorithm via sensitivity analysis . .	.148
A.5	Proof of the MCR lower bound of CFS-SO152
A.6	Proof of Proposition 5152
A.7	Critical fleet size with multiple fleets153
A.8	Path set generation154
A.9	First order optimality conditions for the critical fleet size upper bound linear program155
B	Learning from learners Appendices	159
B.1	Extended summary of test results159
B.2	Increasing probability as the null hypothesis159
	References	163

List of Figures

1.1	Conceptual diagram of this dissertation	6
3.1	Empirical distribution of surge multipliers	35
3.2	Entropy of Uber and Lyft surge multipliers	36
3.3	Temporal correlation of surge multipliers	37
3.4	Graphical representation of the model	38
3.5	Surge multiplier entropy by time of day	43
3.6	Surge multiplier time-of-day clusters	45
3.7	Model advantage by location	48
3.8	Comparison of model error at three locations	49
3.9	Predicted surge multipliers at a single location	51
3.10	Feature importance by type and prediction horizon	53
3.11	Feature selection over time	54
3.12	Importance of event features	56
3.13	Importance of speed features	56
3.14	Importance of past surge multipliers	57

3.15	Parameter clusters	59
4.1	Conceptual diagram of total cost at mixed equilibrium with fleets	68
4.2	Fleet optimality paradox: an illustrative toy network.	72
4.3	Fleet and user flow at CFS-SO (Sioux Falls)	87
4.4	Fleet and user flow at CFS-SO (Pittsburgh)	88
4.5	Empirical CDF of fleet trip flow	88
4.6	Fleet origin-destination demand CDF	89
4.7	Fleet origin-destination demand concentration	90
4.8	Path independence factor drives fleet concentration	92
4.9	Critical fleet size sensitivity to demand level	93
4.10	Fleet volume per lane mile by race-poverty group	103
5.1	Hourly driver GPS activity 2019-2020	115
5.2	Average fraction of TNC driver interactions with tourism-related tags by month in 2019 and 2020	117
5.3	Average fraction of TNC driver interactions with passenger and delivery- related tags by month in 2019 and 2020	118
5.4	E-processes for 10 i.i.d. Bernoulli sequences	120
5.5	E-process for Bernoulli bandit algorithms	121
5.6	Kernel density map of service stops	123
5.7	Average tag incidence over time for the top 10 significant tag-seeking driver-tag pairs	126
5.8	Average tag incidence over time for the top 10 significant tag-avoiding driver-tag pairs	127
5.9	E-processes for an increasing and decreasing sequence	128
5.10	Residential seekers and avoiders	130
5.11	Tag clouds for tag-seeking driver clusters	132
5.12	Equilibrium testing procedure results.	138

List of Tables

3.1	Results of the linear regression of land use data onto model advantage. . .	48
4.1	Fleet, user, and, aggregate demand for the fleet optimality paradox	73
4.2	Demonstration of the fleet optimality paradox	73
4.3	Demonstration that the fleet decreases its total cost at FO	74

4.4	Comparison of critical fleet size on experiment networks.	86
4.5	Coordination discount and average cost by flow class, solution concept, and network	94
4.6	Maximum Likelihood CFS Results	101
5.1	Number and percent of GPS pings by activity type	113
5.2	Bonferroni-corrected p-values for tags from driver-tag pairs exhibiting tag-seeking behavior	124
5.3	Bonferroni-corrected p-values for tags from driver-tag pairs exhibiting tag-avoiding behavior.	125
A.1	Demonstration of the SO paradox	145
A.2	Cases for aggregate, user, and fleet link flow	147
B.1	Summary of test results	161
B.2	Bonferroni-corrected p-values for significant tag-seeking behavior. . . .	162
B.3	Bonferroni-corrected p-values for significant tag-avoiding behavior. . . .	162

1

Introduction

Traffic patterns in transportation networks today arise from aggregate behaviors of heterogeneous network users competing and cooperating with each other to meet their needs. Over the past decade, travelers have increasingly utilized technology to aid their use of transportation systems, from real-time navigation apps like Google Maps, to ride-sourcing apps like Uber and Lyft, to bike and scooter-share apps like Lime and Spin. As technology has increasingly shaped transportation network usage, regulators, legislators, and public advocates have worked to structure a regulatory environment that balances technological change and preservation of public rights of way in the immediate future. Transportation planners, in contrast, have longer horizons in mind; they are tasked with shaping the transportation system that will be in use over a time span of years if not decades. How then should transportation planners account for rapid technological change in transportation? How then can multi-year or decade transportation plans responsibly accommodate not only what future travelers will desire of their transportation system, but what technology will ask of it?

This dissertation distills technological change in transportation to a foundational concept that will underlie further technological change: the coordination of vehicles and travelers via a technological layer we refer to as a *virtual transportation infrastructure*. This layering of virtual and physical transportation infrastructure, and in particular, the *coordinating signals* virtual transportation infrastructure facilitates, complicates transportation planning efforts, which have historically focused on the physical infrastructure. Understanding the relationship between the virtual and physical transportation infrastructures, and moreover how their interaction can produce beneficial or harmful effects on transportation networks is the fundamental question facing transportation planning. This dissertation offers some answers in the form of concrete and novel methodologies to *illuminate* coordinating signals in ride-sourcing networks, *anticipate* benefits and harms of coordinating signals in general transportation networks, and *recover* coordinating signals from observed behavior on the phys-

ical network. Taken together, these methodologies expand the toolbox of the transportation planner to not only accommodate but leverage technological change in realizing the transportation system of the future.

1.1 Premise and motivation

This thesis is premised on the notion that the application of computer technology to mobility and logistics services has fundamentally changed the way transportation networks are utilized. This thesis is motivated by the need for rigorous mathematical and computational frameworks to understand how transportation infrastructure and policy should rise to meet the occasion.

1.1.1 Enabling technologies

Four interrelated technological trends have enabled and driven a revolution within transportation systems, and arguably, within society more broadly.

1. high-speed, high-bandwidth computer networking,
2. available and affordable computational power,
3. small, mobile computing devices, and,
4. computational methods to store, query, and extract patterns from large datasets.

As a motivating example consider the following cycle: location-based data is collected from mobile devices (3), transmitted over the internet (1), stored in large-scale databases, and distilled into actionable insights about urban mobility (2 and 4) which then may be communicated back to the mobile devices over the internet. This process occurs within Transportation Network Companies (TNCs) like Uber and Lyft but also navigation providers like Waze, Google Maps, and INRIX, public transportation operators, on-demand delivery providers like DoorDash and UberEats, bike and scooter share networks like Spin and Lime, and even EZ-Pass, among others.

The interrelation among these enabling technologies and their interaction with the world is best understood through the lens of the “metaverse.” Before the “metaverse” became associated with Facebook, it was a (much more interesting) conceptual framework for understanding “the convergence of 1) virtually enhanced physical reality and 2) physically persistent virtual space” (Smart, Cascio, and Paffendorf 2007). The term “metaverse” originated in Neal Stephenson’s 1992 science fiction novel “Snow Crash,” where it referred to a persistent, shared, 3D virtual world (Stephenson 1992), a 3D internet of sorts. However, the metaverse as a conceptual framework is broader and more abstract, encompassing technological advances in four inter-related categories: virtual worlds, mirror worlds, augmented reality, and life-logging (Smart, Cascio, and

Paffendorf 2007; Koster 2017). Most relevant to this discussion are mirror worlds which are “informationally-enhanced virtual models [...] of the physical world,” and augmented reality which “enhance[s] the external physical world for the individual through [...] interfaces that process and layer networked information on top of our everyday perception of the world” (Smart, Cascio, and Paffendorf 2007). What mirror worlds and augmented reality share is the ability of the user to navigate the physical world via its virtual representation. Augmented reality further allows the user to *interact with* the physical world via its virtual representation. The most meaningful technological shifts in transportation have been driven by what we call *virtual mobility services* (or simply *mobility services*) in these two categories: road and public transit navigation apps like Google Maps, Waze, INRIX and Transit are examples of mirror worlds and ride-sourcing, delivery, and other mobility apps such as Uber¹, Lyft, DoorDash, and Spin, as well as infrastructure-related apps like parking payment apps, or even EZ-Pass are examples of augmented reality.

The mobility services enabled by the layering of a virtual infrastructure on the physical world may also be understood as a *cyber-physical-social system* (Sheth, Anantharam, and Henson 2013), a system of mutually interacting computers, objects, and people. The conceptual origins of cyber-physical-social systems may be traced at least as far back as J. C. R. Licklider (1960) who argued that a “[hu]man-computer symbiosis” would produce results greater than either human or machine could accomplish alone. J. C. R. Licklider (1960) differentiates his “very close coupling” of human and computer (what we might term a cyber-social or cyber-human system) from both a “mechanically extended [hu]man” and “artificial intelligence,” by stressing the interactive communication that must occur between the human and the computer. In 1962 J.C.R. Licklider became the first head of Computer Research at the Advanced Research Projects Agency (ARPA) where he outlined much of what we now recognize as the internet, including cloud computing, in a series of memos (Leiner et al. 1997). His work would directly lead to the creation of ARPAnet, the precursor to today’s internet. In particular, J. C. Licklider and R. W. Taylor (1968) imagined what we now recognize as an internet-enabled cyber-social system wherein a network of computers would mediate and enhance communication among a group of people.² The cyber-physical-social systems of today (Xiong et al. 2015) expand on Licklider’s vision to include the physical systems that affect and are affected by a cyber-social system. Nevertheless, cyber-physical-social systems are fundamentally understood in much the same way as

¹The Uber app was officially launched at the 2010 San Francisco “AppShow” alongside, in a moment of historical synchronicity, Neal Stephenson’s app-based fictional narrative “The Mongoliad” (Terdiman 2010).

²This view was epitomized in “the mother of all demos” given in 1968 by Douglas Engelbart, a collaborator of Licklider’s, showcasing many firsts in computing including screen-sharing and video conferencing over the internet, interactive computing, hypermedia, and the mouse (SRI International 1968).

Licklider understood the human-computer symbiosis: each component of the cyber-physical-social system operates in very close coupling with the other two, allowing behavior to emerge that is not fully explainable by understanding any single component in isolation. Historically, transportation planning has implicitly worked with a physical-social model: human travel desires are mediated by the physical infrastructure and other traveler behavior within it. Technologically enabled mobility services add a “cyber-social” dimension. To return to our motivating example, the cyber component acts as the virtual transportation infrastructure by prompting humans to act, processing their response, and coordinating mobility services within the physical infrastructure.

Both the metaverse and the cyber-physical-social systems frameworks illuminate the fundamental nature of technology-driven changes to transportation systems: the layering of physical transportation infrastructure with virtual transportation infrastructure. People were hiring drivers, renting bikes, boarding buses, and reading maps to navigate their physical world long before digital computers even existed. What *has* changed is the ability to navigate the physical world by first navigating a virtual representation of it.

1.1.2 Implications

Virtual transportation infrastructure makes certain uses of the physical infrastructure easier and less costly which has led to substantially broader adoption of these uses. Many things are easier to do virtually. It is easier to match someone who wants transportation with someone who is able to provide it if both are logged into the same virtual platform. It is even easier if the virtual platform anticipates when and where matches might need to be made. As a result, use of the virtual infrastructure to travel via for-hire drivers has exploded. By the end of 2016, Uber and Lyft vehicles collectively accounted for over 500 million vehicle-miles traveled (VMT) *each month* in the United States (Hensley, Padhi, and Salazar 2017), by 2019, Uber and Lyft accounted for nearly 14% of daily VMT within San Francisco’s urban core (Balding et al. 2019). Similarly, it is easier to use a bike-share network if both the bike and the rider are represented in the same virtual platform. As a result, bike-share and scooter-share networks have grown rapidly in cities world-wide. In 2019, 136 million trips were completed on a bike or scooter share network in the United States, a 60% increase from the year before (National Association of City Transportation Officials 2020). Finally, it is easier for a driver to take local “backroads” when they have confidence it will actually be faster than the main roads they might be more familiar with. As a result, many residential roads have become more heavily used during periods of congestion on nearby arterials (Macfarlane 2019). The layering of virtual and physical transportation infrastructures has and will continue to change how physical infrastructure is used in ways that cannot be understood without understanding the virtual

infrastructure layered on top of it. As a result, planners and policymakers that hope to improve a transportation system must consider both the virtual and physical transportation infrastructures holistically.

Responding to both changes in transportation network utilization and public advocacy, regulators have taken an interest in redefining how the current transportation network and infrastructure can be used, by those using virtual infrastructure or otherwise. For transportation and urban planners, as well as industry leaders and public advocates, however, the question is not how should the rules of infrastructure use be changed to mitigate harm, but rather, how should the infrastructure itself be developed in order to provide equitable and expanded levels of service *in light of* the virtual transportation infrastructure? What transportation systems do we want to have? and how should they be used? These questions cannot be answered without considering both the physical and virtual transportation infrastructures together and in relation to one another.

1.2 Problem statement

If transportation and urban planners, along with industry leaders and public advocates, are to render their vision of the future of transportation systems into infrastructure and policy proposals, they will need tools to understand the extent to which their proposed changes bends the use of transportation systems toward their vision. The study of transportation planning and policy in systems *without* a virtual layer is mature: there are many tools in the transportation or urban planner's toolbox to address problems in transportation systems without a virtual layer. In a layered physical and virtual transportation system, the efficacy of the existing tools is uncertain and the set of tools capable of understanding transportation systems as *cyber*-physical-social systems is small and underdeveloped. The primary task facing today's planners and policymakers is to understand how proposed policy and infrastructure will be used by those interacting with one or more virtual transportation infrastructures. What is currently lacking is a robust set of tools available to planners that are capable of understanding the behavior and impact of cyber-physical-social systems within physical transportation networks and how those same systems might respond to infrastructure or policy changes.

The perspective this dissertation takes of transportation systems in general, and the scope of this dissertation in particular, is shown graphically in Figure 1.1. The left side of Figure 1.1 shows the cyber-physical-social systems view of transportation: multiple virtual systems and multiple human systems all mutually interacting in the physical world. For example a ride-sourcing driver may use both Uber and Lyft to find passengers but also Waze to navigate the road network. A ride-sourcing service itself may interact with multiple other virtual services including Google Maps for traffic

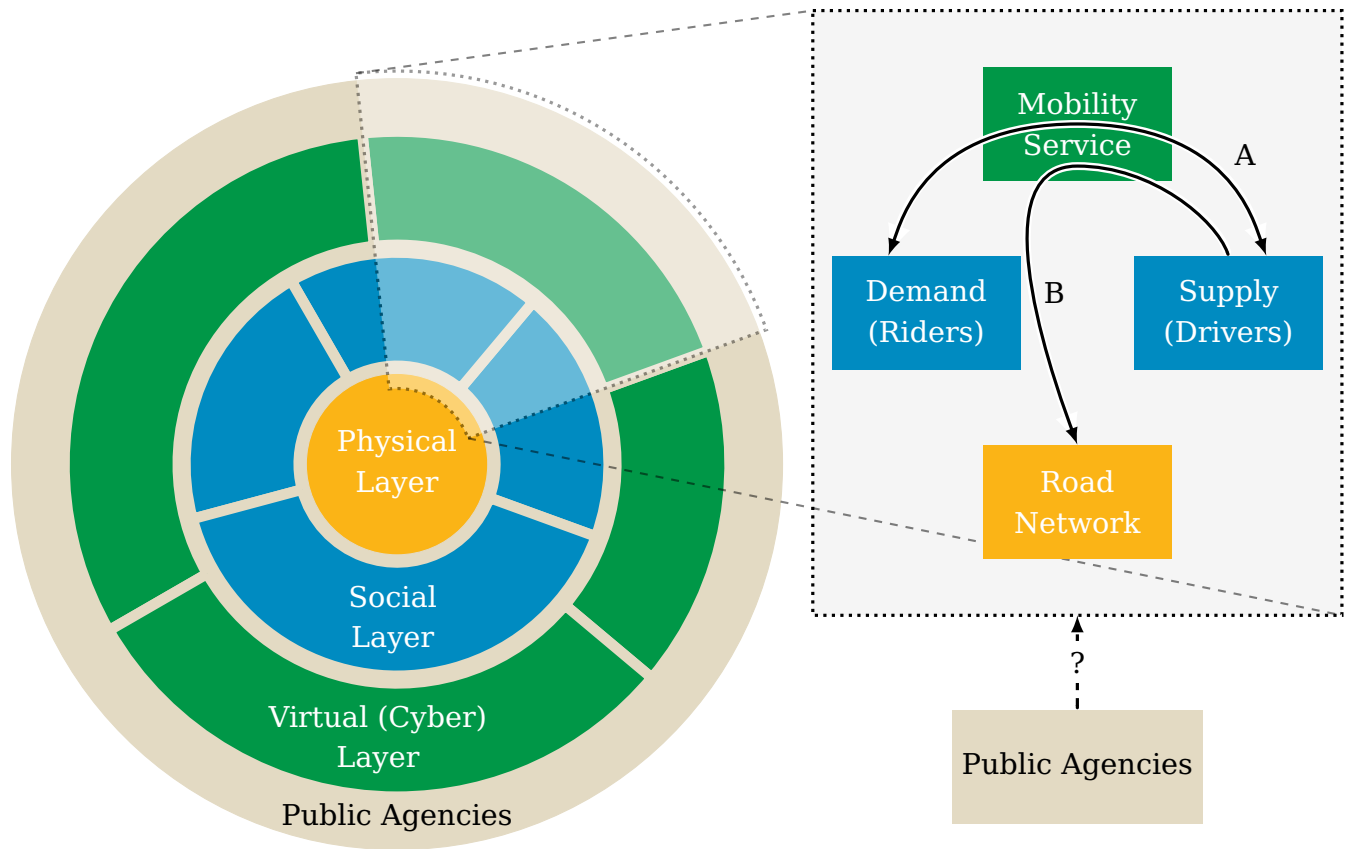


Figure 1.1: Conceptual diagram of this dissertation

estimation or a Department of Transportation traffic incident service. The physical layer is inherently a single entity (the world) but it nevertheless may be viewed as a super-position of various subsystems, the boundaries of which must be defined. The right side of Figure 1.1 shows the scope of this dissertation, which focuses on a “slice” of the cyber-physical-social systems at work within transportation.

The focus of this dissertation is to understand two relationships mediated by a mobility service, marked with arrows A and B in Figure 1.1 and in doing so provide tools for transportation planners and future researchers to determine how best to position public agencies to mediate the relationships between the social, cyber, and physical layers. Relationship A describes how the mobility service mediates the matching between transportation demand and transportation supply. Relationship B describes how the mobility service mediates the supply’s use of the road network. In this dissertation we frequently consider a ride-sourcing service like Uber or Lyft as the mobility service, however relationships A and B appear also in local delivery services like Uber Eats or DoorDash, shared micro-mobility services like Lime and Spin, and (relationship B only) in navigation apps like Google Maps and Waze, among others.

There are certainly other relationships not considered: the mobility service collects data about the road network mediated by the supply and the demand, the demand's use of the road network is mediated by the mobility service and the supply, *et cetera*. Similarly, there are other relevant cyber, social, and physical systems not shown. Certainly, competing mobility services interact with overlapping user bases, other social systems influence transportation demand itself, and physical spaces like the curb and parking are impacted by the use of mobility services. These are all important components that are outside of the scope of this work.

Concretely, the following questions will be addressed within this dissertation;

1. How can we relate the physical and virtual transportation networks?
 - (a) When ride-sourcing fleets are coordinated to reduce supply-demand imbalance via surge-pricing, is it possible to anticipate the imbalance from multi-source urban data?
 - (b) How do characteristics of the physical transportation network mediate the relationship between real-time urban data and supply-demand imbalance?
2. How can we anticipate the effects a virtual transportation network might have on the physical network?
 - (a) When some vehicles on the road network are coordinated by a virtual transportation network, do they benefit or harm network efficiency?
 - (b) Under what conditions could coordination lead to optimal network efficiency?
 - (c) How can transportation planners anticipate fleet coordination effects in network design?
3. How can we discover coordinating signals in a virtual transportation network from behavior on the physical network?
 - (a) When agents are *still learning* to navigate the virtual transportation network in order to maximize their utility, how much can we infer about the utility function they may be attempting to maximize?

1.3 Contributions

This dissertation offers concrete methodologies to understand the behavior and impact of technologically-enabled mobility services on transportation infrastructure along three directions stemming from the use of virtual transportation infrastructure to coordinate use of the physical network:

- *illuminate* the coordinating signals in virtual mobility services,
- *anticipate* benefits and harms of coordinating signals in general physical transportation networks, and

- *recover* coordinating signals from observed behavior on the physical network.

In the first part of this thesis, we demonstrate that spatial-temporal imbalances in supply and demand on ride-sourcing platforms can be predicted from real-time urban data. This work was published in Transportation Research Part C (Battifarano and Z. S. Qian 2019). This work highlights surge pricing as an imperfect market mechanism and also demonstrates that ride-sourcing supply and demand behavior in the aggregate can be recovered from data. What this work also highlights is that ride-sourcing fleets utilize the road network in an idiosyncratic way: what drives supply movement in ride-sourcing fleets is different than the underlying travel demand pattern. The fleets are, in fact, coordinated in that they are responding to a common signal from a shared source: they are moving not only through the physical transportation infrastructure but the virtual transportation infrastructure as well. However, TNCs are not the only entities to centrally coordinate fleets via virtual transportation infrastructure. Logistics services have a long history of careful optimization of their fleets for package delivery. More recently, services such as Google Maps, Waze Carpool, and others offer complex real-time operations within a virtual transportation infrastructure with real impacts on the physical transportation infrastructure.

In the second part of the thesis, we study how the presence of a coordinating signal within a virtual transportation infrastructure impacts travelers, both those who use the virtual network and those who do not, on the physical road network. This work was submitted and is under review in Transportation Science. In pursuit of this we examine a network equilibrium of a mixture of travel behavior on the road network between independently operated individual road users and a fleet of vehicles coordinated by a virtual transportation network. In this setting, we view the virtual transportation network as one that wishes to coordinate the route choices of its members to minimize the total travel cost over its users. Under certain conditions, this mixture of behavior can result in an improvement of network efficiency, but it may also worsen congestion on the network. In the best case, this coordination induces the system optimal network utilization, which implies that optimization in the virtual transportation network and physical transportation network are equivalent even though the virtual transportation infrastructure optimizes over only *some* of the vehicles. We introduce a novel mathematical program to efficiently solve for this state. Taken together this research effort offers the theory of equilibrium with fleets as a valuable analysis tool to both the public and private sectors.

In the third part, we study how we might infer the coordinating signal within a virtual transportation network from observed behavior of its users on the physical network. This work has been submitted to Transportation Research Part C. In pursuit of this goal, we develop techniques to establish when decision-making agents are still learning to balance exploration and exploitation from observed behavior. To achieve this, a novel statistical testing methodology is developed and applied to ride-sourcing

driver behavior during the COVID-19 lockdown in Pittsburgh, a time in which even experienced drivers were faced with a new and uncertain environment. This technique has broad application both within transportation and outside of it. Outside of transportation, it offers a new perspective on inverse reinforcement learning and an alternative to choice modeling when the behavioral data is collected from individuals over time interacting with complex environments. It also provides a real-time change-point detection method with probabilistic guarantees. Within transportation in particular, it lays the groundwork for the ability to assess the validity of network equilibrium models based on a particular notion of travel cost to describe the long-run behavior of a network from data collected on that network over time. We aim to demonstrate that equilibrium theory can provide insight into the preferences of individuals and organizations that are currently using the transportation infrastructure and enable a principled way to subsequently estimate how that usage would adapt in response to network or regulatory changes, insight that is as valuable to transportation planners as it is to business strategy teams.

Concretely, the contributions of this dissertation are summarized as follows;

1. Descriptive analyses of two important and novel data sets:
 - (a) Real-time surge prices across Pittsburgh over time, and
 - (b) GPS trajectories of ride-sourcing drivers during the first nine months of the COVID-19 pandemic lockdown and for the same period of the previous year.
2. A novel real-time predictive model of supply-demand imbalance in ride-sourcing systems leveraging real-time multi-source urban data.
3. Two novel, and particularly important, notions of equilibria of mixed traffic between coordinated fleets and individual users are introduced and examined: critical fleet size for user equilibrium and critical fleet size for system optimal. Efficient numerical algorithms are developed to solve them.
4. A novel non-parametric, non-asymptotic sequential hypothesis testing framework is developed to identify when a sequence of observed rewards may have been generated by an agent who is in the process of learning the reward signal.
5. Open source implementations of original methodology are made available online via:
 - (a) <https://gitlab.com/mbattifarano/surge-prediction/-/tree/classification-regression> (Chapter 3)
 - (b) <https://github.com/mbattifarano/traffic-equilibrium/tree/mb-cfs> (Chapter 4)
 - (c) <https://github.com/mbattifarano/spad/tree/mb-autoencoder> (Chapter 5)

1.4 Thesis organization

This dissertation is organized as follows.

In Chapter 2 we give a detailed background of the ideas and prior work on which this dissertation builds. We first give a concise history of transportation network companies and their disputed role in traffic congestion. We then review transportation network equilibrium theory and relevant extensions. We discuss in detail the dynamic pricing mechanism commonly referred to as “surge pricing,” along with prior academic research into its effects. We briefly discuss predictive modeling, with particular attention to predictive modeling of transportation systems and TNCs in particular. We then review choice modeling with emphasis on applications to transportation research, drawing connections to inverse reinforcement learning and network equilibrium theory. Finally, we discuss some relatively recent advances in sequential statistical hypothesis testing.

In Chapter 3 we demonstrate a real-time predictive model of supply-demand imbalances based on a data set that combines surge pricing data from Uber and Lyft with real-time urban data drawn from several disparate sources. A descriptive and comparative analysis is performed on the surge pricing data which highlights patterns within and between surge pricing on the two platforms.

In Chapter 4 a mixed equilibrium of coordinated (optimized) fleets and individual users is examined in theory and on the Pittsburgh road network. Two particularly important mixed equilibria are identified and a novel algorithm is introduced to solve for them. It is demonstrated that system optimal network flow can be achieved with a combination of fleets and individual users, but that fleets may alternatively worsen congestion. Implications for the design and planning of road networks are discussed.

Finally, in Chapter 5 we address the problem of testing the hypothesis that an agent is learning to optimize a given reward function. A theoretical analysis is grounded in multi-armed bandits and sequential hypothesis testing. The technique is demonstrated on simulated data and on data collected from Pittsburgh ride-sourcing drivers after the COVID-19 lockdown in March of 2020. Broader applications to online change-point detection, continuous monitoring, and network equilibrium theory are discussed.

2

Background and prior work

2.1 Transportation network companies

2.1.1 Overview

Transportation network companies (TNC) refer to private companies that leverage internet-based technologies to offer automated on-demand transportation and logistics services (Moran and Lasley 2017; Clewlow and Mishra 2017). As a regulatory term “Transportation Network Company” was first used by the California Public Utilities Commission (CPUC) in 2013 when they created the category to “apply to companies that provide prearranged transportation services for compensation using an online-enabled application (app) or platform to connect passengers with drivers using their personal vehicles” (California Public Utilities Commission 2013) which at the time explicitly included “Lyft, SideCar, and UberX.” The CPUC found that TNCs fell under their jurisdiction as “charter party passenger carriers.” The passenger mobility services offered by TNCs are referred to as “ride-sourcing” (Rayle et al. 2014) or “ride-hailing” (AP Stylebook 2015). TNCs have since expanded to offerings in shared and micro-mobility, public transportation, freight transportation, and local delivery (Uber Technologies 2021). In this thesis we use the term TNC in the broadest sense to refer to companies that coordinate transportation or mobility services via digital platforms, even if they do not always (or ever) carry passengers.

2.1.2 Surge pricing

Within the passenger mobility segments of TNCs, a dynamic pricing strategy commonly known as “surge pricing” is used to balance supply and demand. When the number of requests for service in a location exceeds the number of nearby vehicles,

a so-called *surge multiplier* is introduced. By increasing the fare, TNCs intend to encourage riders to postpone their trip and reward drivers who relocate to serve the demand.

Initially introduced to Uber’s platform in 2012, surge pricing was aimed at increasing the percentage of requests for rides that are fulfilled, known as the *completion rate* (Gurley 2014). The magnitude of the surge multiplier is driven by the wait times experienced by the users, which in turn is driven by the rate at which demand is out-pacing supply (Gurley 2014; J. Hall, Kendrick, and Nosko 2015). Uber displays the surge multipliers to drivers on an in-app heatmap with the intent of “nudging” them from areas of low demand into areas of high demand (Rosenblat and Stark 2016). Cohen et al. (2016) find that for major cities in the United States, Uber is surging between 14% and 28% of the time. Internal estimates from Uber place this figure at around 10% platform-wide (Gurley 2014). M. K. Chen and Sheldon (2016) characterize Uber surge multipliers; in particular, that Uber surge multipliers are bounded between 1.2 and a per-market maximum and are discrete in increments of 0.1. Further, there is a maximum rate of change for surge multipliers.

Later analysis has suggested that in addition to direct measurements of their own system, Uber is using predictive analytics to anticipate future surges in demand and shortages of drivers (Guda and Subramanian 2017; Rosenblat and Stark 2016; Phillips 2017). In addition to historical platform data, Uber’s demand forecasting relies on social-physical data such as weather, local events, and traffic conditions (Guda and Subramanian 2017). While it is known that Uber uses demand forecasts to provide notifications to drivers, it is unclear whether they are also used to compute the surge multiplier displayed to drivers (Guda and Subramanian 2017; Rosenblat and Stark 2016).

In order to function properly as a signal to re-allocate vehicle supply, surge multipliers provided to each driver should be predictive. That is, the surge multipliers provided to the driver should reflect the surge the driver should expect to see at a location when they arrive should they travel there from their current location. Drivers frequently report arriving to a surge area only to find that the surge has ended (Guda and Subramanian 2017). In fact, one of the most prominent pieces of advice offered by experienced drivers to those new to the platform is “don’t chase the surge” (Rosenblat and Stark 2016). It has been shown that drivers routinely ignore the surge multiplier and other market information provided by Uber. In one study, surge pricing in certain areas of New York City and San Francisco caused drivers to *avoid* the area entirely (L. Chen, Mislove, and Wilson 2015). This behavior is likely caused by driver’s widespread mistrust of the market information provided by Uber (M. K. Lee et al. 2015; Rosenblat and Stark 2016). This mistrust is not unfounded: in certain circumstances the platform operator has an incentive to mis-report the surge multiplier (Guda and Subramanian 2017).

It should be noted that in contrast, several studies have shown that surge pricing is an effective mechanism to increase supply. J. Hall, Kendrick, and Nosko (2015) examine a case study of Uber surging around a specific well-attended event. When surge pricing was in effect, supply rose to meet demand and completion rate remained high. In comparison, when the surge pricing system experienced an outage on New Year’s Eve, completion rate fell dramatically during the outage. Further, M. K. Chen and Sheldon (2016) show that drivers remain on the platform longer and accept more rides when surge pricing is in effect.

Uber has since moved away from explicit fare multipliers. Since 2018, Uber has included the surge multiplier in the total cost displayed to the rider instead of listing the base fare and surge multiplier separately (Uber Technologies 2018). Also starting in 2018, Uber began employing “flat surge pricing” in some markets, offering a fixed dollar bonus instead of a multiple of the base fare (Avedian 2018). Flat surge pricing was met with mixed reviews from drivers, requiring different strategies to maximize earnings during periods of surge (Avedian 2018). In response to a changing regulatory environment, Uber has abandoned flat surge pricing in favor of surge multipliers in California (Uber Blog 2020), indicating that flat surge multipliers may not last as a demand and supply management tool.

Theoretical models of surge multipliers

To our best knowledge, surge multipliers have previously only been modeled from a theoretical perspective. Guda and Subramanian (2017) develop a two-region economic model to examine driver behavior under surge multipliers from a game theoretic perspective. H. Yang and T. Yang (2011) propose a “meeting function” to investigate the equilibrium properties of taxi markets. Similarly, Zha, Yin, and H. Yang (2016) propose a “matching function” to model the e-hailing market. Alonso-Mora et al. (2017) propose an optimization model for matching riders and drivers in a ride-sourcing fleet. Guha, Demirezen, and Kumar (2018) model several aspects of the ride-sourcing platform including surge-pricing in a competitive setting as a differential game.

Several studies view surge pricing from a market equilibrium perspective. X. Wang et al. (2016) and He and Z. J. M. Shen (2015) propose an equilibrium model to examine the effects of several pricing strategies on e-hailing platforms. Similarly, Bimpikis, Candogan, and Saban (2019) propose an time-invariant equilibrium model of riders and drivers on a network, investigating the spatial effect of surge pricing. Zha, Yin, and Xu (2018) leverage a discrete-time geometric matching algorithm to examine the spatial effects of surge pricing under market equilibrium in time and propose a commission rate cap balancing platform revenue and consumer welfare. H. Sun, H. Wang, and Wan (2019) propose an econometric model of labor supply in ride-sourcing markets, using surge pricing as a natural experimental environment.

2.1.3 Congestion

Separately, there has been substantial research into the congestion effects of ride-sourcing fleets. The studies fall broadly into two categories: statistical analysis of ride-sourcing data and simulation-based studies. Among the reasons why ride-sourcing could affect congestion, Erhardt et al. (2019) identifies shared rides, integration with mass transit, and lower rates of car ownership as potentially beneficial and deadhead cruising, impeding traffic flow during pickups and dropoffs, and modal shift away from less congesting modes as potentially harmful. Several empirical studies, including Erhardt et al. (2019), have indicated that ride-sourcing increases congestion. Ward et al. (2019) and J. D. Hall, Palsson, and Price (2018) have further investigated effects of ride-sourcing fleets on public transit ridership, car ownership, and vehicle miles traveled. In contrast, simulation studies have offered ways in which ride-sourcing fleets could possibly be leveraged to reduce congestion. Fagnant and Kockelman (2014) simulate a fleet of shared autonomous vehicles (SAV) to conclude that such a system would require substantially fewer vehicles on the road.

2.1.4 Research gaps

This dissertation addresses several open research questions regarding TNCs. First, little is known about the relationship between how mobility services match supply and demand, and network utilization. In particular, surge pricing acts as a platform-generated signal to riders and drivers specifically designed to influence their decisions. Because the operational details of mobility services are proprietary this is difficult to measure directly. It is nevertheless crucial to understand how a signal within a virtual transportation network is likely to drive aggregate behavior on the physical transportation network. Surge pricing, if effectively predicted and disseminated to both drivers and riders, can be used to more efficiently allocate vehicles, save users money and time, and provide profitable insight to drivers, which ultimately helps the efficiency and reliability of transportation networks. It will also aid public agencies in assessing where and when traffic issues due to ride-sourcing fleets are likely to occur, allowing agencies to focus their efforts to determine what steps are needed to mitigate congestion effects. Our contribution in this area explores the spatio-temporal correlations between the urban environment, traffic flow characteristics, and surge multipliers. We propose a general framework for predicting the short-term evolution of surge multipliers in real-time using a log-linear model with L_1 regularization, coupled with pattern clustering.

Second, we investigate how the presence of a coordinating signal might affect network utilization, not over the next few hours, but over the next several years to decades. To do so we turn to network equilibrium theory, discussed further in Section 2.2. This is a challenging problem because it is not clear how planners should

expect mobility services to operate over the next decade or even how prevalent they will be. This is an important problem because mobility services are a likely fixture in transportation systems over any planning horizon; it is a risk for planners to *not* consider their presence in their plans. Our contribution in this area is to re-frame the problem: if planners wish to have a mobility service that improves network efficiency, what *would* it need to look like? How prevalent *would* it have to be? What risks might they pose to network efficiency? We examine how the route choices of coordinated fleets of vehicles, optimized to minimize the total fleet travel cost, impact the network. In particular, the extent to which fleet optimal behavior may also benefit other road users and network efficiency.

2.2 Transportation network equilibrium theory

2.2.1 Overview

The prevailing notion of network equilibrium was introduced by Wardrop (1952) who advocates for its use as a principled heuristic to estimate the impact of a proposed road network intervention on the future distribution of traffic flow. If travelers are “user-optimal” decision makers (that is they will choose their route in order to minimize their own travel time) we must contend with the fact that travel time both affects and is affected by traveler route choices. The game-theoretic framework proposed by Wardrop, and widely employed by transportation planners to this day (Boyles, Lownes, and Unnikrishnan 2021), acknowledges and reconciles this circular dependency by finding the traveler route choices and congestion levels that are mutually consistent: an equilibrium. This is preferable, Wardrop argues, to the “arbitrary assumptions” one would otherwise have to employ. Understanding network usage from a game-theoretic perspective, as emergent behavior of a population of travelers, offers a concise description of the *cause* of observed network states, and as a result, a principled way of forecasting network activity under hypothetical changes to the network. The ability to answer counterfactual questions of network utilization based on intuitive assumptions remains the main appeal of network equilibrium models. Accordingly, such approaches have sustained academic interest and practical application in transportation planning for the last seventy years.

Concretely, an assignment of travel demand onto paths in the network is an *equilibrium assignment* if and only if no individual traveler can reduce their travel cost by *unilaterally* switching routes. In other words, equilibrium is the state of traffic flow on the network for which all *used* paths between each origin destination pair have equal and minimal travel cost. An equilibrium of user-optimal travelers is referred to as a User Equilibrium (UE).

Mathematically, network equilibrium is most succinctly stated as a *Variational In-*

equality (Smith 1979). We are given a fully-connected directed graph $G = (V, E)$ representing the road network and a vector of travel demand on the network $q_{u,v} > 0$ for $(u, v) \in W \subseteq V \times V$, where W is the set of origin destination pairs. The set of all paths on G is denoted by P and the set of all paths from node u to node v by $P_{u,v}$. A feasible path flow is a non-negative vector $f \in \mathbb{R}^{|P|}$ such that $\sum_{p \in P_{u,v}} f_p = q_{u,v}$ for all $u, v \in W$; denote by F the set of feasible path flow vectors. A feasible path flow may be thought of as an assignment of the travel demand from node u to node v across the paths which connect node u to node v . We define the *flow performance function* as a non-negative vector-valued function $t : F \rightarrow \mathbb{R}^{|P|}$ which returns the travel cost on each path as a function of the flow on all paths. The system is at equilibrium if and only if there exists a feasible path flow $f^* \in F$ (the equilibrium path flow) such that;

$$t(f^*) \cdot (f - f^*) \geq 0 \quad \forall f \in F \quad (2.1)$$

where \cdot denotes the scalar product operator. Conditions for uniqueness and existence of f^* are given in (Smith 1979). The inequality is more intuitively understood as $t(f^*) \cdot f \geq t(f^*) \cdot f^*$: the “cost function”, $t(f^*)$ represents the travel cost of each path at the equilibrium flow, f^* , and $t(f^*) \cdot f^*$ is the total travel time on the network. What the variational inequality expresses then is that if $t^* = t(f^*)$ is known in advance to the travelers, then the best choice they have is to, in aggregate, distribute themselves according to f^* . In other words, f^* is the feasible path flow which uses only shortest paths according to $t(f^*)$.

There has been an active vein of research dedicated to the extension and further applications of the framework proposed by Wardrop (1952) (see Peeta and Ziliaskopoulos 2001; Sheffi 1985; Boyles, Lownes, and Unnikrishnan 2021; Friesz 1985; Boyce 2007) and it remains an unrivaled tool in the transportation planner’s toolkit, even as data-driven computational methods (e.g. machine learning, artificial intelligence) have gained prominence elsewhere. Network equilibrium theory remains relevant precisely because it is able to assess the validity of counterfactual statements about the road network over multi-year (or even multi-decade) planning horizons based on a handful of intuitive assumptions on human behavior. These assumptions are not without flaws, but to echo Wardrop, they seem preferable to the assumptions you would need to make to justify the use of a machine learning model to predict the state of the network several years from now. The assumptions and their flaws will be examined in more detail in Section 2.2.2.

For example, a transportation planner may want a principled way to determine which of two proposed road infrastructure projects offers a better ratio of congestion reduction to cost. This is a counterfactual question at heart: how *would* traffic congestion change over the next 10 years *if* the proposed infrastructure project were to be completed? This is not a question amenable to data-driven methods primarily because a relevant data set would be extremely costly to collect (it would need to span at least 10 years) but also because it’s difficult to assess how well a data-driven

methodology would be expected to generalize to this particular road network at this particular point in time. In contrast, network equilibrium methods assign a notion of travel cost to travelers on the network and let aggregate network behavior emerge from a repeated game in which travelers seek to minimize their travel cost.

2.2.2 Assumptions and implications

There are three key assumptions that justify the application of transportation network equilibrium to transportation planning problems:

1. All travelers share the same notion of travel cost,
2. All travelers have access to perfect information, and
3. All travelers seek to minimize their own travel cost.

Conditions for existence and uniqueness of the solution further require that the travel cost is continuous and monotonic increasing with respect to the flow of travelers on the network.

The assumption that human travelers are perfectly rational with respect to a single notion of cost is certainly unrealistic. The assumption of access to perfect information is simply invalid as such complete network state information (under any useful definition of the term) does not even exist. The critical question is to determine the extent to which these simplifying assumptions result in a useful model of transportation networks.

Several studies have called into question the aforementioned assumptions as well as the conventional assumption that the flow performance function measures travel time. Zhu and Levinson (2015) and Jan, A. J. Horowitz, and Peng (2000) examined GPS traces of individual drivers and found that drivers minimize neither travel time nor distance in a consistent way. Levinson (2003) and Ben-Elia et al. (2013) examine the impact of “advanced traveler information systems” (e.g. Google Maps) on route choice and found that the quality of information has a significant impact on traveler behavior; a result that would be surprising under the assumption that travelers already have all the information they need to make route choices.

Although it is useful to understand how traveler behavior deviates from the model assumptions, it is critical to understand the impact of this deviation on the implications of the theory on *aggregate* network usage over time. To this end, Ma and S. Qian (2019) sought to measure the extent to which the network *currently* deviates from equilibrium and infer this quantity from GPS trajectory data. However, the utility of equilibrium theory lies in its claims on future, hypothetical network usage.

Implicitly, the theory assumes an additional criterion: that the network state described by equilibrium is (approximately) realized by some *adjustment or learning process* of the traveler’s route choices in response to the network change. This follows directly from the interpretation of equilibrium as best estimate of the redistribution

of traffic on a network: if we suppose traffic will eventually resemble this distribution there must exist a process governing the tendency toward equilibrium (Watling and Hazelton 2003). Models that explicitly represent these processes are known as *day-to-day traffic assignment models* (Friesz et al. 1994). Of course, not all adjustment processes lead to equilibrium, or even converge in any meaningful way, but some do (Hazelton 2022). However, if the existence of an adjustment process is required by equilibrium theory, then the fact that a network is not currently in an equilibrium state should not necessarily be taken as evidence against equilibrium theory.

2.2.3 Stochastic user equilibrium

A particularly relevant extension of static network equilibrium to this thesis is known as stochastic user equilibrium (SUE). Stochastic user equilibrium was first introduced by Daganzo and Sheffi (1977) (see also Boyles, Lowmes, and Unnikrishnan 2021, Section 8.3) to relax the assumption that individuals are able to perfectly measure their travel cost. Instead, each infinitesimal traveler chooses a route according to the probability that it is the least cost path available to them. Because each traveler represents an infinitesimal unit of flow, the aggregate network flow reported by SUE is deterministic: the fraction of demand taking a path is equal to the probability that any one traveler chooses the path.

In its original and still most common setting, SUE uses a logit model to represent the route choice probabilities. The logit model and choice modeling more generally will be examined in more detail in Section 2.4. Explicitly, the probability that a traveler chooses path k given a feasible path flow f is proportional to the exponential of the negative of the travel cost of the path, as written below.

$$\Pr[k \mid f] \propto \exp(-t_k(f)) \quad (2.2)$$

2.2.4 Equilibrium of mixed traffic

Another relevant extension of static network equilibrium is the equilibrium of traffic flow comprised of two or more “classes” of flow. First introduced by Dafermos (1972), the equilibrium of mixed traffic seeks to relax the assumption that all travelers share the same notion of cost. As Dafermos (1972) observes,

These [equilibrium] models are based on the common assumption that all travelers on the same link of the network are identical insofar as cost is concerned. However, common experience with highway networks indicates that the above assumption constitutes an over-simplification. Indeed, it is clear that the deviations of individual traveling units from the ‘average’ one are so important that the realism of the existing models for the traffic assignment problem may be questionable.

We may understand travelers desire to avoid congestion, but we may also understand that different travelers experience different congestion to different extents. As a solution, Dafermos offers the concept of a multi-class network equilibrium whereby the flow on each route is decomposed into the flow from each class where each class is permitted to define their own notion of cost as a function of the class flows. Dafermos imagined that “typical” classes might be trucks and passenger cars, or speeders and slow drivers. These classes perceive cost in very different ways. Trucks for example avoid narrow local roads in favor of highways, and thus are less sensitive to congestion. Passenger vehicles, on the other hand, may opt to travel on an alternate route with fewer trucks even though it may take longer.

Harker (1988) arrives at multi-class equilibrium from a different perspective: to model the interaction of coordinated fleets with individual users on the network. In Harker’s setting, all vehicles share the same notion of cost but each fleet coordinates their route choice to minimize their own total fleet cost. Individual travelers, on the other hand, seek to minimize only their own travel time. Harker refers to fleets as “Cournot-Nash players”, as in the firms in a Cournot oligopoly model (Leonard 1994; Moulin 1986) who non-cooperatively determine the quantity of goods produced. The fleets, like the firms in the Cournot oligopoly model, determine the quantity of flow to send over each path in bulk and can therefore coordinate the flow to achieve a lower total cost. In contrast, individual users each control an infinitesimal volume of flow and so their travel costs are completely determined by the aggregate flow. Harker (1988) demonstrates that this form of mixed equilibrium may be expressed as a variational inequality which can be solved via existing methods.

H. Yang, X. Zhang, and Meng (2007) extend Harker’s analysis to examine a network with individuals, fleets, and “system optimal” (SO) users who aim to minimize the total system cost. Although H. Yang, X. Zhang, and Meng (2007) primarily serves to introduce the formulation and solution algorithm for this particular kind of mixed equilibrium, there is an important empirical observation in H. Yang, X. Zhang, and Meng (2007) that sparked a continuing line of inquiry: when individuals control enough of the demand, neither SO nor fleet users may change the total system cost, conversely, if there are few enough individuals on the network a combination of fleets and SO users can achieve SO flow.

There has recently been renewed interest in the line of literature established by Harker (1988) and H. Yang, X. Zhang, and Meng (2007) due largely to the advent of *ad hoc* fleets: a set of vehicles whose behavior is coordinated via a virtual transportation infrastructure. Indeed, to realize the Cournot-Nash player of Harker (1988) in practice requires that the fleet is able to coordinate itself in order to minimize collective cost. By comparison, there is extensive literature on ways to induce SO behavior by manipulating the individual notion of cost via infrastructure (e.g. ramp metering (Sheffi 1985)). In our setting, coordination comes not from the physical infrastruc-

ture but from the digital infrastructure employed by the service provider. In parallel work, Sharon et al. (2018) and Z. Chen et al. (2020) identify a “minimum-control ratio” (MCR) which they define as the smallest volume of SO users capable of inducing SO flow on a network shared with UE users. In both cases, the fleet is centrally routed to minimize total system cost, corresponding to the second half of H. Yang, X. Zhang, and Meng (2007)’s observation: that SO flow can be achieved even if not all users are SO users. Although the system as a whole is better off, the SO users themselves are relatively, if not absolutely, worse off precisely because they prioritized the travel cost of others over that of their own. We examine this idea more precisely in Section 4.3.3. This may not be realistic for private service providers who have no inherent altruistic motivation.

A second line of inquiry addresses the regulations or incentives required to realize the benefits of a centrally routed fleet. K. Zhang and Y. M. Nie (2018) view the fleet as a direct government intervention and balance the benefits of introducing a fleet of SO users against the cost of their deployment. In this view, the fleet is a kind of “mobile actuator” (J. Wang et al. 2020) and becomes another tool in the toolkit of a traffic manager to advance municipal mobility goals directly. Another perspective is offered by Mansourianfar et al. (2021) and Delle Site (2021) who view the fleet as a third party which must be compensated or tolled in order to align their interests with that of the system’s. In contrast, the fleets we examine in this work are assumed to be guided by a self-interested operator who arranges its own fleet in order to optimize the operator’s own objective. Naturally, the operator needs no incentive to optimize their fleet. Lastly, Mehr and R. Horowitz (2019) investigate the impact of platooning autonomous vehicles (AVs) on network equilibrium. When several autonomous vehicles drive in close succession with their driving maneuvers tightly coupled, the resulting AV platoon effectively increases the capacity of the road segment which in turn modifies the link performance function, thus affecting network equilibrium. Similar to our work, they find that a mixed equilibrium of platooning AVs and individual users may increase total delay in the network.

Perhaps most closely related to our work, Cominetti, Correa, and Stier-Moses (2009) investigate the impact of a set of Harker (1988)’s Cournot-Nash players on the network, but do not consider a mix of fleet and individuals.

2.2.5 Research gaps

There are currently several headwinds facing network equilibrium theory which may be grouped into three main challenges, each also offering tremendous research opportunities. This dissertation addresses each of them.

First, virtual transportation infrastructure has dramatically altered how travelers utilize the network, yet much of network equilibrium theory in practice however, remains focused on personal vehicular flow, in particular during the morning

and evening commutes. Extending the theory to account for heterogeneous and/or multi-modal road users remains an open and active area of research. To this end, we examine how a mixed equilibrium approach can be leveraged in the planning process to examine the circumstances under which the presence of a particular heterogeneity will affect traffic patterns and, in particular, when they will improve them. This side-steps uncertainty about which kinds of heterogeneity should be considered by instead focusing on the conditions under which a given heterogeneity will affect the network.

We introduce two novel and particularly important notions of fleet-optimal mixed equilibria: critical fleet size for user equilibrium (CFS-UE) and critical fleet size for system optimum (CFS-SO), capturing the largest fleet for which the resultant mixed equilibrium is user equilibrium, and the smallest fleet for which mixed equilibrium is system optimal. We develop novel mathematical programs and apply them to find critical fleet size on the Sioux Falls and Pittsburgh networks where 33% and 83% of vehicles respectively, must participate in the fleet to achieve system optimum. In Pittsburgh, we find that although fleets permeate all of the links on the network, they accumulate on highways and major arterials. We find the majority of origin-destination pairs to be either exclusively user or exclusively fleet occupied. Critical fleet size offers regulators greater insight into where fleet and system interests align, transportation planners a novel metric to evaluate the benefit of a road improvement, and fleet coordinators a better understanding of their efforts to optimize their fleet.

Second is that the theory has no inherent way to leverage the increasing availability of mobility data. Network equilibrium modeling requires us to specify the rules governing traveler decision-making and in return provides us with a forecast of what network state should emerge from aggregate traveler behavior in the long-run. Mobility data gives us increasingly accurate direct measurements of network state, but as of yet, has little to tell us of the rules governing human travel decision-making. Inverting network equilibrium would then naturally allow insight into travel behavior from network state measurements. Incorporating data from transportation networks in this fashion is currently an open and active area of research. To this end, we introduce a novel statistical procedure to detect when a set of agents appear to be acting in order to learn over time. This methodology has clear applications within choice modeling, but it may also be applied to network equilibrium theory where it could be used to demonstrate that individuals appear to be minimizing a given cost function over time as a way to empirically justify that cost function.

Third is that applications of network equilibrium research remain confined and as a result the scope of the theory appears more limited than we believe it to be. The software companies that provide simulation tools to transportation planners remain the main consumers of research in this niche. Though this application is of paramount importance, it is not the only application. We contend that as both the public and private sectors actively re-imagine the public right-of-way, network equilibrium theory

should be a central component to their planning processes. Our notions of critical fleet size are as important to transportation planners in the public sector as they are to the fleet coordinators in the private sector, giving them insight into how their fleets should be expected to behave in mixed traffic and illuminating ways in which their deployment would be *helpful* to public agencies as well as where their deployment would be unlikely to change network utilization at all.

2.3 Predictive modeling in transportation systems

Predictive models of transportation driven by existing sources of urban social-physical data are well-studied in previous literature, but not with a focus on surge pricing. Several researchers have built predictive models of travel characteristics other than surge multipliers. Each model uses data from a single system to predict its target. Z. Zheng et al. (2016) develop a real-time predictive model of traffic flow from the CO2 sensors in a large office building in Hong Kong. P. Zhang, Zhen, and Qian (2017) use household energy consumption to predict the start and end of morning traffic congestion in Austin, Texas. In contrast, J. Liu et al. (2019) synthesize taxi GPS data and license plate recognition data on the road network to infer vehicle volume and fleet composition at high geospatial resolution. The synthesis of multiple data sources is conceptually aligned with our approach in this work.

In regard to ride-sourcing systems, Ke, H. Zheng, et al. (2017) develop a deep learning based approach to short-term demand forecasting in ride-sourcing systems. Similar to our work, they use both historical measurements and exogenous features such as time-of-day, day-of-week, and weather. They also use a spatially aggregated random forest for feature selection. Similarly, Wei and M. C. Chen (2012) and Y. Li et al. (2017) develop neural networks to predict short-term demand in metro systems based on historical data and temporal factors. Noursalehi, Koutsopoulos, and J. Zhao (2018) predict passenger arrivals at metro stations in real time using dynamic factor models.

There are generally two types of real-time prediction. The first aims to predict only the next time step and updates itself in real-time, often in the Bayesian sense. Such models are usually referred to as short-term forecasting. Fei, Lu, and K. Liu (2011) build a real-time short term travel time model capable of forecasting 5 minutes ahead of time. Mohammad M. Hamed, Al-Masaeid, and Said (1995) predict the short-term evolution of traffic volume over the next minute. W. Zheng et al. (2006) predict the short term evolution of traffic flow, again over the next few minutes. The second category offers predictions over longer time horizons. Relatively few studies fall into this category. Min and Wynter (2011) build a real-time predictive model of traffic up to 1 hour in advance. Vanajakshi and Rilett (2004) compare the performance of neural networks and support vector machines in predicting traffic speed up to 1 hour

in advance.

More recently, a number of studies have applied machine learning methods to predict demand for ride-sourcing services, most with the specific aim of improving ride-sourcing operations. Ke, Feng, et al. (2021) introduce a multi-task machine learning model to predict ride-sourcing demand over multiple ride-sourcing services. Kontou, Garikapati, and Hou (2020) demonstrate that machine-learning-based real-time demand prediction can be applied to reduce dead-heading mileage in ride-sourcing systems. A. Li and Axhausen (2020) construct a graph convolutional neural network to perform short-term demand prediction with applications to ride-sourcing networks. Ke, Qin, et al. (2021) predict origin-destination ride-sourcing demand using a graph convolutional neural network. Yan, X. Liu, and X. Zhao (2020) fit a random forest model to predict zone-to-zone demand for ride-hailing services.

In contrast, relatively few studies have focused on the prediction of ride-hailing supply. D. Wang et al. (2017) and Ke, H. Yang, et al. (2019) apply deep learning methods to forecast the gap between ride-hailing supply and demand in real time. Compared to demand prediction, supply prediction may be more challenging because the data is more difficult to obtain. While demand for transportation can be measured in a variety of ways, ride-sourcing supply is difficult to measure without access to proprietary data, which TNCs are generally reluctant to share. Moreover, supply, unlike demand, was historically not something that needed to be predicted. As a result, “supply forecasting” does not enjoy the same depth of academic community as demand forecasting does. Travel demand has always been the result of the aggregate choices of individuals and hence readily modeled as a random process. In contrast, supply was historically centrally controlled (as in public road networks) or heavily regulated (as in taxi or public transit services) and as a result, ill-suited for forecasting models. As disaggregate supply continues to pervade transportation systems we expect supply forecasting models to attract increasing attention.

Our contribution in this area is a predictive model of surge pricing capable of forecasts up to two hours in the future based on the last hour of observed surge multipliers as well as the last hour of observed urban social-physical features. To our knowledge, real-time predictive models of surge pricing have not been explored in transportation literature. Our model uses a broadly collected, multi-source data set, combining data from multiple disparate systems across the public and private sectors, such as traffic speed, events, road closures, and weather conditions. We employ L_1 regularization to utilize the entire feature set without over-fitting. Most existing work uses data from a single source.

2.4 Choice modeling in transportation systems

2.4.1 Overview

Inferring intent from behavior is an important theme in transportation research. Applications of discrete choice modeling to transportation systems first appeared in the 1960s and were aimed primarily at questions of mode choice, expanding over the course of the 1970s to other dimensions of travel choice, including car ownership, travel demand, and housing (Ben-Akiva 1985; McFadden 1974). The primary methodology in this line of inquiry is the multinomial logit model (McFadden et al. 1973; Bhat 1997) which assigns to each choice in an exhaustive set of mutually exclusive choices a probability:

$$\Pr[Y = j] \propto \exp(\langle X_j, \beta \rangle) \quad (2.3)$$

where Y is a random variable representing the decision reached, X_j represents a vector of attributes of the decision j , and β a vector of *preference parameters*. The dot product between the preference parameters β and the characteristics of the decision X_j is known as the *utility* or *reward* of the choice j . It is assumed that agents wish to take the decision which maximizes their utility. The particular probabilistic model defined by Equation (2.3) arises from an assumption that the observed utility, $U_j := \langle X_j, \beta \rangle$, differs from the true value of utility, V_j , (used by the agent to make the decision) by a Gumbel-distributed random quantity, assumed to be i.i.d. over both the choice set and decisions.

Typically, random utility theory is motivated from the perspective of the statistician: it is assumed the agent has access to the full set of decision features and is able to correctly select the highest utility choice. The randomness that we, the statisticians, observe is due entirely to the fact that we cannot observe the full set of criteria the agents are using. However, the randomness can also be thought of as arising from the environment which motivates random utility theory from the agent’s perspective. In this view, the true utility, V_j , is revealed only after the agent has made a choice, and the observed utility, U_j , represents the agent’s *anticipated* utility from taking decision j . In order to arrive at Equation (2.3), the difference between the reward the agent anticipated and what they actually received is Gumbel distributed, assumed to be i.i.d. over both the choice set and decisions.

When viewed from the agent’s perspective, it is clear that the agent is assumed to know quite a lot about how the environment will respond to their action. In fact, the agent has a “correct” model of the environment in that their uncertainty about the reward is entirely *aleatory*, stemming from inherent randomness rather than from misspecification. When the agent is interacting with a complex environment with a large action space, it may not be reasonable to assume that the agent has no epistemic

uncertainty about their environment. In such cases, the agent takes actions based on their *beliefs* about the reward which are either substantiated or refuted by the actual reward they receive and may lead the agent to adjust their beliefs for subsequent decisions. As a result there is a temporal dependence between actions which standard choice modeling cannot account for.

In typical transportation applications of choice modeling, the dataset of decisions is aggregated across individuals at a point in time; any temporal dependence that may exist within an individual’s policy over time is not detrimental to the application. However, as disaggregate data, for example individual GPS trajectories, becomes both more prevalent and more detailed, the application of choice modeling to datasets of individual’s behavior over time becomes more enticing. In perhaps its most general form, this approach is termed *inverse reinforcement learning* (IRL).

2.4.2 Inverse reinforcement learning

Inverse reinforcement learning observes a *trajectory* of behavior and produces a reward function which best rationalizes the observed behavior. The dataset of behavior in IRL is a temporal sequence of state, action pairs for each observed individual. Maximum Entropy IRL (MaxEnt IRL) (Ziebart et al. 2008) effectively extends the Equation (2.3) to the Markov decision process (MDP) on which reinforcement learning is based:

$$\Pr[\zeta = (s_1, a_1, \dots, s_T, a_T)] \propto \exp \left(\sum_{t=1}^T r_{\theta}(s_t, a_t) \right) \quad (2.4)$$

where ζ is a random variable representing a trajectory through the MDP, r_{θ} is the reward (parameterized by θ) as a function of the state at time t and the action taken at time t .

In all cases the reward function (and by extension, the policy) is explicitly assumed to be time independent, as the parameters of the reward function are fixed over time. In other words, the agent’s beliefs about the relationship between state, action, and reward are unchanged by their previous observations of rewards. The assumption of a time-independent policy in IRL is typically motivated from the perspective of *behavioral cloning* (Sutton, Barto, et al. 1998) wherein the behavior is assumed to be an “expert demonstration” of the task. The experts, it is assumed, have already learned everything they need to know about the task. In this sense, inverse reinforcement learning as a term used to describe the process of inferring a reward function from demonstrations of a fixed expert policy can be considered somewhat of a misnomer: if (forward) reinforcement learning is the process of efficiently balancing exploration and exploitation to learn a policy that maximizes a known reward function, then inverse reinforcement learning should refer to the process of inferring the reward func-

tion from the behavior of agents who are in the process of efficiently learning policies to maximize it.

2.4.3 Choice modeling and transportation network companies

Choice models have an important role to play in understanding the impact of TNCs on transportation systems. On the demand side, where much of the research has been performed, choice modeling helps transportation planners understand what motivates travelers to use ride-sourcing over other modes of transportation. One of the first demand-side choice models specifically aimed at ride-sourcing systems demand was given by Dias et al. (2017) who estimated a probit choice model to explain the influence of socio-economic and demographic factors on the frequency of ride-sourcing use. Alemi et al. (2018) fit a binary logit model to estimate the extent to which socio-demographic, attitudinal, and environmental factors influence traveler adoption of ride-hailing services. Similarly, Lavieri and Bhat (2019) construct a latent-variable model relating socio-demographic and revealed transportation preferences to the adoption and utilization of ride-hailing services.

Choice models have also been applied to study ride-sourcing drivers. J. V. Hall and Krueger (2018) provide the first comprehensive study of the labor market for Uber drivers. Though their findings offer a descriptive account of Uber drivers, J. V. Hall and Krueger (2018) identify important features that distinguish Uber drivers from the general population. Their findings lay the foundation for subsequent research aimed at explaining driver behavior. Berliner and Tal (2018) estimate an ordinal logit model to understand how socio-demographic and attitudinal factors influence willingness to drive for a TNC. Ruijter et al. (2022) model the emergent behavior of decentralized driver decision making on the evolution of supply in ride-hailing networks. Bansal et al. (2020) estimate preferences of both TNC riders and drivers to understand both TNC usage and vehicle purchasing behavior.

There is also a small but growing body of research in applying choice models to understand TNCs in the context of urban delivery. Miller, Y. Nie, and Stathopoulos (2017) model traveler willingness to work as a delivery driver on a TNC platform. Punel and Stathopoulos (2017) model the preferences of package senders engaged with TNC delivery platforms.

2.4.4 Research gaps

Our contribution in this area is to extend choice modeling to apply in circumstances where agents engaged in sequential decision making are likely changing their decision-making criteria over the course of the observation period in pursuit of maximizing some reward or utility function. We introduce a novel statistical framework to detect

whether (and when) learning is occurring and, if it is, identify regions of the decision-criteria space that the agents are drawing towards over time.

2.5 Sequential hypothesis testing with non-negative super-martingales

2.5.1 Overview

Sequential hypothesis testing extends the notion of classical hypothesis testing to the setting in which data may be observed, and the hypothesis test may be performed, in sequence. The value of sequential hypothesis testing methods is that, unlike conventional hypothesis testing methods, they are valid under arbitrary data-dependent stopping times (Ramdas, Ruf, Larsson, and W. Koolen 2020). That is, the researcher can inspect the data and results of the hypothesis test over time and decide when to stop collecting data based on any function of the data collected so far, including the value of the test statistic itself. When using conventional hypothesis testing, using the data to determine when to stop the test and draw conclusions invalidates the results and is known as “p-hacking” (Head et al. 2015). The pursuit of sequential hypothesis testing and analogous confidence *sequences* is the pursuit of “anytime-valid” tests: tests that retain their validity in the face of arbitrary data-dependent stopping times.

An intuition for sequential hypothesis testing can be developed from gambling. In fact some of the earliest academic work in probability theory developed in order to define what it meant for a gamble to be “fair” (Shafer and Vovk 2019, Preface). This legacy of gambling remains when we use a “fair coin” or “fair dice” to exhibit some aspect of probability theory, but as Shafer and Vovk (2019) expound fairness is not an axiom, but rather a result of the game being played. Fairness in gambling is fundamentally a question of pricing. If a gambler is being offered a ticket that may be redeemed for \$1 if the outcome of the next coin toss is heads, what is the most the gambler should be willing to pay for it? The “fair” price in this instance is exactly the probability that the coin toss results in heads. If the price of a ticket is more than the “fair” amount, we would think the gambler is being cheated; if the price is less than the “fair” amount, we would think the house is being cheated.

Shafer and Vovk (2019) refine this idea and give it a rigorous mathematical foundation by developing a game-theoretic theory of probability. However, what Shafer and Vovk develop is also true in the more conventional measure-theoretic probability. We suppose that a gambler will play this coin toss game repeatedly, beginning with \$1 and without the ability to borrow money or go into debt at any point. We can represent the gambler’s wealth over the rounds of the game (time) as a sequence of random variables (a stochastic process). If the gambler is allowed to purchase or short-sell the

ticket for heads in any quantity then Shafer and Vovk (2019, Proposition 1.2) prove that the gambler has a strategy over repeated plays of this game which guarantees that either the price of the game is fair (i.e. the price matches the probability) or the gambler’s wealth will grow infinitely.

When the coin toss game is fair, the gambler’s wealth over time is a non-negative (super-)martingale (NSM). A martingale is a sequence of random variables whose value at any time is the conditional expectation of its value at any future time (Ramdas, Ruf, Larsson, and W. Koolen 2020). A non-negative martingale is a martingale for which every random variable in the sequence is non-negative with probability 1. For example, when the coin toss game is fair, the expected value of the gambler’s wealth after the next coin toss given their wealth now is simply the gambler’s wealth now. In other words, the expected value of the gambler’s winnings on each round is equal to the price they paid for the ticket: it is a fair game. In the language of mathematical probability, let \mathcal{F}_t be a filtration and P be a measure on \mathcal{F}_∞ . Following Ramdas, Ruf, Larsson, and W. Koolen (2020), a stochastic process $(M_t)_{t=0\dots}$ is a **martingale** if M_t is \mathcal{F}_t -measurable, P -integrable, and

$$\mathbb{E}_P[M_t \mid \mathcal{F}_s] = M_s \text{ for any } t \text{ and } s \leq t. \quad (2.5)$$

$(M_t)_{t=0,\dots}$ is a **super-martingale** when instead,

$$\mathbb{E}_P[M_t \mid \mathcal{F}_s] \leq M_s \text{ for any } t \text{ and } s \leq t. \quad (2.6)$$

Finally, $(M_t)_{t=0,\dots}$ is a **sub-martingale** when instead,

$$\mathbb{E}_P[M_t \mid \mathcal{F}_s] \geq M_s \text{ for any } t \text{ and } s \leq t. \quad (2.7)$$

The filtration \mathcal{F}_t can be described in conceptual terms as the set of all information available at time t . The *canonical filtration* for a stochastic process M_t is usually taken to be the σ -algebra generated by the sequence $(M_i)_{i=1}^t$, which is to say the set of all possible realizations of the sequence $(M_i)_{i=1}^t$. Although the filtration is an important concept to the theory of stochastic processes, in practice we typically do not have to think much about it so long as we are careful that the information available at time t does not include information from future times. In this thesis we will abbreviate the conditional expectation $\mathbb{E}_P[M_t \mid \mathcal{F}_s]$ as $\mathbb{E}_s[M_t]$ when both the probability measure P and the filtration \mathcal{F}_s are clear from context.

Non-negative super-martingales are at the heart of several converging lines of statistical research in sequential hypothesis testing and sequential estimation (i.e. confidence sequences) (Howard and Ramdas 2019). Shafer, A. Shen, et al. (2011) refer to “test martingales”, Shafer and Vovk (2019) and Shafer (2019) to “betting scores”, and Vovk and R. Wang (2021) and Grünwald, Heide, and W. M. Koolen (2020) to “safe e-values”. Underlying all of this work are non-negative super-martingales. In fact,

Ramdas, Ruf, Larsson, and W. Koolen (2020) show that *every* sequential hypothesis test or estimation method must either itself rely on non-negative super-martingales or be strictly out-performed by one.

2.5.2 Ville's inequality

The workhorse of hypothesis testing with non-negative super-martingales is *Ville's Inequality* (Ville 1939) (see also Howard, Ramdas, et al. 2020, Section 2.3) which extends Markov's inequality to stochastic processes, namely non-negative super-martingales. Let Y_t be a non-negative super-martingale adapted to the filtration \mathcal{F}_t ; Ville's inequality allows us to conclude that

$$\Pr[\exists t : Y_t \geq \alpha \mathbb{E}[Y_0]] \leq \frac{1}{\alpha} \quad (2.8)$$

where Y_0 is the initial value of the NSM. What Ville's inequality states is that the event that a NSM exceeds its initial (expected) value by a factor of α *at any point in time* occurs with probability at most $1/\alpha$. The core idea of sequential testing with NSMs is to construct a test statistic, itself a stochastic process, that is (upper bounded by) an NSM with initial (expected) value less than or equal to 1 under the null hypothesis. Ville's inequality then allows us to reject the null at confidence level $1 - \alpha$ if we observe the sequence *at any point in time* to exceed $1/\alpha$.

2.5.3 E-processes

An e-process is a generalization of non-negative super-martingales similarly suitable for hypothesis testing. Concretely, for a set of distributions \mathcal{P} , a \mathcal{P} -safe e-process is a non-negative stochastic process E_t with the following property;

$$\sup_{P \in \mathcal{P}} \sup_{\tau} \mathbb{E}_P[E_{\tau}] \leq 1 \quad (2.9)$$

where the inner supremum is taken over all possible stopping times τ . A stopping time is any time that can be selected using *only* information available up to that point in time, that is, the decision to stop at time t cannot depend on data collected after time t . This property implies, via Markov's inequality, the stopping time defined as the earliest time t such that $E_t \geq 1/\alpha$ results in a level α sequential test against the null \mathcal{P} (Ramdas, Ruf, Larsson, and W. M. Koolen 2022). The above property guarantees that stopping an e-process at any valid stopping time results in a valid sequential test statistic.

E-processes are related to the more recognizable p-value via simple inversion. For any \mathcal{P} -safe e-process E_t , the sequence $p_t := \inf_{s \leq t} 1/E_s$ is an any-time valid \mathcal{P} -safe p-process (Grünwald, Heide, and W. M. Koolen 2020). As its name suggests, a \mathcal{P} -safe p-process produces p-values against the null \mathcal{P} at every stopping time.

Intuitively, e-processes can be thought of as a generalization of the likelihood ratio to handle a composite null and composite alternative. For a single null P and single alternative Q , Wald (1945) develops the sequential likelihood ratio test:

$$L_t := \frac{\ell_Q(X_1, \dots, X_t)}{\ell_P(X_1, \dots, X_t)} \quad (2.10)$$

The ratio measures the weight of evidence that the data was generated by Q rather than P . When the true data-generating distribution is P , the expectation of the P likelihood function evaluated at the data will be at least as great as the expectation of the Q likelihood function. In other words, $\mathbb{E}_P[L_t] \leq 1$. Here, L_t is also a NSM under P .

In general, e-processes are intimately related to test NSMs. When \mathcal{P} is a singleton every NSM is also an e-process. For general sets, \mathcal{P} , that is when the null is composite, there may not exist a single process that is an NSM for every $P \in \mathcal{P}$. However, as Ramdas, Ruf, Larsson, and W. Koolen (2020) show, every \mathcal{P} -safe e-process, E_t , is dominated by a \mathcal{P} -safe e-process of the form

$$E_t \leq E_t^* := \inf_{P \in \mathcal{P}} M_t^P \quad (2.11)$$

where M_t^P is a non-negative martingale on P with initial expectation (at most) 1. Neither E_t nor E_t^* is itself a non-negative (super-)martingale but it is upper bounded by one so Ville's inequality is still applicable.

Ramdas, Ruf, Larsson, and W. M. Koolen (2022) demonstrate one particularly useful and simple construction of e-processes for a composite null and a composite non-parametric alternative which we leverage in this work. Inspired by Wasserman, Ramdas, and Balakrishnan (2020), they define the following \mathcal{P} -safe e-process;

$$\inf_{P \in \mathcal{P}} \frac{\prod_{s=1}^t g_s(X_s)}{\ell_P(X_1, \dots, X_t)} = \frac{\prod_{s=1}^t g_s(X_s)}{\sup_{P \in \mathcal{P}} \ell_P(X_1, \dots, X_t)} \quad (2.12)$$

where g_s is any “non-anticipating” probability mass (or density) function. The function g_s is “non-anticipating” if it depends only on the first $s - 1$ data points. As a result, g_s can be constructed as data is being collected, learning from past data in any way it sees fit. This construction opens the door for the direct application of machine learning methods in sequential testing. In the language of Shafer and Vovk (2019), g_s is a “bet” against the null which, we as the statistician, are free to generate. Of course, the choice of g_s implicitly defines an alternative and therefore the power of the test. We can use this to our advantage to tune the sensitivity of the test toward preferred regions of the alternative space.

3 **What drives a ride-sourcing fleet?: real-time prediction of supply-demand imbalances in ride-sourcing networks**

3.1 Introduction

Ride-sourcing companies like Uber and Lyft, also known as Transportation Network Companies (TNCs) as a more general term, have become a common presence in cities worldwide. Despite their growing popularity, much of their operational metrics remain opaque. In particular, little is known about *surge pricing*: the dynamic pricing mechanism employed to ensure sufficient vehicle supply during periods of high rider demand. Though it is unclear how surge pricing is updated at each time point, forecasting the surge prices in the next few minutes to hours would appear to benefit all parties. If effectively predicted and disseminated to both drivers and riders, future surge prices can be used to 1) help transportation managers understand the evolution of ride-sourcing supply and demand during demand surges, 2) help ride-sourcing companies navigate the changing transportation landscape to improve real-time operations, 3) provide profitable insights to drivers, and 4) save riders' money and time.

Surge-pricing is a signal that exists only within the virtual transportation system which serves to coordinate ride-sourcing supply in response to an imbalance of supply and demand. However, it affects and is affected by the physical system of transportation infrastructure and the social systems generating demand for and supply of ride-sourcing services. As a signal within the virtual transportation system, surge-pricing mediates the relationship between supply and demand as well as the relationship between supply and the physical infrastructure. Learning the relationship between urban social-physical systems and surge-pricing is an important step in understanding how this particular virtual signal is steering ride-sourcing supply on the network. The associations learned by the model are helpful not only to anticipate when and where

excess demand might occur, but also inform policy makers how best to target TNC regulations and roadway infrastructure improvements.

The existence of surge price multiplier implies a spatio-temporal imbalance of supply and demand in a ride-sourcing system. With this in mind, surge pricing can be viewed as a market correction due to information not already priced into the cost of a ride (Gurley 2014). In other words, it is a measure of unanticipated demand for ride-sourcing services and, in certain circumstances, a proxy for unanticipated travel demand more broadly. Predicting surge multipliers is therefore a way of characterizing the nature of unanticipated demand within a ride sourcing system. Under the assumption of perfect information, surge pricing should be random and brief. If it can be predicted then there is room for greater efficiency in the coordinating signals used to allocate vehicle supply. Our work does not intend to reverse engineer the underlying mechanism of surge pricing offered by ride-sourcing companies. Rather, we use machine learning methods to predict surge prices minutes or hours ahead of time in selected locations in a city, regardless of surge pricing strategies. The prediction is performed with a data-driven approach. In fact, the surge prices become a holistic measure of the relation between vehicle supply and rider demand, both temporally and spatially. As much as riders and drivers would like to forecast surge prices for trip or business planning, it would also benefit ride-sourcing companies to proactively manage fleet operations with knowledge of the distribution of vehicles and riders in advance.

Specifically, a real-time predictive model of surge multipliers over time and locations is of interest to several parties in the ride-sourcing ecosystem.

- **Drivers** may use the model to increase their earnings by acting on better forecasts of when and where surge pricing is likely to occur.
- **Riders** may use the model to plan their departures to avoid surge multipliers. In particular, riders could know whether the surge multiplier is soon likely to increase or how long they would likely need to wait for the surge multiplier to decrease.
- **Ride-sourcing companies or TNCs** may use the model to more intelligently dispatch their fleet or offer additional incentives to drivers to better balance supply and demand.
- **Third parties**, such as Gridwise Inc., who sell ride-sourcing intelligence to drivers, may use this model to improve their products. Public agencies may use this model to assess gaps in public transit operations or to implement targeted TNC policy.

A reliable, transparent model enables drivers to make choices with confidence. Drivers do not always trust market information provided by Uber since Uber's strategies may not be in alignment with individual drivers (Guda and Subramanian 2017;

Rosenblat and Stark 2016). If drivers can not trust that the reported surge multiplier is reflective of the fare multiplier they will actually receive, then surge pricing will not function as intended. In particular, by having an estimate of future surge multipliers, drivers would be better able to assess the opportunity cost of relocating. Predictions help drivers at a range of timescales. In the short-term, it may help drivers relocate in time to catch a surge. At longer horizons it may help drivers choose which trip to take now. For example, suppose a driver knows it will likely surge in 1 hour at location B 20 miles away. In choosing their next trip, they may opt, all else equal, to take the one in the direction of location B so that they are set up to meet the surge.

If potential riders had access to near-term forecasts of surge multipliers, they would be able to better plan their departure times or their mode choice. When faced with a surge multiplier, many users who have opened the app choose to postpone their trip or seek an alternate mode. Broadly, this indicates that demand for ride-sourcing services is elastic (J. Hall, Kendrick, and Nosko 2015). If riders have access to a surge price forecast they may decide to take a trip in time to avoid the surge price. In this way, surge prediction not only saves the rider money but offers the TNC a more efficient use of their vehicles.

Accurately forecasting surge multipliers is more valuable to TNCs than forecasting demand because it highlights precisely the travel demand that TNCs are not currently able to serve effectively. Because TNCs like Uber and Lyft do not employ drivers directly, they may only incentivize, rather than compel, drivers to relocate to areas of high demand. A predictive model of surge multipliers could then be utilized to help direct drivers, in real-time, to areas before the imbalance of supply and demand materializes. It may also be used, in tandem with their own proprietary data, to generate hypotheses as to why surge multipliers are predictable at all and what interventions may be necessary to better allocate the fleet.

A better understanding of the dynamics of demand and supply within ride-sourcing services would enable the development of better tailored policies for regulating them. Government agencies have adopted a wide-range of policies for regulating these now-ubiquitous services. A 2017 report by the San Francisco County Transportation Authority found that Uber and Lyft accounted for 15% of all vehicle trips in San Francisco (San Francisco County Transportation Authority 2017). The governments of London, Denmark, and Germany have banned them entirely (Satariano 2017). In 2015, New York City mulled a licensing freeze (Lapowsky 2015).

This chapter introduces a real-time spatio-temporal predictive model of Uber and Lyft surge multipliers based on measurements of urban social-physical systems and traffic characteristics in the recent past. We build a location-dependent log-linear model that takes a comprehensive set of spatio-temporal features extracted from multiple sources over the past hour to predict the future surge multiplier at a fixed time horizon.

Our goal is to investigate the extent to which the evolution of the imbalance of supply and demand in ride-sourcing systems in the immediate future could be recovered from features of the urban social-physical systems at a specific time, so that it helps inform both riders and TNC drivers for their better decisions without necessarily engaging TNCs. This goal should be contrasted with those of theoretical models which aim to explain the decision making process of individuals or groups by identifying variables which cause them to, on average, change their behavior. This is an important vein of research that can help economists and policymakers analyze the ride-sourcing market. However, a keen understanding of driver and/or rider behavior on average is of limited use in forecasting how surge pricing will actually change over the next few hours.

Our work departs from prior work in five ways:

1. We build a data-driven model of surge multipliers, an operational characteristic of ride-sourcing services. Other work either builds theoretical models of surge pricing or builds data-driven models of characteristics of general traffic, such as demand or traffic flow.
2. Our model uses a broadly collected, multi-source data set, combining data from multiple disparate systems across the public and private sectors, such as traffic speed, events, road closures, and weather conditions. We employ L_1 regularization to utilize the entire feature set without over-fitting. Most existing work uses data from a single source.
3. Our model uses spatio-temporal features available in real-time from a regional network. The feature set contains measurements at multiple locations and time points, fully exploring the spatio-temporal relations among those features and surge prices.
4. Our model is capable of being run in real-time and is able to predict surge multipliers up to two hours in advance—a relatively long prediction horizon compared to other real-time predictive models.
5. This chapter offers a head-to-head comparison of Uber and Lyft surge pricing.

This chapter predicts the surge multiplier up to two hours in the future based on the last hour of observed surge multipliers as well as the last hour of observed social-physical features. To our knowledge, real-time predictive models of surge pricing have not been explored in transportation literature. Moreover, synthesis of data across real-time APIs including Uber, weather services, and municipal services to serve as features in a predictive model is also unexplored.

The rest of the chapter is organized as follows. Section 3.2 describes the data used in this chapter. Section 3.3 presents the model formulation. Section 3.4 presents the quantitative and qualitative results of the predictive model on real-world data. Section 3.5 discusses several implications of our results for policymakers. Finally, Sec-

tion 3.6 summarizes the results and contributions of this chapter as well as potential extensions.

3.2 Descriptive analysis

3.2.1 Surge multipliers

Surge multipliers are a rare occurrence; over 90% of 10-minute windows over all locations during the study period had no surge (e.g. had a surge multiplier of 1.0). The reported surge multiplier is discretized from an underlying continuous process (M. K. Chen and Sheldon 2016). Three artifacts of the discretization process are visible in the histogram in Figure 3.1.




Figure 3.1: Probability of surge multiplier occurrence by surge multiplier value in log-scale.

First, Uber and Lyft have different rounding preferences. Uber prefers to round to the first decimal place whereas Lyft prefers to round to increments of 0.5 although both use a finer discretization for small surge multipliers. Second, despite their respective preferences, other values of surge multipliers do occur, but much less frequently than otherwise expected. For example Uber surge multipliers of 1.2 and 1.3 occur approximately two orders of magnitude more frequently than a multiplier of 1.25. Lastly, both Uber and Lyft surge multipliers are capped: Uber at 4.9 (with two specific exceptions) and Lyft at 8.

Uber surge multipliers greater than 4.9 occurred only twice: once in the early hours of Sunday October 30 2016 and once in the early hours of New Year’s Day 2017 (also a Sunday). In both instances nearly half of the 49 locations experienced a surge greater than 4.9. Interestingly, the locations experiencing large surges were

nearly the same on both days. This suggests that there is a maximum surge value for Pittsburgh that may be lifted in certain circumstances. The data from these two instances were omitted from the training data.

Entropy of Uber and Lyft surge multipliers





Figure 3.2: Entropy of the empirical surge multipliers distributions over day of week and time of day over all locations. Daily patterns in the early morning and during the mid-day during the week are clearly visible.

Surge multipliers exhibit some daily and weekly temporal periodicity. Entropy is a measure of spread in a distribution which is used in this setting to identify times when large surge multipliers are more likely. In this case, low values of entropy correspond to periods of low surge, whereas high values correspond to times when larger surge values are more likely. Several patterns can be seen in the entropy of the surge multipliers shown in Figure 3.2. First, the mid-day during the week (Monday through Friday) experiences the least amount of surging. Second, the AM peak surge around 8:30 AM is visible on weekdays. Third, surge multipliers increase around 2:00 AM, particularly on the weekend. Fourth, the PM peak is visible during the week, but not as pronounced as the AM peak. This is consistent with traditional characterizations of commuter demand in literature, for example in Cain, Burris, and Pendyala (2001). Lastly, Friday and Saturday evenings are marked with a higher propensity to surge.

Some locations are simply less likely to experience surges than others: for both Uber and Lyft the top six highest surging places will surge around twice as often as average. In general, Lyft and Uber are likely to surge at similar locations and at similar times. For both Uber and Lyft, surge multipliers tend to occur in densely populated areas near the urban core. Both services are between 2 to 3 times as likely to surge in urban areas than in suburban areas. The surge multipliers of two services are also well-correlated in time (Figure 3.3).

Compared to Uber, Lyft is slightly more likely to surge in the central business district (CBD) and slightly less likely to surge in populous areas just outside the CBD. Moreover, Lyft surges are well correlated at small negative lag times, implying that

Temporal correlations of surge multipliers



(a) Cross-correlation of Uber and Lyft surge multiplier at three locations. The surge multipliers are best correlated at 0 lag. However, the asymmetry indicates that Lyft surges more often lag Uber surge than not.

(b) Auto-correlation of Uber and Lyft surge multiplier at the Benedum Center. In general, Lyft surge multipliers are more linearly predictive of future surge multipliers than Uber surge multipliers are.

Figure 3.3: Auto and cross correlations of Uber and Lyft surge multipliers.

Uber surges are generally a leading indicator of Lyft surges.

We summarize our comparison of surge pricing between Uber and Lyft below.

- The surge multipliers on the Uber and Lyft platforms are spatially and temporally well-correlated. Uber surge multipliers are a leading indicator of Lyft surge multipliers over short time horizons (< 10 minutes).
- Uber and Lyft prefer different discretization. Uber prefers to round their surge multipliers to the nearest 0.1 whereas Lyft prefers to round to the nearest 0.5. As discussed in Section 3.4, we hypothesize that this choice has a large impact on model performance.
- Both Uber and Lyft appear to have a maximum surge price in the Pittsburgh market: Uber at 4.9 and Lyft at 8.

3.2.2 Traffic speed measurements

Traffic management channels (TMCs) measure road segment speed. Roadways have different characteristics which impact the speed at which vehicles may travel them. To compare speed measurements across TMCs, the measured speeds from each TMC have been (separately) normalized to zero mean and unit variance. The first observation is that once normalized, speed measurements across TMCs are highly temporally correlated. And on average, daily congestion cycles can be readily observed during

the AM and PM peaks. As expected, congestion is most severe in the urban core and less severe in the suburbs.

3.2.3 Road closures

Road closures as reported by the PennDOT Road Condition Reporting System (RCRS) contain regularly scheduled road work as well as unexpected incidents such as collisions. For the purposes of this study, all closures are treated equally. On average, road closures are expected during the mid-day during weekdays. Moreover, certain tracts are far more likely than others to experience road closures. The urban core, relative to much of the surrounding area is slightly more likely to experience closures. Roadways near the airport are particularly prone to road closures.

3.3 Methodology

Graphical representation of the model




Figure 3.4: Conceptual graphical model of our predictive surge multiplier model. Solid arrows represent dependence of variables. The dashed arrow represents the data-driven model we build in this chapter.

Regardless of whether or not ride-sourcing companies use predictive models in their surge pricing algorithm, surge multipliers are fundamentally a spatio-temporal function of vehicle supply and user demand. However, vehicle supply and user demand in ride-sourcing systems are currently considered trade secrets and are not made available publicly in any form. Therefore, we may not consider them directly, nor train a model to estimate them. Succinctly, this model represents a choice made in the face of a trade-off between estimating the direct factors of surge multipliers and measuring indirect ones. The principal drawback of estimating the supply and demand individually is that the errors in estimating each compound in estimating the

surge multiplier. By skipping the intermediate modeling step, we remove the source of compound error which leads to more accurate predictions. Figure 3.4 provides a graphical representation of the predictive model introduced in this chapter. Our model leverages the environmental features which have been shown to influence supply of and demand for travel in ride-sourcing systems.

Of course, supply and demand in ride-sourcing systems also likely depend on socio-demographic, land-use, and other location-specific features. Our methodology implicitly models location-dependent features by training a separate model for each location at which data has been collected. Any spatial feature which is not also time-varying would thus be a constant in the training set and could not meaningfully contribute to the model.

As a result of fitting a separate model for each location, the model trained at a location generalizes to new data collected at that location: ours is a *transductive model*. In order to predict well on a new location the model must be trained with data collected about that location. It may not simply use the fitted parameters from another location. There are two main reasons for this. First, many features used in the model are not stationary over long periods of time. The distributions of weather features, for example, are not stationary over the year. Events are similarly non-stationary: sports have on and off seasons and outdoor events are less common in winter. On short time scales however, the assumption of stationarity is reasonable for many features including weather measurements. By limiting the training set to only data collected to a finite time period prior to the held-out test data, we are able to apply a model which assumes stationarity. Second, locations may exhibit idiosyncratic relationships among social-physical features and ride-sourcing supply and demand. By limiting the training set to only one location per model, we implicitly model the complex and multi-faceted relationship between location-specific characteristics (e.g. land use, demographics).

Because the surge multipliers are bounded below by 1, linearity is likely an unrealistic assumption. Specifically, the transition between no-surge and surge is likely not linear. We address this issue in three ways. First, we predict the log of the surge multiplier. Even though surge multipliers are bounded below by 1, predicted surge multipliers between zero and one can be thought of as a discount when supply outpaces demand. In practice, instead of setting a discount to boost demand when supply is too high, Uber and Lyft rely on drivers to either relocate or log off of the system when they are having a hard time finding passengers. Predicted surge multipliers less than zero, however, have no meaningful interpretation. By using a log-linear model predicted surge multipliers are always positive. This has the additional effect of reducing the impact of training error on those samples for which the model predicts surge multipliers less than one.

Second, we compare the performance of the log-linear model to non-linear methods. We find that the log-linear model generally out-performs the non-linear meth-

ods. In this case the linearity constraint combined with the L_1 regularization yields a model that generalizes better than more complex non-linear models despite the fact that surge multipliers likely do not truly depend log-linearly on the features.

Third, we cluster the surge multiplier time series at each location and use the distance to cluster centroid, based on the last hour, as a variable in the model. This helps differentiate no-surge from surge in the relatively common case when there is simply no surge multiplier for most of the day at a location.

In short, we claim that our methodology is able to ingest five month’s worth of data about a specific location, and produce a fitted model capable of reliably predicting the short-term evolution of the surge multiplier at that location given data collected in the following week.

We outline the full procedure for training the model below, followed by full descriptions in each subsection.

1. *Data collection.* All data used in this model are available in real time from web APIs. Data is collected using a straightforward python script¹ which requests data at regular intervals.
2. *Data processing.* The API response data must first be parsed then re-indexed to a common time series—in our case, even 10 minute intervals. Spatial data is aggregated to its containing US Census Tract.
3. *Temporal segmentation.* To account for time-of-day heterogeneity, data is segmented into heuristically defined time-of-day windows.
4. *Clustering.* The time series of surge multipliers is clustered within each time of day window. Clusters capture natural patterns in the data over longer time periods. In particular, clustering obviates the need to include day-of-week variables. Distance to each of the cluster centroids is then included as a feature in the model.
5. *Training.* A linear model with L_1 regularization is trained on Uber and Lyft data, respectively, using the scikit-learn Python package (Pedregosa et al. 2011). A separate linear model is trained for each time-of-day window, location, and prediction horizon (Δt). Two-level cross validation is used to select the L_1 penalty and evaluate the model. Each training set consists of five consecutive months of data. Each training example contains the last hour of data (from time $t - 1h$ to time t) and is used to predict the surge at time $t + \Delta t$. The following week is held out for validation. Mean-squared percentage error is used to measure the performance of the fitted model on the validation data.
6. *Prediction.* Predictions for a full day are made by concatenating the predictions of each time-of-day model.

¹e.g. <https://github.com/mbattifarano/mac-data>

3.3.1 Data collection

The Uber ride pricing data contains data on 5,769,456 requests on the Uber platform from September 2016 through the end of March 2017 from 49 locations in the greater Pittsburgh area. Those requests were queried from real-time Uber APIs².

Event data contains 19,052 public events within a 20-mile radius of Pittsburgh from September 2016 through the end of March 2017 collected from four event aggregation websites (Eventbrite³, Eventful⁴, LotaData⁵, and SeatGeek⁶). Some cleaning was applied to the public event data to remove erroneous or irrelevant entries. For each time point, events were then aggregated as counts within their containing census tracts.

Pennsylvania Department of Transportation road condition reporting system⁷ (RCRS) data contains road conditions reports from 89,399 incidents in the State of Pennsylvania from August 2015 through mid-May 2017. Only road closure incidents are used. Counts of road closures are aggregated to their containing tract.

The INRIX road segment speed data contains 133,958,103 speed measurements from 1,923 traffic management channels (TMCs) in Allegheny County from September 2016 through the end of April 2017. The data were obtained through the Pennsylvania Department of Transportation (PennDOT). Speed measurements were averaged to every five minutes. All records have either a speed measurement or an average speed and nearly 99.5% of records have both.

Weather Underground is an online weather forecasting service which provides a web-based application programming interface (API) to historical weather conditions⁸. Hourly measurements of temperature, dew point, humidity, wind speed, wind gust speed, visibility, pressure, windchill, heat index, precipitation, fog, rainfall, snowfall, and qualitative weather condition were collected within the study period. The weather data is measured hourly from a single station so that the measurements do not vary spatially in this study.

3.3.2 Data processing

First, the study period was defined as the six-month period from October 2016 through the end of March 2017. The Uber and Lyft surge responses within the study period were aligned to even 10-minute increments by rounding the timestamp to the nearest

²<https://developer.Uber.com/docs/riders/references/api>

³<https://www.eventbrite.com/developer/v3/endpoints/events/>

⁴<http://api.eventful.com/docs>

⁵<https://docs.lotadata.com/apis.html>

⁶<http://platform.seatgeek.com/>

⁷<https://www.penndot.gov/Doing-Business/OnlineServices/Pages/Developer-Resources.aspx>

⁸<https://www.wunderground.com/weather/api/>

10-minutes. Samples that were rounded to the same timestamp were resolved by taking the maximum of the surge multipliers. Timestamps that had no samples rounded to them were interpolated by taking the maximum of any samples in a 20-minute window centered at the timestamp. The convex hull of the locations of the INRIX Traffic Management Channels (TMCs) was computed and determined the study region; only those surge locations inside this convex hull were considered. The precipitation, wind speed, and temperature data were collected hourly, then joined to the surge data that occurred during the same hour.

The Road Condition Reporting System (RCRS) data contains several records for each incident, describing the progression of the road incident from the initial closure through the road or lane’s reopening. For each timepoint in the study period, counts of ongoing road closure events were aggregated to the census tract level. For each timepoint in the study period, counts of ongoing public events were aggregated to the tract-level.

The traffic speed data contains speed measurements from TMCs throughout the Pittsburgh region. Speed data was first rounded to the nearest 10 minutes. For each timepoint in the study period, TMCs were aggregated to the tract level and the mean speed was computed over the TMCs in each tract.

For each timepoint in the study period, counts of ongoing public events were aggregated to the tract-level.

Because we train one model per location, the feature dataset is indexed by place label l and timestamp t . Each feature is then scaled to zero mean and unit variance. Δt represents the prediction horizon (such as 10 min, 1 hour, or 2 hours). Each example contains the following features,

- The historical mean of the surge multiplier at each of the 49 places at the hour and minute of time t taken over all prior days.
- A vector containing surge multipliers at each of the 49 places at time t . (49 variables)
- A vector containing the average speed over the TMCs in each census tract at time t . (503 variables)
- A vector containing the number of ongoing road closures over the segments in each census tract at time t . (503 variables)
- A vector containing the number of ongoing public events over the venues in each census tract at time t . (503 variables)
- The temperature in Allegheny county at time t . (scalar)
- The wind speed in Allegheny county at time t . (scalar)
- A boolean indicating whether or not is it precipitating in Allegheny county at time t (scalar)

- A one-hot encoding of the weather condition in Allegheny county at time t . (10 variables)

Note that none of these features actually depends on the place label, l , at which we are trying to predict the surge multipliers. In other words the same exact data are used to train the models at each location. This is advantageous. Since we are fitting the parameters for data collected at each place separately, we can directly compare the values of the parameters to gain a more robust understanding of how surge multipliers behave at the network level.

3.3.3 Temporal segmentation

Surge multipliers exhibit a highly non-linear relationship with respect to time of day. To prevent this non-linearity from degrading the performance of the linear model, time-of-day windows were defined to segment the day into behaviorally distinct periods. The boundaries of each segment were based on the entropy of surge multipliers by time of day. Entropy is a particularly useful statistical property of surge multiplier distributions: larger values of entropy correspond to longer-tailed surge multiplier distributions. As a result, entropy is a proxy measure for surge activity and window boundaries are defined at the approximate times at which the entropy of the surge multiplier distribution increases or decreases rapidly. Figure 3.5 shows the entropy of the surge multipliers at each time point over all locations and all days as well as the chosen time of day window breakpoints. This process resulted in six time of day windows: Early Morning (03h-06h), AM Peak (06h-09h), Mid-day (09h-16h), PM Peak (16h-18h), Evening (18h-21h), and Late Night (21h-03h+1day).




Figure 3.5: Daily surge multiplier entropy labeled by time of day segments. Segments were chosen to divide the day into segments of distinct patterns in entropy.

When fitting the model, the training set is first split by the time of day of target. Concretely, suppose we have features for a location l indexed by time t . We are trying to predict the surge multiplier at location l at a horizon of Δt . Our training set is then given by $\{(y_{l,t+\Delta t}, X_t)\}_t$. This set of (target, feature) pairs is partitioned by the time of day window in which $t + \Delta t$ resides. For example, if we wish to make predictions 2 hours in advance, then a prediction for the surge at 18:30 would use the Evening model, even though the features would come from the hour between 15:30 and 16:30.

3.3.4 Clustering

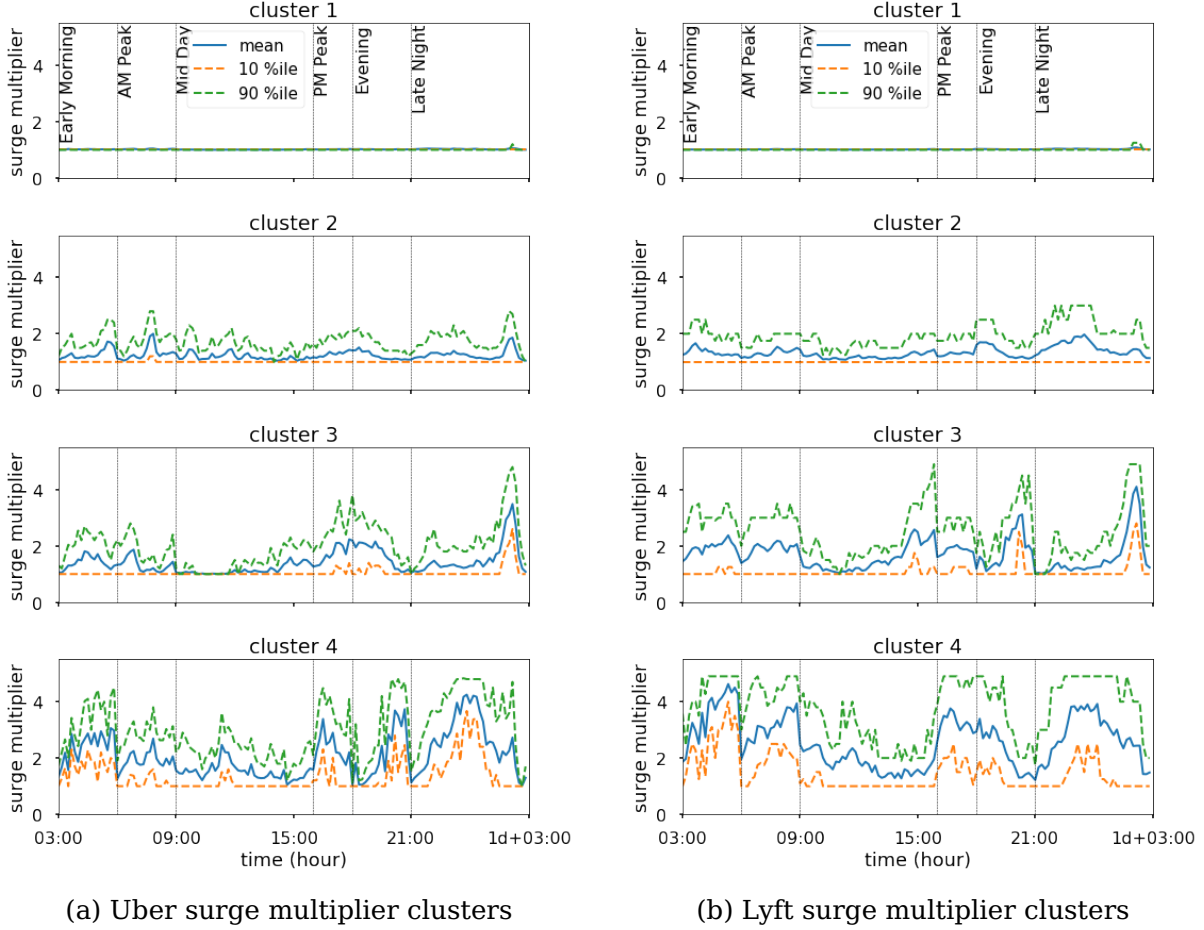
Surge multipliers exhibit non-linear behavior with respect to the day of the week caused by changes in underlying activity patterns. In particular, distinct modes exist for weekdays and weekends. Instead of modeling day of week explicitly, we use unsupervised learning to partition the time course of surge multipliers into behaviorally distinct modes. The resulting clusters will implicitly model the day-of-week modes as well as their exceptions (e.g. holidays and special events). K-means clustering was performed on the time-course of surge multipliers separately for each time window using the scikit-learn Python package (Pedregosa et al. 2011). The clustering reliably recovers the no-surge mode as the largest cluster in each time window. Clustering within time windows allows time-window specific modal behavior to be recovered. For example, the clusters in the late night time window (the right-most time window in Figure 3.6) recover three distinct behaviors that are wholly different from behaviors observed at any other time of day. By incorporating cluster information into the model, non-linear longer-term temporal surge behavior can be exploited to refine linear predictions.

Interestingly, cluster membership is nearly independent across time of day bins in the same day. In other words, knowing the cluster of today’s surge multipliers during the mid day tells us almost nothing about which cluster this evening’s surge multipliers are likely to fall in. We establish this result by noting that adjusted mutual information between cluster labels of different time windows on the same day is small (< 0.3). This implies that predictions can be made independently across time of day windows and offers a heuristic justification for this particular temporal segmentation.

3.3.5 A log-linear model with L_1 regularization

Equation (3.1) defines the linear model for predicting the log of the surge multiplier at location l with prediction horizon Δt where t indexes time. Cluster distance measures the euclidean distance between each cluster centroid and the time course of the surge multiplier up until time t . Window elapsed measures the fraction of the time-of-day window that will have elapsed at time $t + \Delta t$. The historical mean is the average surge multiplier at each location at the hour and minute of $t + \Delta t$ over all previous days. The

Figure 3.6: Surge multiplier time-of-day clusters



K-means clustering ($k = 4$) performed on each time window. Clusters of daily Uber surge multipliers are shown in 3.6a from largest cluster (top subfigure) to smallest (bottom subfigure). Similarly, 3.6b shows clusters of daily Lyft surge multipliers from largest (top subfigure) to smallest (bottom subfigure). Blue line shows the mean of each cluster and dashed lines show 10-th and 90-th percentiles. We can assess the relative quality of clusters by noting where the 10 and 90 percentile lines differ from the no-surge event represented in each time window by the largest cluster (top subfigure).

final term can be interpreted as the state of the urban social-physical system over the last hour. It includes, for each of the last six timepoints (i.e. the last hour), the surge multipliers at all locations, the weather conditions, the traffic speed measurements, and the ongoing events aggregated to the tract-level.

$$\begin{aligned}
\log(\text{surge}_{t,t+\Delta t}) \sim & \sum_{i=1}^4 \left[\text{cluster distance}_{t,i} \right] + \text{window elapsed}_{t+\Delta t} + \sum_{l'=1}^{49} \text{historical mean}_{l',t+\Delta t} \\
& + \sum_{t'=t-1\text{hr}}^t \left[\sum_{l'=1}^{49} \text{surge}_{l',t'} + \sum_{i=1}^{503} (\text{mean speed}_{t',i} + \text{n events}_{t',i} + \text{n closures}_{t',i}) \right. \\
& + \text{temperature}_{t'} + \text{wind speed}_{t'} + \text{is precipitating}_{t'} \\
& \left. + \sum_{j=1}^{10} \text{weather condition}_{j,t'} \right]
\end{aligned} \tag{3.1}$$

In total, there are 8898 features. In order to avoid overfitting, L_1 regularization is employed during training (Park and Hastie 2007). L_1 regularization encourages the model to concentrate the weights on only a small number of the most relevant features. The number of selected features can be tuned by adjusting the L_1 penalty hyper-parameter.

Two-level time series cross validation is employed during training. The outer level evaluates the model performance by training the model on moving windows of twenty consecutive weeks and evaluating the trained model on the (held out) subsequent week. The inner level selects the best L_1 penalty weight by cross validation on the training set (Baraniuk 2007). The two-level cross validation ensures that the model is evaluated on data that was held out from both training *and* L_1 penalty tuning.

time series cross validation differs from other methods of cross validation in that it ensures that the held out data occurs *after* the training data. This is a particularly important consideration for this data set as the features contain past surge multipliers. In short, training a model on the future and evaluating it on the past is not fair. Moreover, we found that time series cross-validation was critical in the inner level to select a L_1 penalty term that was large enough to generalize to the validation data.

The performance of the log-linear model is then compared to two naïve methods: the overall mean and the historical mean. The overall mean is simply the average surge multiplier at the location being predicted over all previous days. The historical mean is the average surge multiplier at the location being predicted at the hour and minute of the time being predicted taken over all previous days.

3.4 Results

We train and evaluate the log linear model on Uber and Lyft surge data for all 49 locations in the greater Pittsburgh area. We focus our discussion of the results on three locations: the Benedum Center, South Side, and Shadyside. These locations

were selected because they are likely to surge and are geographically representative of different parts of the urban areas of Pittsburgh. Predictions under the three models are compared to actual surge multipliers in Figure 3.9. At these locations we also fit the following three non-linear models and show that the lasso model generalizes the best to unseen data.

1. Random Forest trained on the features selected by Lasso.
2. Random Decision stumps trained on all of the features with the maximum tree depth hyper parameter chosen by cross validation.
3. Neural Network with two ReLu hidden layers of width 32 and 16 respectively and with L2 regularization parameter chosen by cross validation trained on the features selected by Lasso.

We find that the proposed methodology generally out performs both the naïve and non-linear methods.

3.4.1 Model performance on Uber surge multipliers

When trained on Uber data, lasso regression is able to out-perform both the historical mean and overall mean for 46 of the 49 locations in the greater Pittsburgh area when predicting surge multipliers two hours in advance. Further, it out-performs all three non-linear methods in 28 of the 49 locations. We evaluate the performance of each model on unseen data by the *mean squared percentage error* defined in Equation (3.2) below. The model out-performs naïve methods in all locations in the urban areas of Pittsburgh and out-performs naïve methods in all but 3 of the 18 suburban locations as shown in Figure 3.7. To quantify the relative performance of two methods we compute the *model advantage*, defined in (3.3), of the linear model with respect to the naïve predictor as the difference between the mean squared percentage errors of the two predictors on held out data. However, suburban locations both are far less likely to surge and far less likely to ever see large surges than locations in urban areas. Specifically, the model strictly out-performs naïve methods for all locations at which there is greater than a 2.5% overall chance of a surge. From the perspective of potential applications, this is acceptable; being able to predict larger, more likely surge multipliers is more valuable to drivers and riders than being able to predict smaller, less likely surges.

$$\text{mspe}(y_{\text{predicted}}, y_{\text{true}}) = \frac{100}{n} \sum_{i=1}^n \left(\frac{y_{\text{predicted}} - y_{\text{true}}}{y_{\text{true}}} \right)^2 \quad (3.2)$$

$$\text{model advantage} = \text{mspe}_{\text{naïve}} - \text{mspe}_{\text{lasso}} \quad (3.3)$$

To characterize how well we expect our methodology to perform on a new location,

Model advantage by location

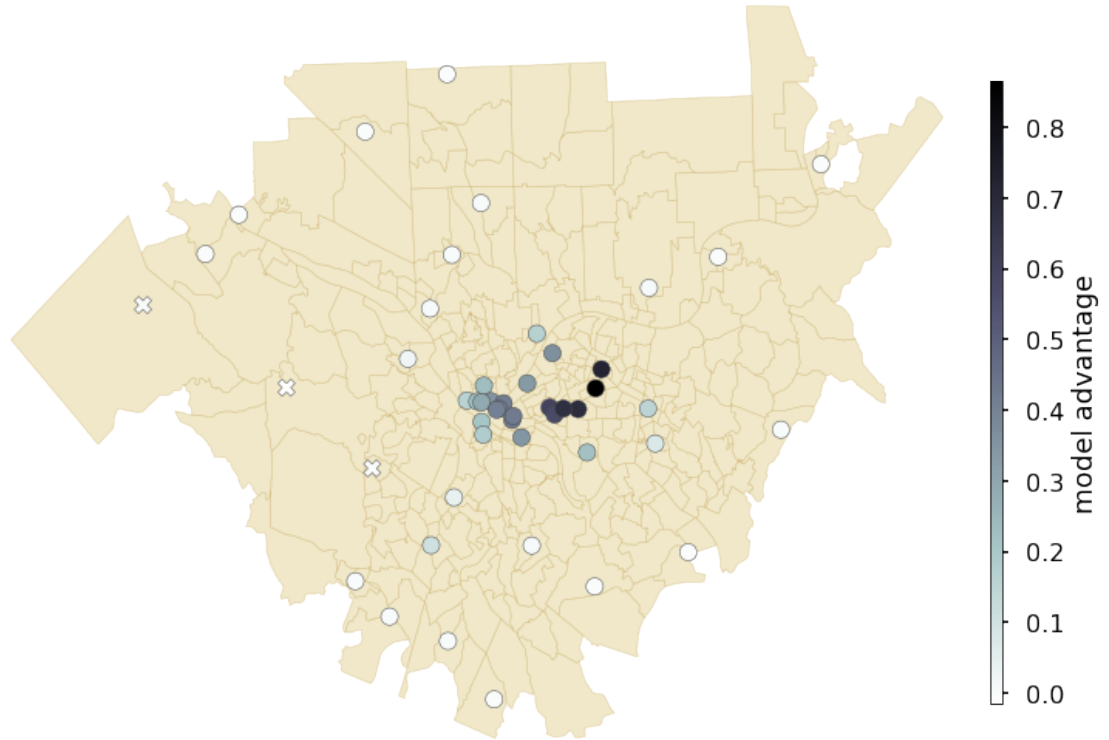


Figure 3.7: Model advantage by location at a 120 minute prediction horizon. Model advantage is defined as the difference between the mean squared percentage error (mspe) of the overall mean and the mspe of the linear model. Positive values mean that the model out-performed the overall method. X shaped markers denote locations where the model strictly under-performs the overall mean.

Table 3.1: Results of the linear regression of land use data onto model advantage.

	R^2	0.782		Adjusted R^2	0.772	
	F-statistic	80.52		Prob. F-statistic	1.36e-15	
	coefficient	std err	t	$P \geq t $	[0.025	0.975]
Low-Density Residential	-0.1848	0.015	-12.328	0.000	-0.215	-0.155
Commercial	0.2218	0.023	9.458	0.000	0.175	0.269

socio-demographic data from the US Census and Land Cover data⁹ were used in a

⁹<https://data.wprdc.org/dataset/allegHENY-county-land-cover-areas>

linear model to explain the spatial variation of model advantage. We find that the composition of land use in the surrounding area determines how well the model will perform. Our model will tend to perform better than naïve methods for locations with a large percentage of commercial land and worse on locations with a larger percentage of low-density residential land. We found that population density, per capita income, and the racial composition of the tract containing each location had a statistically insignificant impact on model advantage. Intuitively, this makes sense as we have no reason to believe most Uber trips originate near the user’s home. Moreover, several locations are in tracts in which very few people actually live. For example, the tract containing Heinz field in the north side of Pittsburgh has no residents according to the US Census. In contrast, as Table 3.1 reports, 78% of the variance in model advantage was explained, with statistical significance at the 0.001 level, by just two land cover variables: commercial land area and low-density residential area, both as fractions of the tract containing the location. In line with our intuition, the fraction of low density residential land negatively impacts model advantage while the fraction of commercial land positively impacts model advantage.

Comparison of model error at three locations




Figure 3.8: Comparison of the mean squared percentage errors of Uber models on held-out data in three active surge locations in Pittsburgh over seven prediction horizons. For Uber surge multipliers, the model predictions outperform the historical and overall mean as well as the non-linear methods.

Model performance on unseen data varies by location and prediction horizon, as shown in Figure 3.8. Intuitively, mean squared percentage error tends to increase with prediction horizon. In each of the shown locations—each in a different part of the urban core of Pittsburgh, the error remains below those of both naïve methods.

Interestingly, the overall mean out-performs the historical mean for these locations. This is due to the overwhelming predominance of times at which there is no surge multiplier and emphasizes the lack of strong daily periodicity in surge multipliers.

The proposed methodology also out-performs the three non-linear methods for 28 of the 49 locations. However, lasso remains the best choice overall: when non-linear methods out-perform lasso, they do so with small margins. On average, when non-linear methods out-perform lasso, they reduce mean-squared percentage error by an average of 0.018%. In contrast, when lasso out-performs any of the non-linear methods it reduces mean squared percentage error by 0.08%. We hypothesize that the relatively poor performance of the non-linear methods on held-out data is due to over-fitting. The increased expressiveness of the non-linear methods, even when regularized, ends up learning noise in the training set. In certain cases, the non-linear model fails to out-perform naïve methods. The strict regularization of the L_1 penalty paired with the lower complexity of a linear model lends this method to more robust predictions.

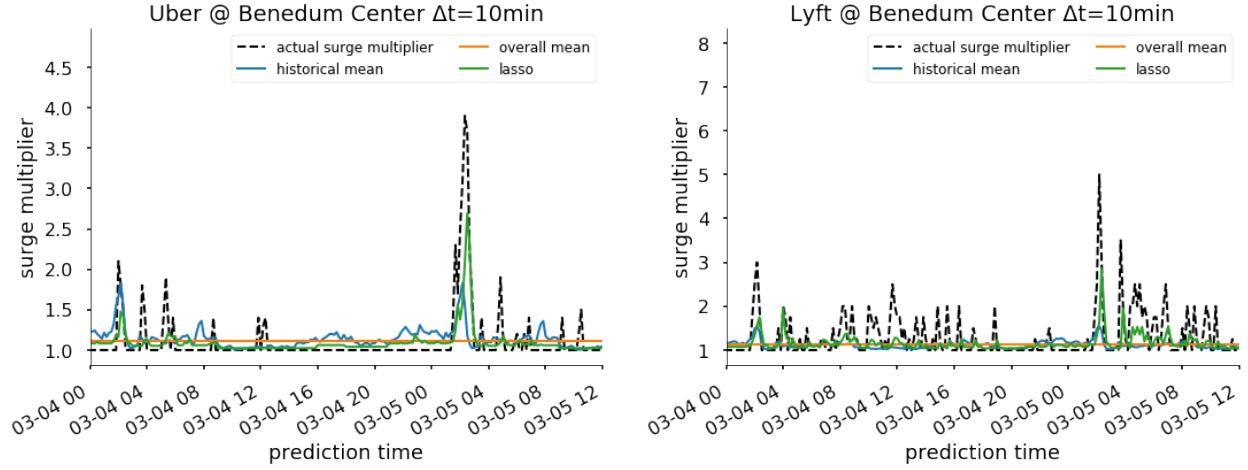
Model performance also varies by time window. Because the surge multiplier distribution changes over time windows, the performance of the naïve methods vary as well. This is particularly evident in the mid day time window, where the overall mean achieves a mean squared prediction error of nearly 2%. The same prediction yields around 3.75% error during the AM peak window preceding it. The historical mean robustly out performs the overall mean in the AM peak, PM peak, and Late night, when surges are most active, but under-performs the overall mean during the mid day and evening windows when large surge multipliers are less likely.

3.4.2 Model performance on Lyft surge multipliers

In contrast to our model’s performance on Uber data, when trained on Lyft data, the model only slightly out-performs the naïve methods. This is surprising considering that Uber and Lyft surge multipliers are well-correlated. The proximal cause is that the Lyft model has a higher prediction error compared to naïve methods on low surge multipliers (less than 1.5) than the Uber model does. Since both Uber and Lyft are not surging in over 90% of samples, poor performance on low surge multipliers is devastating to prediction error overall. In fact, the Lyft model performs as well or *better* than the Uber model for larger surge multipliers. This suggests that the non-linearity in the transition between no-surge and surge is more pronounced in Lyft than in Uber. In practice, this could be caused by differences in the sizes of the two services’ driver or user pools, different promotional practices, or any number of operational or business decisions.

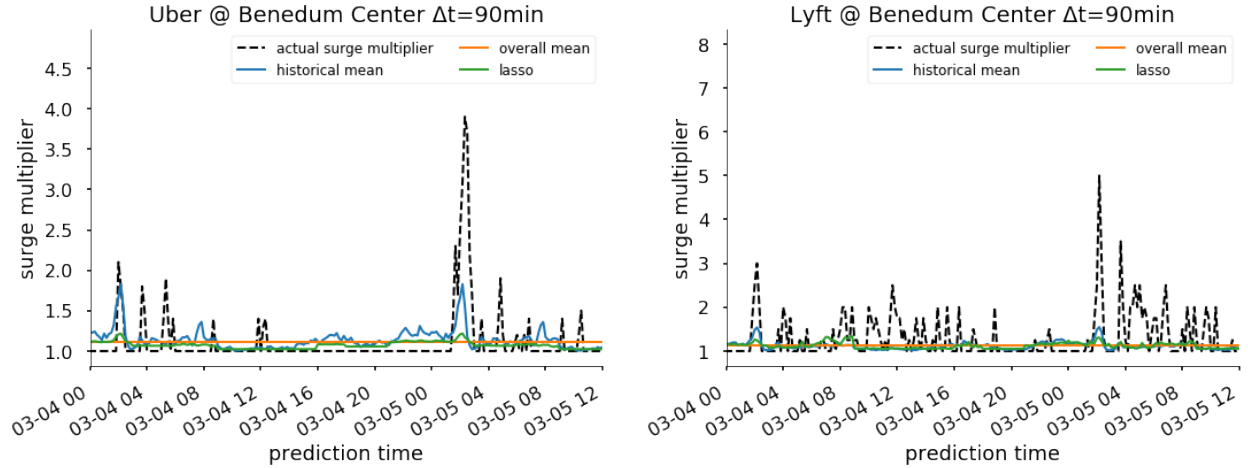
We also hypothesize that the coarser discretization employed by Lyft makes it harder to predict. Intuitively, this makes sense: the error one gets when fitting a line to a step function increases with step size. The coarser the discretization, the more

Predicted surge multipliers at a single location



(a) Predicted Uber surge multipliers at the Benedum Center 10 minutes in advance.

(b) Predicted Lyft surge multipliers at the Benedum Center 10 minutes in advance.



(c) Predicted Uber surge multipliers at the Benedum Center 90 minutes in advance.

(d) Predicted Lyft surge multipliers at the Benedum Center 90 minutes in advance.

Figure 3.9: Predicted Uber surge multipliers over three days at the Benedum Center with $\Delta t = 10\text{min}$ (Figure 3.9a) and $\Delta t = 90\text{min}$ (Figure 3.9c). Predicted Lyft surge multipliers over three days at the Benedum Center with $\Delta t = 10\text{min}$ (Figure 3.9b) and $\Delta t = 90\text{min}$ (Figure 3.9d). All models tend to over-estimate low surges and under-estimate large surges; a phenomenon which becomes more pronounced when the prediction horizon increases to 90 minutes.

difficult it is to recover a linear relationship. To validate this claim, we re-discretized the recorded Uber surge multiplier to match the discretization used by Lyft and then re-trained the model. When trained on the re-discretized data, the model performance

degrades in both absolute value and relative to naïve methods. Overall, prediction error increases by around 1%. In particular, the model performance degrades most on surge multipliers less than 1.5. Although one might expect that a random forest model would be more capable in this regard, our experiments suggest otherwise.

3.4.3 Selected features

On average, L_1 regularization selects just 6 of the 8898 features. Within each fold of the cross-validation and within each time-of-day window however, L_1 regularization selects a different set of features. There are 7 folds of the cross validation, 6 time-of-day windows, 7 prediction horizons, and 49 locations. This means there are 14406 ($= 7 \times 6 \times 7 \times 49$) fits performed and evaluated in our experiments.

To measure the significance of each variable, we compute the p values of the variables selected by each fit. There is a small body of work dedicated to statistical inference based on the Lasso estimator (J. D. Lee et al. 2016; Lockhart et al. 2014; Loftus and J. E. Taylor 2014). However, there is no well-established procedure for significance testing on the variables selected by Lasso. Because the data itself has been used to select the features, performing the usual tests for significance on the training data will result in anti-conservative p-values (Loftus and J. E. Taylor 2014). In this chapter we take a conservative approach and test the significance of the features selected by Lasso by using them in an ordinary linear model trained on the held-out data. We consider any feature with a p-value less than 0.05 to be significant. We caution that statistical significance in this setting is less meaningful than performance. Because location-specific relationships are implicitly modeled in this setting, it’s less important that we identify those variables whose influence on surge multipliers is statistically significant and more important that on average our method is able to predict surge multipliers better than other methods. That said, significant variables may give us insight into potentially reliable location-specific pre-indicators of surge multipliers. To measure overall influence of each variable, we compute *feature importance* as the average absolute weight for each feature over all folds.

In general, surge multipliers exhibit location-specific dependence on the features. As a result, a model trained in one location should not be expected to perform well in another. That each location requires a different model is an advantage of this approach because it abstracts away harder-to-measure interactions between ride-sourcing supply and demand, as well as socio-economic, demographic, and land use variables each of which contributes to the dynamics of surge multipliers. As such, we can not generalize the dependence revealed in the fitted models to other cities or even other locations within Pittsburgh. Rather, we may use these models to reveal hyper-local dynamics of the ride-sourcing market.

Across locations, L_1 regularization consistently selects from the current surge multiplier variables and the weight placed on the current surge decreases with the pre-

diction horizon. This makes sense given the strong autocorrelation observed in the surge multiplier time series at small lags. However, other features types are selected with different importance at different locations, re-affirming the need to fit them separately. For example, Figure 3.10 shows that at the Benedum Center and South Side, event variables are consistently selected and their weight increases with prediction horizon. In contrast, the Shadyside model places less weight on events and more on historical averages. The patterns in feature type importance across locations is consistent with high-level neighborhood characteristics: the Benedum Center is an event center located in Downtown Pittsburgh in proximity to several other event centers. The South Side is home to bars, clubs, concert venues, and theaters. Shadyside has many restaurants and bars but fewer event venues.





Figure 3.10: Relative feature importance in the Uber model by feature type and prediction horizon. Overall, the surge multipliers were the most frequently selected feature across all three locations. However, as the prediction horizon increases, current surge multipliers tend to become less important, and more weight tends to be placed on traffic speed, events, and historical average surge multipliers. Feature importance for each feature is computed as the mean of the absolute value of the feature weight over folds. Feature importance of a feature type is computed as the maximum of the feature importance of all features of that type.


Of the cluster distance features, the distance to the first (largest) cluster was selected with the most weight. The first cluster was the “no-surge” cluster and is helpful in distinguishing periods in which there is a lot of activity from periods of no surge. Taken together with the weight placed on historical averages, this would seem consistent with the pattern of nightlife characteristic of areas like Shadyside: somewhat regular bursts of activity unrelated to events that may be going on.

To understand how robust the feature selection is, we can look at how consistent the set of selected features is across the folds of the cross validation. First, we define

Feature selection over time



(a) **Feature consistency.** Frequency with which features are selected in the Uber model over folds in the cross validation. The consistency tends to decrease with prediction horizon.



(b) **Feature drift.** Correlation coefficient between the temporal distance between two subsets of the data and the hamming distance between the features selected by models trained on those subsets.

Figure 3.11: Feature consistency and feature drift at three locations.

consistency as the mean fraction of folds in which a feature will be selected given that it has been selected in at least one fold. A consistency score of 1 means that the same features were selected in each fold, and 0 means that no feature was selected more than once. Consistency is undefined when no features are selected in any fold. Second, we examine a measure of feature drift to understand the extent to which selected features are changing over time. We define feature drift as the hamming distance between two feature sets, measuring how different any pair of selected features are. By computing the correlation coefficient between hamming distance and time we can understand the extent to which selected features are changing over time. Because the underlying data is temporal, the folds of the cross validation use chronologically sequential subsets of the data. As a result, we can meaningfully determine the amount

of time between the data used in each fold. A high correlation coefficient means that the further apart in time two subsets of the data are, the more different the two sets of selected features will be.

Although there is no theoretical relationship between statistical significance and whether or not a variable is selected by Lasso, in our experiments, we find a statistical relationship between consistency and statistical significance. Features that are significant in one fold are slightly (but statistically significantly) more likely than non-significant features to be selected in at least one other fold. Around half of non-significant features are selected in a different fold compared to between 55% and 65% of significant features, depending on location.

In general, two trends are clear. First, consistency is low and tends to decrease slightly with prediction horizon as in Figure 3.11a. Second, the feature drift is high across all prediction horizons as shown in Figure 3.11b. That feature drift is large may suggest that the importance of certain features is likely to change over long time periods (i.e. months).

The weight placed on each feature type changes slightly over folds. Event and surge variables slightly gain importance over folds whereas cluster and closure variables tend to lose weight. Within each feature type, however, feature weights shift between individual features. Within a type, this shift could be due to redundant information: for example, two tracts may have similar patterns of events with one being slightly more useful to the regression in one fold than the other.

The statistical significance of features varies by location and changes over prediction horizon. Current and historical surge multiplier features however, were consistently found to be significant across locations at short prediction horizons. The significance of current surge multipliers decreases over prediction horizon. Other features vary substantially. For example, nearly 60% of traffic speed features were found to be significant for the Shadyside location, whereas less than 10% of traffic speed features were significant at both the Benedum Center and South side. These findings support the idea that there are latent, complex interactions that vary by location. Some interactions may be idiosyncratic to a location. By fitting each location separately, these complex interactions are implicitly modeled, enabling relatively simple techniques to characterize non-trivial behavior.

Here we analyze the spatial distribution of dependencies for the three highlighted locations to demonstrate how location-specific models may be used to better understand the dynamics of the ride-sourcing market at a hyper-local level. In general, spatial dependencies tend to be geographically disparate, highlighting the importance of including non-local spatial data in the model.

Overall, each of the three location models place weight on events in similar, geographically disparate areas of the city. As previously noted, both the Benedum Center

Importance of event features

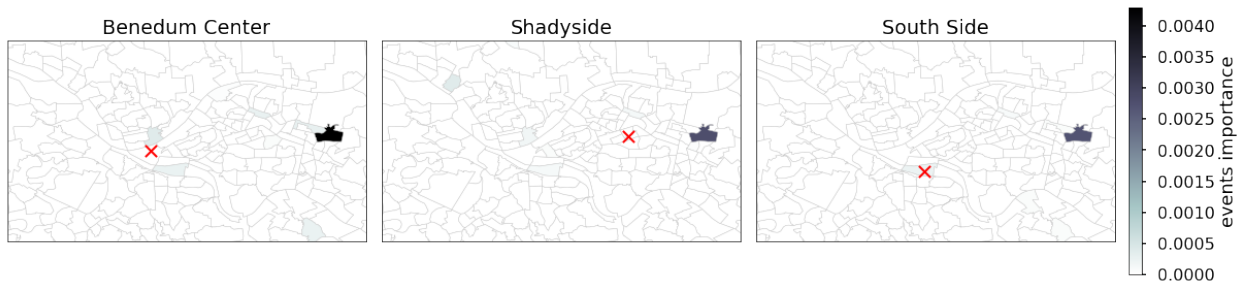


Figure 3.12: Importance of event variables in each of the models for the three selected locations at a 120-minute prediction horizon. The red 'x' indicates the geographic position of the modeled location.

and South Side leverage event variables to predict surge multipliers. Interestingly while both models place weight on the tract containing South Side, the Benedum Center model does not place weight on its own tract, preferring instead geographically disparate tracts which are less likely to host events as shown in Figure 3.12. These locations might be particularly useful indicators precisely because they are *not* event centers. In other words, events occurring downtown are rarely a surprise: drivers may even assume there is always some event going on. This view is supported by the data. The tract containing the Benedum Center is also host to several other event centers which has the effect of there almost always being at least one ongoing event making the variable completely useless, a problem that finer spatial aggregation might help to solve. In contrast, an event in East Hills (the tract at the center right of Figure 3.12) might be completely unanticipated.

Importance of speed

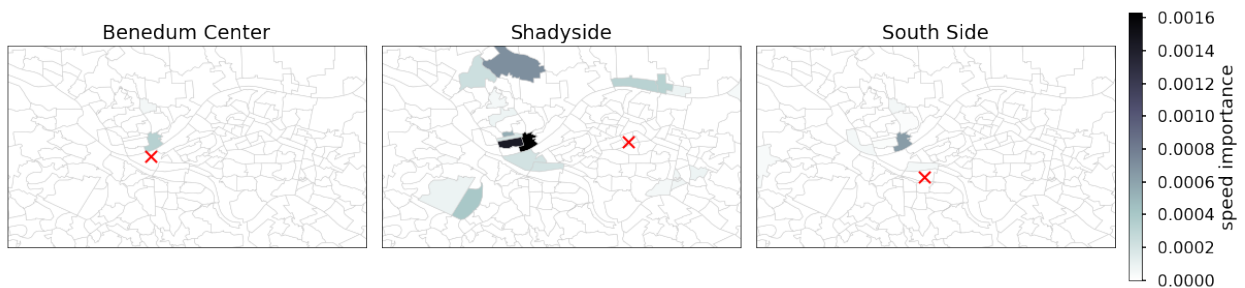


Figure 3.13: Importance of speed variables in each of the models for the three selected locations at a 120-minute prediction horizon. The red 'x' indicates the geographic position of the modeled location.

Similar to events, the models for each of the three locations place weight on traffic

speed in geographically disparate areas of the city, as shown in Figure 3.13. Each of the models for the representative locations place importance on traffic speed on the north side. This could indicate that traffic in the north side is a useful proxy for traffic conditions more broadly. Surge multipliers in Shadyside, whose model is more sensitive to traffic speed than the other two models, depend on traffic speed in a larger, geographically disparate set of tracts.

Importance of past surge multipliers

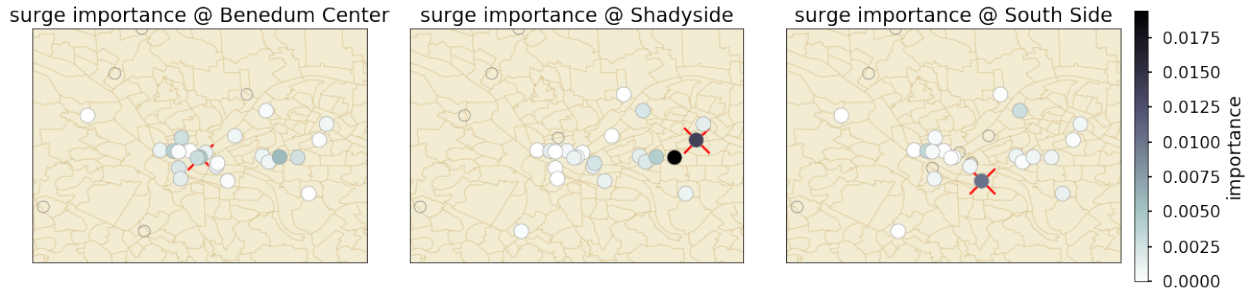


Figure 3.14: Importance of surge variables in each of the models for the three selected locations at a 120 prediction horizon. The red 'x' indicates the geographic position of the modeled location. Surge variables with 0 importance are represented as hollow circles.

The dependence of surge multipliers at one location on past surge multipliers in other locations varies the most by the location being predicted. More so than the importance of speed and events, the importance of surge variables tends to concentrate geographically in areas near to the location being predicted. However models also place weight on farther-flung locations in many directions. The three selected locations exhibit markedly different spatial dependence patterns as can be seen in Figure 3.14.

In short, because our proposed model abstracts harder-to-measure features impacting surge multipliers by training the model separately for each location using simple features, each fitted model does not generalize to other locations. However, the fitted parameters of each model reveal local relationships between surge multipliers at that location and real-time urban features across the city. As a result, the values of the fitted parameters, and their significance, may be used to better understand what can lead to a mismatch of supply and demand in ride-sourcing systems in a specific area of a city.

3.4.4 Parameter clusters

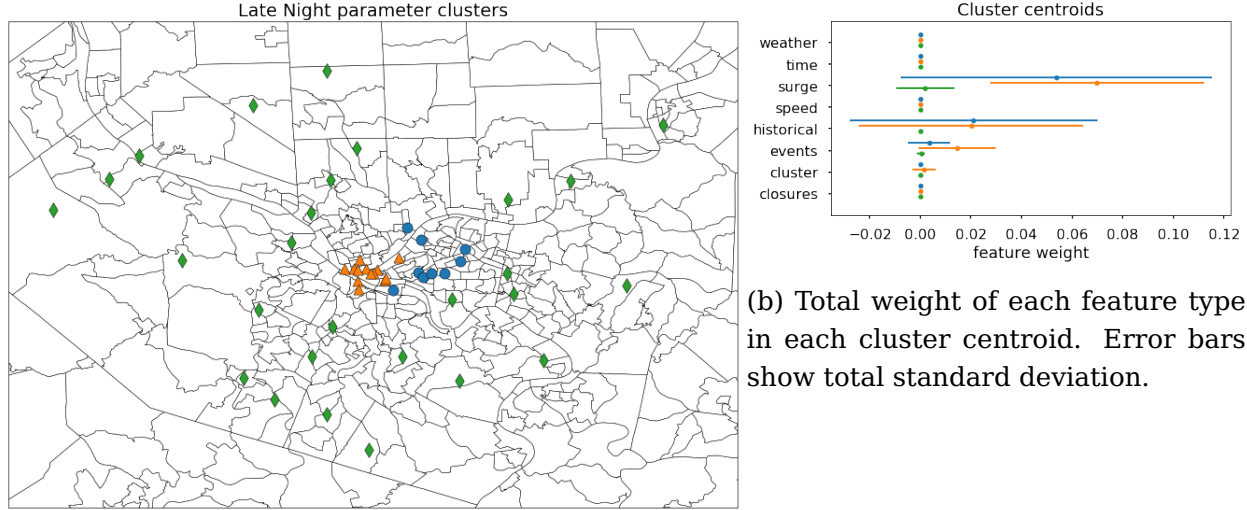
Having analyzed the selected features at three specific locations, we draw more general conclusion by exploring the structure of the fitted parameters within the param-

ter space. Structure in the parameter space indicates that the models are related and may be represented more succinctly than one model per location. As previously discussed, the fitted parameters differ between locations indicating location-specific relationships between features and surge multipliers. The nature of these relationships may be mediated by a variety of location variables including land use, socio-economic, and demographic features. Characterizing these relationships is no doubt interesting but outside the scope of this chapter. Fitting each location separately serves to improve prediction accuracy, but this choice ignores two important facts. First, surge multipliers are correlated in space, meaning that we should expect locations that are near to one another to have similar models. Second, if surge multipliers are in fact responding to location-specific characteristics, then locations of similar character should have similar models. Taken together, this perspective illuminates a possible extension to this work in which the similarity of locations in some latent space is utilized to jointly fit all locations. In particular, a multi-task learning or mixture model approach might work well.

In light of this, we cluster the separately fitted parameters of all location models in order to examine first the structure of the fitted parameters within the parameter space and second the potential of parameter sharing in this setting. From a probabilistic perspective, we are now treating the fitted parameters of each location model as random vectors drawn from a distribution in the space of parameters. Structure in the parameter space means that the parameter vectors are drawn from a distribution conditional on a smaller number of latent factors. Agglomerative clustering was applied to the fitted parameter values for each location and each time-of-day (Pedregosa et al. 2011). To improve clustering performance, principal component analysis (PCA) was first applied to reduce the dimensionality of the parameter space while preserving at least 99% of the variance (Pedregosa et al. 2011).

When rendered geographically, the clusters labels are intuitive and largely group locations by geography and land use patterns. In Figure 3.15a the green cluster can be seen to represent suburban areas while orange represents the business, entertainment, and commercial districts while the blue cluster contains mostly dense residential areas with commercial main streets. This characterization can also be seen in the cluster centroids. In Figure 3.15b, the centroid in each cluster is represented by the total feature weight within each feature type. Note that the centroid of the orange cluster places weight on events more so than the blue cluster centroid. The green cluster places almost no weight at all on events. This representation is necessary visually to avoid displaying a large number of features at the expense of some nuance. In particular, that the total weights in the centroids for the blue and orange clusters are similar does not mean that the centroids are close together. In fact the centroids of the blue and orange clusters respectively place surge and historical surge weights on different geographic areas. This indicates that both clusters represent a different

Parameter clusters



(a) Cluster labels of fitted parameters in the late night model.

(b) Total weight of each feature type in each cluster centroid. Error bars show total standard deviation.

Figure 3.15: Parameter clusters of the late night lasso models over all locations.

spatial relationship with similar magnitude between past and future surges.

Parameter clusters of the other time-of-day windows also largely group locations by geography and land use. However, because land use patterns also contribute to activity temporally, the clusters and their centroids differ between time-of-day windows. For example, there are clusters in the AM and PM peak, in contrast to the Late Night cluster centroids shown in Figure 3.15b, that highly weigh speed but not events. Clusters in the early morning weigh past surge activity almost exclusively as do clusters in the mid day.

The changing relationships also have geographic implications. Suburban locations tend to always be in the same cluster, but depending on the time-of-day, certain locations may behave similarly to suburbs from the surge prediction perspective. In the mid day for example, nearly every location looks suburban except for the areas populated by several large universities, principally the University of Pittsburgh and Carnegie Mellon University. In contrast to the late night clusters, during the AM and PM peaks, the business district separates itself into its own cluster while the commercial and entertainment districts join the suburban cluster. These observations suggest that the relationship between features of the urban environment is mediated by land use and in the temporal patterns of activity it produces.

From a probabilistic perspective, the clustering results show that the parameters of each of the linear models can be understood as having been drawn from a distribution conditioned on a low dimensional latent space representing a mixture of clusters.

Ultimately what this means is that nearly 9,000 features spanning 49 locations can be described by a handful of linear models each using on average 6 features. These linear models not only forecast surge multipliers across the city but also reveal what useful signals the driver network is failing to incorporate into its collective decision making process. Taken together this suggests that a mixture model approach or multitask learning might work well here. From a multi-task perspective the separate lasso models could be trained jointly by asserting that locations that are close in land-use space (or geographic space) should have similar parameters. This additional source of regularization could improve generalization error. Similarly, from a mixture model approach, we consider the parameters to be drawn from a distribution conditioned on the location’s cluster membership probabilities in the latent space associated with locations.

3.5 Policy implications

Although ride-sourcing riders and drivers are clear beneficiaries of real-time surge-price predictions, policymakers may also leverage the knowledge generated by our model to mitigate the impact of TNC fleets on the physical infrastructure. In most policy settings, the real-time nature of the model is not very useful because most of the actions policymakers can direct agencies to make are not performed in real-time. However, our analysis of our model demonstrates some general associations between the physical-social features and surge-pricing giving policy makers important information to tailor their interventions.

A large surge multiplier indicates not only an imbalance between those seeking a ride on the platform and those providing one, but a mismatch between demand and supply of transportation in general at that time and location. If riders had other modes available to them surge-pricing should make those modes relatively more attractive. This may in fact explain the correlation between surge multipliers on the Lyft and Uber platforms. More broadly, a large surge multiplier implies that there are riders who would rather pay the premium than find another way to complete their trip, for example, by taking public transportation. In this light, demand during periods of large surge multipliers is demand with particularly high mobility cost which may make good candidates for alternative modes.

Policies aimed at either diverting demand to alternative modes (demand-side intervention) or increasing the efficiency of ride-sourcing pickups or drop-offs (supply-side interventions) have two main effects. First, TNC-related congestion will decrease as fewer vehicles, (a result of demand-side interventions) are making pickups and drop-offs more efficiently (a result of supply-side interventions). Second, surge multipliers should decrease and/or become shorter as riders will prefer alternative modes at a lower price threshold. In motivating this work we discussed the advice given to new

drivers: “don’t chase the surge” implying that surge multipliers frequently cause more drivers to relocate than necessary. This fact implies that congestion effects of TNC fleets may be inflated during periods of surge due not only to the increase of vehicles needed to meet demand, but the increase of vehicles *in excess* of the demand.

Surge-price prediction may reveal gaps in existing public transit service. Although it is certainly infeasible for public transit agencies to respond to demand in real time, there may exist some corridors where transit frequencies could be better tuned to address predicted surges. Large events at sports arenas for example are scheduled well in advance and could be a good candidate for additional transit service before and after the event. Neighborhoods with many smaller event venues may also be attractive candidates to transit agencies because demand will be more evenly distributed in time.

In locations where transit solutions are infeasible because surges are not regular enough to meet them cost-effectively via adjustments to transit scheduling, it may be possible to leverage shared micro-mobility systems to absorb demand surges and mitigate TNC congestion. Bike and scooter share systems offer greater flexibility than public transit systems to address mobility demand. Placing stations at locations that frequently surge and ensuring that they are well-stocked when surges are anticipated may help mitigate TNC congestion during demand surges. In this setting, our model is most useful in identifying stations or areas that may need additional bikes or scooters in the next two hours. Depending on the re-balancing methods used by the shared micro-mobility service provider this may be enough time to reposition the supply of bikes.

Public transit and shared micro-mobility solutions address the demand side of surge-pricing. Where those solutions are infeasible, infrastructure changes may be used to address the supply side of surge-pricing. In particular, time-dependent curb space regulation could be targeted to areas with a tendency toward large surge multiplier predictions at certain times of day. Dedicated curb space offers more efficient pickups and drop-offs which not only reduces the congestion effect of ride-sourcing vehicles, but may also serve to reduce surges in the first place by allowing higher throughput. Dedicated curb space can open up a new revenue stream for local governments, who can charge TNCs for the service of rededicating the public right-of-way.

Policy solutions aimed at addressing surge demand may be effective in decreasing congestion and increasing ridership but may not be particularly equitable. These solutions must be balanced against equity goals in assessing their feasibility. Although the demand-side policy solutions discussed benefit everyone seeking mobility services, the fact that TNC demand is used to identify the times and locations for these interventions introduces a bias toward better mobility for those who use TNC services. Compared to those who do not use TNC services, TNC riders tend to be younger, wealthier, are more likely to be White or Asian, and less likely to have a physical

disability (Grahn et al. 2020; Cochran and Chatman 2021). Deriving mobility interventions based on TNC riders data is therefore likely to overlook the elderly, low income, and minority populations as well as those with physical disabilities; exactly the populations equitable transit service should address. Any intervention based on TNC riders should be paired with policies aimed to increase mobility access among non-TNC users. Additionally, the associations between surging demand and the social-physical features may remain helpful in identifying locations and times where greater mobility services are wanted even in regions where TNC utilization is low. Whether or not the model or insights derived from it may be used in this way is left for future work.

3.6 Conclusion

Data describing the current state of several social-physical systems are collected from web APIs and used as features of a log-linear model to predict surge multipliers in Pittsburgh. L_1 regularization is used to allow a small number of important features to be selected in a data-driven way from the large number of spatio-temporal features describing the urban state. To allow the linear model to describe non-linear behavior, temporal segmentation and clustering are employed. Days are segmented into time-of-day windows and separate models are trained for each. Clustering extracts temporal patterns from the data and distance to the centroids of these clusters are included as features in the model.

Overall, log-linear regression with L_1 regularization is able to predict Uber surge multipliers up to 2 hours in advance with smaller mean squared error than naïve and non-linear methods using only on average 6, and at most 25, of the 8898 measurements from the current and recent past. In each of the time windows the model is able to outperform naïve methods when predicting surge multipliers up to 120 minutes in advance. The fact that it is a linear model, of course, lends straight-forward interpretations to the fitted parameters. In addition to offering greater insight into surge multipliers, this model, without any modification, could be deployed to produce real-time predictions of Uber surge multipliers. Each of the features used in this model is available in real-time from various web APIs.

A linear model with L_1 regularization, commonly known as “LASSO” is indeed a classical methodology, but one that, in combination with feature engineering and unsupervised learning, out-performs more complex methods in many applications. In this case, LASSO was the parsimonious choice. Moreover, more complex learning methods lend themselves less easily to analysis of their features than LASSO does. Taken together, LASSO enabled us to achieve better and more interpretable results than the more complex methods we applied.

We find that the model performs best to predict Uber surge multipliers in urban ar-

eas. For both Uber and Lyft the model tends to under-perform naïve methods for surge multipliers between 1.1 and 1.5 and out-perform naïve methods for larger values. This effect is particularly pronounced in the Lyft model, whose poor performance on low surge multipliers substantially degrades its overall performance. Lyft’s coarser surge discretization is at least partially to blame, but other operational differences between the two organizations likely cause most of the discrepancy.

The model consistently places weight on the variables representing the current and recent surge multipliers; this makes sense given the strong auto-correlation observed in the surge multipliers. Outside of surge multiplier variables, feature selection varies by location consistent with high-level neighborhood characteristics. The spatial distribution of feature weights also varies by location but tends to be geographically disparate, highlighting the importance of including non-local spatial data in the model. The model sometimes selects seemingly counter-intuitive features, for example, the traffic speed in a distant location, or the number of events in a relatively un-eventful tract. However, since surge multipliers are a proxy for unanticipated demand, the occurrence of an event that is not widely known might be more likely to influence a surge multiplier than a well-known event.

If this model were employed as a decision tool, awareness of the predictions would affect driver and rider behavior, which would affect observed surge multipliers. However, real-time prediction is critical to ride-sourcing system operation and management. Many of the features used in this chapter do not meet the stationarity assumptions required by linear regression. It is for this reason that we train the model on a fixed period (approximately 5 months) and evaluate the fitted model on the subsequent week. Behavioral changes similarly violate stationarity assumptions by changing the relationship between the features and surge multipliers over time. Using a moving window of training data from the recent past solves this problem as well. In short, if such a model were employed in a real-time application, current data would be incorporated and past data dropped on a weekly basis, slowly capturing changing behavior over time. We expect the model would remain effective even under changing surge multiplier dynamics caused by the dissemination of its predictions.

There are three broad classes of extensions to this model. First, a multi-task or mixture model approach could be applied to this modeling framework to both improve general insights into the behavior of surge multipliers at a city scale and to improve the generalization error through additional regularization. Second, the feature set could be extended to include more real-time measurements of the urban environment, including public transit data from GTFS, the wait time from Uber or Lyft’s API, or more specific attributes of location, for example the number of open businesses. Third, the same feature set used in this model might be used to predict other characteristics of urban mobility, such as public transit availability.

4

On the impact of optimized fleets in transportation networks

4.1 Introduction

The central claim of this chapter is that groups of vehicles coordinated to achieve a group-level objective have significant and diverging impacts on total travel cost on the road network. On the one hand, vehicle coordination may increase total travel cost on the network. On the other hand, if deployed on certain origin-destination pairs in sufficient volume, a fleet can reduce system cost on the network, up to achieving the system optimal (SO) network flow. In this work, we consider the static network effect of a single fleet that routes its vehicles in a fleet optimal (FO) manner to minimize the total fleet travel cost in mixed traffic with individual vehicles who seek to minimize their own travel cost.

Technological advances have allowed mobility and information service providers to influence traveler decision-making in real-time and on an unprecedented scale. ride-sourcing vehicles, connected and autonomous vehicles, vehicles using real-time navigation devices, and vehicles using carpool matching services, are all examples of what we call *ad hoc* fleets currently operating on road networks. We refer to the platforms who coordinate *ad hoc* fleets interchangeably as ‘service providers’ or ‘fleet coordinators’.

It is often in the interest of the service to coordinate the vehicles in their fleet to advance a service-level goal which may not always align with the goal of an individual user. In this work we examine the goal of minimizing average fleet travel cost or, equivalently, travel time. This is perhaps the simplest fleet-level goal lying in the intersection of plausible and interesting fleet behavioral principles; it is by no means the only one.

Ride-sourcing platforms like Uber and Lyft coordinate drivers in a variety of ways to benefit the platform. Matching riders and drivers on the Uber platform, for ex-

ample, is optimized in batches over a local fleet (Uber Technologies 2022). With the same information an individual driver could likely find a better match than the one they were assigned precisely because the matching was done to *minimize* a fleet level metric rather than *equalize* a driver level metric. Minimizing travel time in particular is an important existing goal of strategy at Uber. For a fixed demand and fixed driver pool, quicker service means increased capacity and higher level of service, which in turn can demonstrate the value of the platform to new riders and drivers. The dynamic pricing strategy known as “surge-pricing” is designed specifically to reduce rider wait time when it increases too quickly (Gurley 2014). In effect, surge-pricing acts as an incentive to compensate drivers for relocating (without a fare) in order to reduce the total wait time of the riders. In another example, Uber employs “pickup spots” to reduce congestion *among Uber vehicles* in small areas of high demand. Many airports, for example, have dedicated entire wings of near-terminal parking structures to coordinate ride-hailing pickups. This is fundamentally a fleet (and likely system) optimal rather than user optimal solution: for each individual rider-driver pair terminal curbside pickup would offer travel-time savings, but measured across the fleet, organized parking lot pickups offer less total time wasted.

Real-time navigation systems with substantial user pools like Google Maps also have the means and motivation to induce FO behavior. Conventional reasoning among network equilibrium theorists is these systems will lead to a user equilibrium (UE) by providing travelers with accurate day-to-day traffic information to find their least cost route. However, the goal of Google Maps is to generate revenue. It will only help users find least cost paths insofar as it drives ad revenue: by keeping users on the platform. In this light, UE behavior is but one of many possible results of Google Maps use; FO is another possibility with compelling rationale. Google Maps currently balances several factors in selecting routes including travel time, future predicted traffic, emissions (i.e. “green routes”) (Alcántara 2021), road quality, directness, and safety, among others (Lau 2020). The fact that Google Maps is predicting traffic while simultaneously directing a subset of it means that it has, intentionally or not, answered the question: how will our traffic predictions take into account the directions provided? If the predictions do not take into account the directions at all, then the predictions will be inaccurate when a large volume of travelers use the service. If they do, then the two should be mutually consistent: the traffic predictions used to generate the directions remain accurate when the traffic volume of users following those directions is fully incorporated into traffic prediction. There are many models which achieve mutual consistency and FO is a particularly compelling one. Under FO, because the fleet rather than the individual cost is minimized the directions are not always fair: some users end up on better routes than others, but its users are, on average, better off than they otherwise would have been. However, Google Maps has a great deal of flexibility in how it presents routes to users so that it may not be obvious to a user, and

is in fact very difficult to validate that their route is not best. FO in this setting offers a competitive advantage to real-time navigation providers with large user pools: in the same way that large firms negotiate lower unit prices, these services can extract lower unit travel costs via coordination. For users then, Google Maps may be able to offer them lower travel costs on average than individuals and competitors with smaller market share, thereby keeping users on its platform.

In short, fleet coordinators may, and in many cases already do, influence the behavior of their fleets in order to improve a metric computed over the fleet for the benefit of the service. In this work we take this metric to be the total travel cost of the fleet, and present the system-level implications of its use. This chapter illustrates and highlights the importance of considering fleet's goals in system-level planning and operation, and the methodology and solutions can be extended to incorporate other system-level metrics in future work.

If *ad hoc* fleets are to remain a fixture on road networks, how should transportation planners understand and anticipate their use of transportation infrastructure? For example, in analogy to oligopoly models in economics, if a service provider wished to leverage its market power to extract a better deal on travel cost from the network via route choice coordination of its fleet, what would happen to network efficiency? Further, are there network designs which align the interests of the fleet and society at large? By that same token, fleet coordinators also have an interest in understanding the network impacts of coordinating their fleet: what discount on travel cost does their market power allow them to extract? Perhaps there is an opportunity for service providers to align their goals with traffic managers, or pass congestion-relief incentives to riders and drivers. Both the identification of such opportunities and the measurement of their benefits rely on a framework for understanding how vehicles with fleet-level goals and those with individual goals interact on a network.

The remainder of this chapter is organized as follows. Section 4.2 introduces the notation and fundamental concepts used throughout this chapter. Section 4.3 presents an example of the “fleet optimality paradox” in which the presence of a fleet increases total system travel time relative to UE. It is followed by examination of conditions under which fleets *do* improve total system travel time. In Section 4.4, we introduce two important mixed equilibria: the smallest fleet to induce SO, termed the critical fleet size for SO (CFS-SO), and the largest fleet to induce UE, termed the critical fleet size for UE (CFS-UE). CFS-SO and CFS-UE are examined analytically in a parallel network. Solution methodologies are then developed to solve both CFS-SO and CFS-UE in general networks. Section 4.5 presents the critical fleet size solution on two networks and provides an analysis of the results. Section 4.6 discusses our findings, outlines potential areas of future research, and discusses the relevance of CFS-SO and CFS-UE as a practical tool for transportation planners, traffic managers, and fleet

coordinators to understand and tune the impact of fleets on road networks. Section 4.7 extends the CFS-SO and CFS-UE programs to accommodate prior knowledge of fleet demand patterns. Section 4.8 analyzes the socio-demographic impacts of critical fleet size. Finally, in Section 4.9 we offer closing remarks and areas for future extensions of our work.




Figure 4.1: Conceptual diagram of system cost at mixed equilibrium with one fleet as fleet penetration, as a percent of total demand, increases. In this chapter we will show fleets can worsen system cost relative to UE and that UE and SO may be achieved at certain penetration levels. Between these critical levels, however, it is unclear how total system cost behaves. Since demand is generally high-dimensional, this diagram does not directly map to general networks on which oftentimes CFS-UE is larger than CFS-SO.

This chapter aims to fill in some key details of the conceptual diagram in Figure 4.1 and extends the literature in the following ways;

1. We demonstrate that there exist networks on which mixed equilibrium with fleets will increase system cost over UE. This is marked on the conceptual diagram as “congestion upper bound?”. We leave open the question of whether there exists an upper bound, but if there is one it must be, at least for some networks, strictly greater than UE.
2. We show that when there is a single fleet, there must exist a fleet demand pattern for which the resultant mixed equilibrium is SO; in this work we are concerned with the smallest fleet to induce SO (“CFS-SO” in Figure 4.1). We also show that there must exist a fleet demand pattern for which the resultant mixed equilibrium is UE; we are interested here in the largest fleet to induce UE (“CFS-UE” in Figure 4.1).
3. We formulate novel mathematical programs to find CFS-SO and CFS-UE. Exact

and heuristic algorithms are developed to efficiently solve those programs in large-scale transportation networks.

4. We solve for CFS-SO and CFS-UE on two real-world networks, finding that not all vehicles need to participate in an optimized fleet for the system to attain its minimum system cost. Moreover we show that such a fleet would need no external subsidy to benefit from optimizing itself.
5. We demonstrate the practical relevance of critical fleet size to regulators, transportation planners, and fleet coordinators.

4.2 Preliminaries

In this section we define our notion of network equilibrium which we term ‘mixed equilibrium with fleets’ and show that this mixed equilibrium has UE and SO as special cases.

We consider a road network represented as a graph \mathcal{G} with nodes N and edges (or links) A . On this network, there are a volume of travelers each seeking to travel from one node to another. We refer to the set of all such ordered node pairs as the set of origin-destination (O-D) pairs, $W \subseteq N \times N$. The travel demand is segmented into *flow classes*: the individual user flow class, denoted (typically as a superscript) by u , and $k \geq 1$ fleet flow classes, denoted c_i for $i = 1, \dots, k$. Although in this chapter we will consider only one fleet, we define mixed equilibrium with multiple fleets here for completeness. The set of all fleet flow classes is written \mathcal{F} . The travel demand for each flow class is represented as a vector of O-D travel volume: $\mathbf{q}^u \in \mathbb{R}_+^{|W|}$ for the user flow class and $\mathbf{q}^{c_i} \in \mathbb{R}_+^{|W|}$ for each fleet flow class $c_i \in \mathcal{F}$. For each flow class and O-D pair, travel is represented by the assignment of travel demand across the paths connecting the O-D pair. The set of all paths on the network is denoted by P . This assignment may be represented using a trip path incidence matrix, $\mathbf{M} \in \{0, 1\}^{|W| \times |P|}$, where the element at (w, p) is 1 if and only if path p starts and ends at the origin and destination, respectively, of the pair w . A user path assignment $\mathbf{f}^u \in \mathbb{R}_+^{|P|}$ is feasible if $\mathbf{M}\mathbf{f}^u = \mathbf{q}^u$. Similarly for each fleet $c_i \in \mathcal{F}$, the fleet path flow $\mathbf{f}^{c_i} \in \mathbb{R}_+^{|P|}$ is feasible if $\mathbf{M}\mathbf{f}^{c_i} = \mathbf{q}^{c_i}$. The relation between path flow and link flow is represented using a link path incidence matrix, $\mathbf{D} \in \{0, 1\}^{|A| \times |P|}$, where the element at (a, p) is 1 if and only if link a lies on path p . A user link assignment $\mathbf{x}^u \in \mathbb{R}_+^{|A|}$ is feasible if there exists a feasible user path flow \mathbf{f}^u such that $\mathbf{D}\mathbf{f}^u = \mathbf{x}^u$. Similarly for each fleet $c_i \in \mathcal{F}$, the fleet link flow $\mathbf{x}^{c_i} \in \mathbb{R}_+^{|A|}$ is feasible if there exists a feasible fleet path flow $\mathbf{f}^{c_i} \in \mathbb{R}_+^{|P|}$ such that $\mathbf{D}\mathbf{f}^{c_i} = \mathbf{x}^{c_i}$. Note that feasibility is always with respect to an O-D demand vector.

Travelers incur a non-negative travel cost on each link traversed represented as a link-separable monotone non-decreasing and differentiable function of aggregate link flow $\mathbf{t} : \mathbb{R}_+^{|A|} \rightarrow \mathbb{R}_+^{|A|}$.

We assume that individual users each wish to minimize the cost of their own travel, corresponding to Wardrop's first principle (Wardrop 1952). Each fleet, as the Cournot-Nash players in Harker (1988), is assumed to minimize the average travel cost over the fleet. We may now define mixed equilibrium with fleets.

Definition 1 (Mixed Equilibrium with Fleets). *Let Ω^u denote the set of q^u -feasible user link flows and Ω^{c_i} for each $c_i \in \mathcal{F}$ the set of q^{c_i} -feasible link flows for fleet c_i . The tuple of feasible link flows $(\mathbf{x}^{u*}, \mathbf{x}^{c_1*}, \dots, \mathbf{x}^{c_k*})$ is a mixed equilibrium if the following holds:*

$$\langle \mathbf{t}(\mathbf{x}^*), \mathbf{x}^u - \mathbf{x}^{u*} \rangle \geq 0 \quad \forall \mathbf{x}^u \in \Omega^u \quad (4.1)$$

$$\langle \tilde{\mathbf{t}}(\mathbf{x}^*, \mathbf{x}^{c_i*}), \mathbf{x}^{c_i} - \mathbf{x}^{c_i*} \rangle \geq 0 \quad \forall \mathbf{x}^{c_i} \in \Omega^{c_i} \quad \forall c_i \in \mathcal{F} \quad (4.2)$$

where $\mathbf{x}^* = \mathbf{x}^{u*} + \sum_{c_i \in \mathcal{F}} \mathbf{x}^{c_i*}$ represents the aggregate link flow at the mixed equilibrium and $\tilde{\mathbf{t}}(\mathbf{x}, \mathbf{x}^{c_i}) = \mathbf{t}(\mathbf{x}) + \mathbf{x}^{c_i} \mathbf{t}'(\mathbf{x})$ represents the marginal cost of fleet travel (referred to as fleet marginal cost), and $\mathbf{t}'(\mathbf{x})$ is the element-wise derivative of the link cost function.

It can easily be seen via the Beckmann transformation (Sheffi 1985; Beckmann, McGuire, and Winsten 1956) that an equilibrium of the fleet marginal cost in (4.2) is equivalent to a minimization over the total fleet cost, $\mathbf{x}^{c_i} \mathbf{t}(\mathbf{x})$. It is also useful to point out that the difference between our fleet users and the system optimal users in H. Yang, X. Zhang, and Meng (2007), K. Zhang and Y. M. Nie (2018), Z. Chen et al. (2020), and Sharon et al. (2018) is that the system optimal users seek to equalize the *system marginal cost*, expressed as $\mathbf{t}(\mathbf{x}) + \mathbf{x} \mathbf{t}'(\mathbf{x})$, or in our notation, $\tilde{\mathbf{t}}(\mathbf{x}, \mathbf{x})$.

From Definition 1 we immediately see that SO and UE are special cases of mixed equilibrium with fleets. In particular, if all fleet demand is zero, then mixed equilibrium is UE and if one fleet accounts for all of the demand, then mixed equilibrium is SO.

We are interested in demonstrating that on some networks when neither individual users nor the fleet control all of the demand, SO can be achieved in the aggregate flow. It is straightforward to imagine how a fleet might *reduce* congestion on the network: the fleet minimizes the total cost for a subset of the flow so one would hope that this effort also reduces travel cost for non-fleet users. We can generate an intuition for why we should expect a mix of selfish behaviors to ever induce SO by considering H. Yang, X. Zhang, and Meng (2007) in a simpler setting: adding SO users to aggregate UE on the Braess network. The Braess network (Sheffi 1985) contains three paths, "upper", "lower", and "shortcut". At UE, 2 units of flow use each path, at SO, 3 units will take each of the upper and lower paths, with no flow on the shortcut. If we were to replace some UE demand with SO users, they would simply replace UE flow on the upper and lower paths; all remaining flow would still utilize the shortcut path. So despite the fact that the SO users are choosing routes to benefit the system (not themselves), they still choose paths they would have chosen as users. At small SO penetration these two notions of cost are aligned: the least marginal cost paths are

also least cost paths.

Now suppose we are at aggregate SO and we replace some SO flow with a fleet. The total marginal cost and fleet marginal cost differ by the product of the user flow and the link cost derivatives. If this difference is uniform enough across paths at aggregate SO, then the two notions of cost may be aligned: least fleet marginal cost paths are also least total marginal cost paths. We would then expect a fleet to make the same route choices that they would have made as SO flow and therefore the aggregate network state will not change when SO flow is replaced with fleet flow. If we can do the same with UE users, finding paths where cost and marginal cost align and replace SO users with UE users, then UE users would similarly make the same decisions SO users would. In effect we have partitioned the network at SO into paths for which cost and marginal cost are aligned and paths for which fleet marginal cost and marginal cost are aligned so that the combination of FO and UE behaviors in aggregate achieves SO on the network.

4.3 Fleet optimality paradox

In this section we first analyze total system cost at mixed equilibrium in general. We then demonstrate a concrete example of an optimized fleet which, at mixed equilibrium, *increases* total system cost relative to UE. We refer to this phenomenon as the fleet optimality paradox.

4.3.1 Total delay under mixed equilibrium with fleets

On the one hand, total delay at mixed equilibrium with fleets has a trivial but tight lower bound of the total delay at SO. Only under certain conditions, however, is the total travel time at UE an upper bound of mixed equilibrium with fleets. We give a sufficient condition in (4.3) the proof of which is provided in Appendix A.1.

$$\sum_{a \in A} \mathbf{x}_a t'_a(\mathbf{x}_a)(\mathbf{y}_a - \mathbf{x}_a) \geq \sum_{a \in A} \sum_{i \in \mathcal{F}} \mathbf{x}_a^{c_i} t'_a(\mathbf{x}_a)(\mathbf{y}_a^{c_i} - \mathbf{x}_a^{c_i}) \quad (4.3)$$

where, $\mathbf{x}_a = \mathbf{x}_a^u + \sum_{i \in \mathcal{F}} \mathbf{x}_a^{c_i}$ and $\mathbf{y}_a = \mathbf{y}_a^u + \sum_{i \in \mathcal{F}} \mathbf{y}_a^{c_i}$ are the link flows at mixed and user equilibrium respectively. This condition is not of practical use since it depends on the link flows at UE and mixed equilibrium: one could simply compare the total system costs. It does however yield some helpful theoretical insight.

The left hand side of (4.3) is a first order estimate of the change in total cost when the aggregate link flow shifts from mixed equilibrium to UE; the right hand side is the first order estimate of the change in total *fleet* cost when the network shifts *fleet* link flow from mixed equilibrium to UE. The coefficients, $\mathbf{x}_a t'_a(\mathbf{x}_a)$ and $\mathbf{x}_a^{c_i} t'_a(\mathbf{x}_a)$, measure the impact of a given change in flow. Roughly speaking, (4.3) is fulfilled if links with

high total impact that receive additional aggregate flow are also high *fleet* impact and receive additional fleet flow. When UE behavior shifts fleet flow onto high fleet impact links which are not also high total impact links, then the right hand side can exceed the left, possibly by enough so that the total system cost at mixed equilibrium exceeds that at UE. This idea will be exploited in the following example.

4.3.2 An example of the fleet optimality paradox

In this section we will examine the network in Figure 4.2 as a concrete example of the fleet optimality paradox. Intuitively, the paradox arises when the fleet chooses a path with high system marginal cost but low *fleet* marginal cost. In the example network this is caused by an imbalance of fleet flow between two alternative paths. The fleet only wishes to avoid interfering with other vehicles in the fleet and so does not consider the effect of its route choice on the individual users. As a result, the path with fewer fleet vehicles has low *fleet* marginal cost even though its many users induce a large system marginal cost, causing the total system cost to increase over UE.




Figure 4.2: Fleet optimality paradox: an illustrative toy network.

Consider the network in Figure 4.2. The O-D pairs are $a \rightarrow b$, $c \rightarrow d$, and, $e \rightarrow f$. The only real route choice in this network exists for users travelling from a to b . The other O-D pairs have only one route available to them. The links leaving a and entering b have 0 cost to simplify the arithmetic, but this choice neither fundamentally changes the problem nor the paradox.

We define the link costs t for cd and ef as follows:

$$t_{cd}(x_{cd}) = 1 \cdot x_{cd} + 50 \quad (4.4)$$

$$t_{ef}(x_{ef}) = 10 \cdot x_{ef} + 0 \quad (4.5)$$

We consider the mixed equilibria resulting from the two O-D demand scenarios shown in Table 4.1. The left hand column under each O-D pair in Table 4.1 present the O-D demand for first scenario (UE); the right-hand columns (in gray) show the

O-D demand for the mixed equilibrium which produces the paradox, the fleet optimal scenario (FO). Note that the total volume of vehicles between each O-D pair, the aggregate demand, remains unchanged.

Table 4.1: Fleet, user, and, aggregate demand for the fleet optimality paradox

	$c \rightarrow d$		$e \rightarrow f$		$a \rightarrow b$	
fleet	0	1	0	2.75	0	0.05
user	13	12	2.75	0	1	0.95
aggregate	13		2.75		1	

We summarize the UE and FO equilibria in Table 4.2.

Table 4.2: Demonstration of the fleet optimality paradox

path via	UE			FO							
	aggregate (users)			aggregate			user		fleet		
	flow	cost	total cost	flow	cost	total cost	flow	total cost	flow	total cost	
$c \rightarrow d$	13.00	63.0	819.000	13.05	63.05	822.8025	12.00	756.60	1.05	66.2025	
$e \rightarrow f$	3.75	37.5	140.625	3.70	37.00	136.9000	0.95	35.15	2.75	101.7500	
total	16.75		959.625	16.75		959.7025	12.95	791.75	3.80	167.9525	

At the UE link flow, given by the left half of Table 4.2, the link cost of ef is minimal: all demand from a to b selects the route through link ef . At the FO link flow, given by the right half of Table 4.2, the users and fleet traveling from a to b prefer different paths. The individual users still prefer the path through ef as it remains least cost. The fleet, however, prefers the path through cd because it is least *fleet marginal cost*. By selecting the least fleet marginal cost path, the fleet *reduces* its total cost relative to its total cost under UE as shown in Table 4.3 (bottom row, bold), however, the total system cost increases from UE to FO (Table 4.2 bottom row, in bold). The fleet operator makes routing choices to minimize the travel delay impact only on members of its own fleet, regardless of the impact on other travelers. In this case, the fleet has so few vehicles on link cd that the fleet marginal cost of the path through cd is lower than the fleet marginal cost of the path through ef , where the fleet already has vehicles.

To see why this occurred we note that not only is the path through ef user optimal, it is also system optimal. Which of course means that the SO link flow and UE link flow are the same for the given aggregate travel demand on this network. What is then also true of this example is that the fleet is better off then it would have been at SO, this is related to a more general phenomenon we call the system optimal paradox discussed further in Section 4.3.3.

Table 4.3: Demonstration that the fleet decreases its total cost at FO

		UE		FO	
OD pair	fleet demand	OD cost	total cost	OD cost	total cost
$c \rightarrow d$	1.00	63.0	63.000	63.05	63.0500
$e \rightarrow f$	2.75	37.5	103.125	37.00	101.7500
$a \rightarrow b$	0.05	37.5	1.875	63.05	3.1525
total	3.80		168.000		167.9525

4.3.3 System optimality paradox

The fact that the fleet may actually reduce its total travel cost relative to SO in the fleet optimality paradox results from a property of SO traffic assignment which we call the system optimality paradox. For a network with non-unit price of anarchy, the following is true of SO traffic flow:

1. there is *always* volume that is *relatively* worse off compared to available paths, and
2. there is *possibly* volume that is *absolutely* worse off compared to the travel cost of the same O-D pair at UE.

In short, when SO path flow is partitioned, the guarantees that SO provides for the network as a whole may not apply to each partition individually.

The first point follows directly from the UE principle which requires that all used paths between each O-D pair have equal and minimal travel cost. Any traffic assignment that is not UE must therefore have at least one O-D pair for which some volume occupies a path that is not least cost: this volume is worse off *relative* to some other available path.

The second point occurs only on some networks. On the Braess paradox network, for example, the cost of every used path at SO is strictly less than the cost of every used path at UE. Thus, even though some users at SO are *relatively* worse off (there is a lesser cost path available), they are all better off than they would have been under UE. However, we can easily construct a network where this is not the case. A concrete example is given in Appendix A.2, but the essential intuition is that when UE is different from SO on a parallel network one path gains volume (and therefore travel cost) from the other in moving from UE to SO. In this setting, not only are some users not using the least cost path (they are relatively worse off) they are also experiencing higher absolute travel costs than they would have under UE (they are absolutely worse off). This solution is SO because the users that are better off (both relatively and absolutely) accrue total travel cost savings that exceed the increased travel cost of the worse off travelers so that on average the system is better off.

4.3.4 Concluding remarks

Taken together, this section shows that at extreme penetration levels of 0% and 100% mixed equilibrium with fleets achieves a total system cost equal to the total system cost at UE and SO respectively. However, as penetration levels increase from 0% to 100%, total system cost will not necessarily monotonically decrease, and may in fact become larger than the total system cost under UE. This has strong policy implications: under certain demand and roadway network conditions, increasing fleet penetration (e.g. individuals shift to use a mobility service that is centrally planned by a private entity) may increase congestion. Only if the fleet penetration is adequately high would its presence be guaranteed to reduce network congestion, up to achieving SO. Transportation network companies, such as Uber/Lyft, who decide to optimize their fleet may paradoxically lead to more congestion in some areas in their initial development stage (when the penetration is low), even if they would have replaced some private driving trips and O-D demand remains the same.

4.4 Critical fleet size

In this section we will examine how, in certain networks, SO can be achieved through mixed equilibrium where not all vehicles are fleet vehicles. The smallest fleet size for which mixed equilibrium achieves aggregate SO traffic assignment is expressed as a fraction of total demand and termed the **critical fleet size for system optimum** (CFS-SO). There is a complementary notion, termed the **critical fleet size for user equilibrium** (CFS-UE), which measures the largest fleet size for which mixed equilibrium achieves aggregate UE traffic assignment.

CFS-SO is bounded by 0 and 1; a value of 1 indicates that all vehicles *must* be in the fleet in order for mixed equilibrium to result in SO and a value of 0 indicates that no fleet vehicles are required. A network has a CFS-SO of 0 if and only if its UE and SO assignments are the same, in other words, the network has a price of anarchy (Roughgarden 2005) of 1. Similarly, CFS-UE is bounded by 0 and 1; a value of 0 indicates that the presence of any fleet vehicles on the network “breaks” UE (induces a mixed equilibrium whose aggregate link flow is different from UE). A CFS-UE of 1 indicates that all vehicles may participate in the fleet without changing UE. A CFS-UE of 1 implies that UE and SO assignments are identical on this network and as a result, CFS-UE is 1 if and only if CFS-SO is 0. Outside of this case, the relationship between CFS-SO and CFS-UE on a network is not at all clear and is left for future research.

CFS-UE is different from the notion of the smallest fleet for which mixed equilibrium is different from UE. As a somewhat counter-intuitive result with respect to the simplified one-dimensional Figure 4.1, CFS-UE can be larger than CFS-SO. In representing demand as a uni-dimensional quantity, Figure 4.1 is in a sense overly simplis-

tic: when demand is higher dimensional, CFS-UE can occupy an entirely different set of O-D pairs than CFS-SO and simply have larger magnitude. In one dimension, this is only possible when, as we discuss in the next paragraph, the price of anarchy on the network is 1. The smallest fleet to ‘break’ UE may also be of interest, but is strictly a different question than the one CFS-UE seeks to answer.

The example network in Figure 4.2 on which we demonstrated the fleet optimality paradox also provides instructive examples of CFS-SO and CFS-UE. We noted that the UE and SO traffic assignments are the same on the example network. As a result, CFS-SO is 0 because users alone achieve SO, and CFS-UE is 1 because when all of the volume on the network is an optimized fleet, the network is at SO, which is also UE. The paradox presented in Section 4.3 is thus also an example of the following three phenomena:

1. CFS-UE can be greater than CFS-SO,
2. fleet O-D demand that is element-wise greater than the fleet O-D demand at CFS-SO is **not** necessarily SO, and,
3. fleet O-D demand that is element-wise less than the fleet O-D demand at CFS-UE is **not** necessarily UE.

What CFS-SO and CFS-UE *do* imply however, is that at mixed equilibrium with fleets,

1. SO cannot be achieved with a fleet smaller than CFS-SO, and
2. UE cannot be achieved with a fleet larger than CFS-UE.

4.4.1 Critical fleet size on a parallel network

In this section we will analyze CFS-SO and CFS-UE on a parallel network. We derive analytical results that hold for all separable, monotonic increasing link performance functions on a parallel network. Though we should not expect these results to generalize, the parallel network is useful in developing an intuition for critical fleet size and for mixed equilibrium more generally. In particular, we are interested in whether or not critical fleet size results in paths, O-D pairs, or entire networks that are exclusively fleet vehicles or individual users. In what follows, consider a parallel network with one O-D pair connected by n parallel links.

Mixed equilibrium at aggregate UE

Let $\bar{x}^{ue} \in \mathbb{R}_+^n$ denote the UE link flow on the network. We wish to find the mixed equilibrium with the largest fleet share such that UE link flow is preserved in aggregate. Proof of Propositions 1 to 4 can be found in Appendix A.3.

Proposition 1. *Fleet flow at any mixed equilibrium where fleet demand is strictly*

positive on a link is positive if and only if aggregate link flow is positive. That is,

$$\bar{x}_a^{\text{ue}} > 0 \iff x_a^c > 0 \quad \forall a \in A \quad (4.6)$$

Proposition 2. Given a fleet demand $q^c \in \mathbb{R}_+$ such that $0 < q^c \leq q$ we can find the mixed equilibrium on the parallel network analytically;

$$\tilde{t}^* = \frac{q^c}{\sum_{a \in A_+} \frac{1}{t'_a}} + t^* \quad (4.7)$$

$$x_a^c = \frac{\tilde{t}^* - t^*}{t'_a} \quad (4.8)$$

provided that the following holds,

$$\tilde{t}^* \leq \min_{a \in A} t_a + \bar{x}_a^{\text{ue}} t'_a \quad (4.9)$$

where A_+ denotes the set of all links where aggregate flow is strictly positive, t_a is the link cost on link $a \in A$ evaluated at mixed equilibrium, t'_a is the derivative of the link cost on link $a \in A$ evaluated at mixed equilibrium, \tilde{t}^* is the minimum fleet marginal link cost at mixed equilibrium, and t^* is the minimum link cost at mixed equilibrium.

Proposition 3. The largest fleet demand that preserves UE is given by,

$$\sum_{a \in A_+} \frac{\tilde{t}^* - t^*}{t'_a} = \sum_{a \in A_+} x_a^c = q^c \quad (4.10)$$

where \tilde{t}^* satisfies

$$\tilde{t}^* = \min_{a \in A} t_a + \bar{x}_a^{\text{ue}} t'_a \quad (4.11)$$

Mixed equilibrium at aggregate SO

We now wish to find the smallest fleet demand such that SO is preserved at mixed equilibrium.

Proposition 4. The minimum fleet demand required to induce SO flow on a parallel network with aggregate demand q is either

1. 0 if the UE flow is the same as the SO flow; or
2. q if UE flow is different from SO flow (CFS-SO=1).

In short, Proposition 4 ensures that for any parallel network CFS-SO is either 0 or 1.

Extending to a general network

While the parallel network provides a useful demonstration of fleet optimal mixed equilibrium, Proposition 4 ensures that it will not be an interesting one. In general networks, any value of CFS-SO is possible. Perhaps the simplest, albeit unsatisfying, way to demonstrate this fact is to consider a network composed of two parallel networks one of which is entirely fleet and the other, entirely individual users. Any value of critical fleet size can be achieved by varying the demand on the two sub-networks.

Nevertheless, the parallel network does provide us some valuable intuition for CFS on general networks. Key to the proof of Proposition 4 is the realization that if the set of least cost and least marginal cost paths are not the same, the fleet will need to fill those paths that are least marginal cost, but not least cost. However, in order to maintain fleet optimal assignment, the fleet marginal cost must be equalized over all paths in use by the fleet, which in turn requires each path to be filled by the fleet. In a general network, since many path flows may induce the same aggregate link flow, it is not necessarily the case that user flow removed from a path must be replaced by fleet flow on *that specific path*, rather, the fleet may occupy a completely different set of (fleet-optimal) paths which still induces the same aggregate link flow. It is for this reason that we should not expect general networks, even at the O-D level, to only have the two extremal values of critical fleet size.

In total, it is possible, but difficult to imagine, that O-D pairs may not be exclusively user or exclusively fleet at critical fleet size. Precise conditions under which paths, O-D pairs, or entire networks are exclusively user or fleet are left for subsequent work. In this chapter, we will attempt to shed light on the question of exclusivity empirically on real world networks in Section 4.5.

4.4.2 Critical fleet size as a mathematical program with equilibrium constraints

The fleet vehicles we are examining wish to minimize the fleet travel time: $z(\mathbf{x}^c, \mathbf{x}^u) = \sum_a \mathbf{x}_a^c t_a(\mathbf{x}_a^u + \mathbf{x}_a^c)$. Given fleet O-D demand \mathbf{q}^c and the link flow of the individual users \mathbf{x}^u , the fleet path (link) flows may be computed as the solution to the following mathematical program (Sheffi 1985). The program is parameterized by the user link flow

and the fleet demand, which are considered constant within this program.

$$\text{FO}(\mathbf{x}^u, \mathbf{q}^c) = \arg \min_{\mathbf{f}^c} \mathbf{z}(\mathbf{x}^c, \mathbf{x}^u) \quad (4.12a)$$

$$\text{s.t.} \sum_{p \in P_w} \mathbf{f}_p^c = \mathbf{q}_w^c \quad \forall w \in W \quad (4.12b)$$

$$\mathbf{f}_p^c \geq 0 \quad \forall p \in P_w, w \in W \quad (4.12c)$$

$$\sum_{p \in P} \mathbf{f}_p^c D_{a,p} = \mathbf{x}_a^c \quad \forall a \in A \quad (4.12d)$$

We now consider the critical fleet size problem which imposes the external constraint that aggregate link flow at mixed equilibrium must match a given aggregate link flow. For CFS-SO, that aggregate link flow, denoted $\bar{\mathbf{x}}^{\text{so}}$, is link flow at SO. CFS-SO is expressed as the solution to mathematical program with equilibrium constraints (MPEC) given by Program 4.13 and CFS-UE is expressed as the solution to the Program 4.14, with $\bar{\mathbf{x}}^{\text{ue}}$ as the UE link flow.

$$\min_{\mathbf{f}^u} \sum_{w \in W} \mathbf{q}_w^c \quad (4.13a)$$

$$\text{s.t.} \sum_{p \in P_w^*} \mathbf{f}_p^u = \mathbf{q}_w^u \quad \forall w \in W \quad (4.13b)$$

$$\mathbf{f}_p^u \geq 0 \quad \forall p \in P \quad (4.13c)$$

$$\sum_{p \in P^*} \mathbf{f}_p^u \delta_{a,p} = \mathbf{x}_a^u \quad \forall a \in A \quad (4.13d)$$

$$\mathbf{x}^c = \mathbf{D} \text{FO}(\mathbf{x}^u, \mathbf{q}^c) \quad (4.13e)$$

$$\mathbf{x}^c + \mathbf{x}^u = \bar{\mathbf{x}}^{\text{so}} \quad (4.13f)$$

$$\mathbf{q}^c + \mathbf{q}^u = \bar{\mathbf{q}} \quad (4.13g)$$

$$\max_{\mathbf{f}^u} \sum_{w \in W} \mathbf{q}_w^c \quad (4.14a)$$

$$\text{s.t.} \sum_{p \in P_w^*} \mathbf{f}_p^u = \mathbf{q}_w^u \quad \forall w \in W \quad (4.14b)$$

$$\mathbf{f}_p^u \geq 0 \quad \forall p \in P \quad (4.14c)$$

$$\sum_{p \in P^*} \mathbf{f}_p^u \delta_{a,p} = \mathbf{x}_a^u \quad \forall a \in A \quad (4.14d)$$

$$\mathbf{x}^c = \mathbf{D} \text{FO}(\mathbf{x}^u, \mathbf{q}^c) \quad (4.14e)$$

$$\mathbf{x}^c + \mathbf{x}^u = \bar{\mathbf{x}}^{\text{ue}} \quad (4.14f)$$

$$\mathbf{q}^c + \mathbf{q}^u = \bar{\mathbf{q}} \quad (4.14g)$$

Because the aggregate link flow is fixed under UE or SO and the link performance function depends only on the aggregate link flow, the travel cost (and marginal travel cost) for each link is a constant with respect to both the user and fleet link flows. As a result, the set of least cost paths between each O-D pair, P^* , is fixed and known *a priori*. Constraints 4.13d and 4.14d therefore are sufficient to describe the UE condition in their respective programs completely: individual users may only use least cost paths.

4.4.3 Critical fleet size heuristic solution algorithm

Typically, MPECs are solved via a heuristic solution methodology known as sensitivity analysis (Tobin and Friesz 1988). However, for this specific MPEC, sensitivity analysis

is not the best choice. In this section we outline how one would apply sensitivity analysis before identifying its limitations in this setting. A detailed treatment of our sensitivity analysis approach is given in Appendix A.4.

Sensitivity analysis is typically applied to an MPEC for which the objective depends on the network at equilibrium and the network equilibrium, in turn, is affected by the decision variables. Sensitivity analysis provides the gradient of the equilibrium link flow with respect to the decision variables and is used to compute the descent (ascent) direction. In our Programs 4.13 and 4.14, the objective is an input of the equilibrium operator rather than its output. To apply sensitivity analysis we solve the programs in their augmented Lagrangian forms (Hestenes 1969), treating the equilibrium constraint as a penalty term. The algorithm, given by Algorithms 4 and 5, alternates between descent steps on the augmented lagrangian via sensitivity analysis and sequential linear programming (Wright, Nocedal, et al. 1999) known as the primal update, and ascent steps for the penalty weights, known as the dual updates.

The main theoretical deficit of solving CFS as an MPEC via sensitivity analysis is that the MPEC itself is non-convex and as such the solution algorithm is guaranteed only to return a local optimal solution. In contrast, our algorithm takes advantage of the problem structure to not only guarantee a global optimum solution but provide upper and lower bounds on the optimal value.

The main practical limitation is that this particular application of sensitivity analysis requires many more equilibrium computations than a typical application: once every iteration within each primal update which itself is simply one descent step. This can take significant amounts of time. Our methodology avoids computing mixed equilibria entirely and instead forms mathematical programs to descend directly in the space of feasible equilibria.

Taken together, on neither practical nor theoretical grounds does solving CFS via sensitivity analysis offer benefits over our methodology.

4.4.4 Exact critical fleet size solution algorithm via mixed integer programming

As an alternative to sensitivity analysis, we can leverage the structure of the problem to develop an exact solution methodology. Conceptually our approach reverses the order that sensitivity analysis imposes on the solution algorithm: instead of searching the space of mixed equilibria for those that happen to meet an aggregate link flow constraint, we search the space of feasible user-fleet link (path) flow partitions that satisfy the conditions of mixed equilibrium. In particular, the constraint that aggregate link flow for CFS-SO is fixed to be the aggregate SO link flow and the aggregate link flow for CFS-UE is fixed to be aggregate UE link flow implies two important facts:

1. The set of paths the fleet *could possibly use* is known, finite, and substantially

smaller than set set of all (simple) paths, and,

2. The fleet marginal link (path) cost is linear in fleet link (path) flow.

First, fact 2: when we consider the aggregate link flow fixed, the fleet marginal link cost is a linear function of fleet link flow; $\tilde{t}_a(\mathbf{x}_a^c) = \mathbf{t}_a + \mathbf{x}_a^c \cdot \mathbf{t}'_a$. If we consider the fleet marginal path cost as a vector and denote by $\text{diag}(\mathbf{t}')$ the $|A| \times |A|$ diagonal matrix with the elements \mathbf{t}'_a on the diagonal, then we may write the vector of fleet marginal path costs as a linear function of fleet flow \mathbf{f}^c : $\tilde{\mathbf{c}}(\mathbf{f}^c) = \mathbf{c} + \mathbf{D}^T \text{diag}(\mathbf{t}') \mathbf{D} \mathbf{f}^c$, or equivalently as a function of user flow \mathbf{f}^u , $\tilde{\mathbf{c}}(\mathbf{f}^u) = \mathbf{c} - \mathbf{D}^T \text{diag}(\mathbf{t}') \mathbf{D} \mathbf{f}^u$. Where \mathbf{c} and \mathbf{c} represent the path cost and marginal path cost at the given aggregate link flow.

From a computational perspective, the form of fleet marginal cost that depends on the user path flow is attractive because the user path flow vector will generally be sparse. By the equilibrium principle we know that individual users may only take least-cost paths so we need only specify user path flows for paths that are least cost at SO, a path set that in practice we expect to be known before solving CFS and that is substantially smaller than the set of all simple paths.

Next, fact 1. Because aggregate link flow is fixed, it is also true that the set of least marginal cost paths \tilde{P} , the set of usable paths for SO users, is fixed. When aggregate flow is SO any path flow vector that satisfies both the demand conservation constraint and produces the SO link flow must only use paths in \tilde{P} . As a result, at CFS-SO, the fleet may use only paths in \tilde{P} . For CFS-UE where the aggregate link flow at mixed equilibrium is the aggregate UE link flow, a similar argument applies: the fleet may only use least cost paths at UE, P^* .

Combining these two facts we can rewrite the CFS-SO and CFS-UE MPECs as mixed integer linear programs (MILP) in Programs 4.15 and 4.16 below.

$$\min_{\mathbf{f}^u \geq 0, \mathbf{f}^c \geq 0, \lambda, z} \sum_{p \in \tilde{P}} \mathbf{f}_p^c \quad (4.15a)$$

$$\text{s.t. } \tilde{\mathbf{c}}_r(\mathbf{f}^u) \geq \lambda_w \quad \forall r \in P_w \quad \forall w \in W \quad (4.15b)$$

$$\tilde{\mathbf{c}}_r(\mathbf{f}^u) \leq \lambda_w + m_1 \cdot (1 - z_r) \quad \forall r \in \tilde{P}_w \quad \forall w \in W \quad (4.15c)$$

$$\mathbf{f}_r^c \leq m_2 \cdot z_r \quad \forall r \in \tilde{P}_w \quad \forall w \in W \quad (4.15d)$$

$$\mathbf{D}^* \mathbf{f}^u + \tilde{\mathbf{D}} \mathbf{f}^c = \bar{\mathbf{x}}^{\text{so}} \quad (4.15e)$$

$$\mathbf{M}^* \mathbf{f}^u + \tilde{\mathbf{M}} \mathbf{f}^c = \bar{\mathbf{q}}^{\text{so}} \quad (4.15f)$$

$$z_k \in \{0, 1\} \quad \forall k \in \tilde{P}_w \quad \forall w \in W \quad (4.15g)$$

where m_1 and m_2 are large constants, \tilde{P} and P^* are the sets of least marginal cost and least cost paths at SO, respectively, $\tilde{\mathbf{D}}$ and \mathbf{D}^* are the link path incidence matrices restricted to paths in \tilde{P} and P^* , respectively, $\tilde{\mathbf{M}}$ and \mathbf{M}^* are the O-D path incidence matrices restricted to paths in \tilde{P} and P^* , respectively. The binary variable z_k encodes

whether the path k is used by the fleet (1) or not (0). The first two constraints ensure that the fleet flow is fleet optimal: the fleet marginal cost on all used paths are equal and minimal. We note that MPECs may always be written and solved as mixed integer programs; this is not often done in practice because the number of integer variables scales with the number of paths in the network and the cost function is typically non-linear. In our problem, the cost function is linear and the number of integer variables is substantially smaller than the number of paths.

The CFS-UE MILP is constructed analogously, notably as a maximization and using the set of least cost paths at UE where the CFS-SO program uses least marginal cost paths at SO.

$$\max_{\mathbf{f}^u \geq 0, \mathbf{f}^c \geq 0, \lambda, z} \sum_{p \in P^*} \mathbf{f}_p^c \quad (4.16a)$$

$$\text{s.t. } \tilde{\mathbf{c}}_r(\mathbf{f}^u) \geq \lambda_w \quad \forall r \in P_w \quad \forall w \in W \quad (4.16b)$$

$$\tilde{\mathbf{c}}_r(\mathbf{f}^u) \leq \lambda_w + m_1 \cdot (1 - z_r) \quad \forall r \in P_w^* \quad \forall w \in W \quad (4.16c)$$

$$\mathbf{f}_r^c \leq m_2 \cdot z_r \quad \forall r \in P_w^* \quad \forall w \in W \quad (4.16d)$$

$$\mathbf{D}^*(\mathbf{f}^u + \mathbf{f}^c) = \bar{\mathbf{x}}^{\text{so}} \quad (4.16e)$$

$$\mathbf{M}^*(\mathbf{f}^u + \mathbf{f}^c) = \bar{\mathbf{q}}^{\text{so}} \quad (4.16f)$$

$$z_k \in \{0, 1\} \quad \forall k \in P_w^* \quad \forall w \in W \quad (4.16g)$$

Mathematical programs for CFS-SO and CFS-UE with multiple fleets are discussed in Appendix A.7.

4.4.5 Critical fleet size bounds via linear programming

Although the MILP formulations of CFS-SO and CFS-UE offer substantial computational efficiencies relative to MILP formulations of general MPECs they can still present challenges on large networks. For such cases we present linear program relaxations of the CFS-SO and CFS-UE MILPs. The LP relaxation of the CFS-SO MILP is given by Program 4.17 and the LP relaxation of the CFS-UE MILP is given by Program 4.18.

$$\min_{\mathbf{f}^u \geq 0, \mathbf{f}^c \geq 0, \lambda} \sum_{p \in \tilde{P}} \mathbf{f}_p^c \quad (4.17a) \qquad \max_{\mathbf{f}^u \geq 0, \mathbf{f}^c \geq 0, \lambda} \sum_{p \in \tilde{P}} \mathbf{f}_p^c \quad (4.18a)$$

$$\text{s.t. } \tilde{\mathbf{c}}_r(\mathbf{f}^u) = \lambda_w \quad \forall r \in \tilde{P}_w \quad \forall w \in W \quad (4.17b) \qquad \text{s.t. } \tilde{\mathbf{c}}_r(\mathbf{f}^u) = \lambda_w \quad \forall r \in P_w^* \quad \forall w \in W \quad (4.18b)$$

$$\tilde{\mathbf{c}}_r(\mathbf{f}^u) \geq \lambda_w \quad \forall r \in P_w \quad \forall w \in W \quad (4.17c) \qquad \tilde{\mathbf{c}}_r(\mathbf{f}^u) \geq \lambda_w \quad \forall r \in P_w \quad \forall w \in W \quad (4.18c)$$

$$\mathbf{D}^* \mathbf{f}^u + \tilde{\mathbf{D}} \mathbf{f}^c = \bar{\mathbf{x}}^{\text{so}} \quad (4.17d) \qquad \mathbf{D}^*(\mathbf{f}^u + \mathbf{f}^c) = \bar{\mathbf{x}}^{\text{ue}} \quad (4.18d)$$

$$\mathbf{M}^* \mathbf{f}^u + \tilde{\mathbf{M}} \mathbf{f}^c = \bar{\mathbf{q}} \quad (4.17e) \qquad \mathbf{M}^*(\mathbf{f}^u + \mathbf{f}^c) = \bar{\mathbf{q}} \quad (4.18e)$$

Formulating the Critical Fleet Size in this way allows us to relate it directly to the Minimum Control Ratio (MCR) introduced by Z. Chen et al. (2020) and Sharon et al. (2018). In particular, because fleet vehicles must use a subset of the SO paths, CFS-SO is bounded below by the MCR, proof of which is given in Appendix A.5.

Proposition 5 (Bounding Critical Fleet Size via Linear Programs). *The heuristic linear programs for CFS-SO and CFS-UE provide the following bounds:*

1. *The linear program 4.17 provides an upper bound of CFS-SO.*
2. *The linear program 4.18 provides a lower bound of CFS-UE.*

The proof is given in Appendix A.6.

It is worth noting that the solution returned by the linear program will typically have paths that are unused by the fleet but remain least fleet marginal cost, known in the literature as a *degenerate* equilibrium. Degeneracy occurs in the LP because the set of paths over which fleet marginal cost must be equal and minimal is fixed. Allowing the fleet to abandon unused paths, thereby releasing them from the constraint that their fleet marginal cost is minimal, could allow for a smaller fleet. As a result, it may be possible to refine a bound returned by the LP by iteratively removing unused paths and re-solving the LP. Our preliminary work into such a procedure suggests that only modest gains can be achieved without a careful heuristic to select the paths to abandon each iteration. Broadly, the path abandonment procedure can be viewed as a member of the family of MILP heuristics. It is a very simple heuristic that always preserves feasibility by recognizing the structure of the problem but does not form an overall scheme that will find the global optimum. In contrast, MILP methods and heuristics available via state-of-the-art solvers do not recognize the structure of the problem but will progress toward the global optimum. A better option which we leave for future work would be to leverage path abandonment within the MILP solver.

4.4.6 Solution algorithm via column generation

In both the MILP formulation and their LP relaxations it is assumed that we have access to the full path set, in addition to the least cost and least marginal cost path sets at UE or SO. The latter two sets are not difficult to obtain (see Appendix A.8). However, the full set of simple paths may not be feasible to enumerate. With only a subset of the paths, we cannot guarantee that the fleet is using paths that are minimal fleet marginal cost. To circumvent this difficulty a column-generating approach is presented in Algorithm 1. Each iteration, the program is solved with an approximate path set \hat{P} , and least fleet marginal paths are computed at the candidate solution. If the least fleet marginal cost paths are not in the set of known paths, the solution assigned fleet flow to paths not in the fleet-usable path set. These paths are then added to the known path set and the updated linear program is re-solved. On re-solving the linear program, the fleet marginal costs of these paths are constrained to be no smaller than λ_w , the least fleet marginal cost over all paths connecting the OD pair.

Algorithm 1 Critical Fleet Size Solution Algorithm with Column Generation

```

1: procedure CriticalFleetSize( $\mathcal{G}, \bar{q}$ )
2:    $\bar{x}^{so} \leftarrow \text{SystemOptimalLinkFlow}(\mathcal{G}, \bar{q})$ 
3:    $P^*, \tilde{P} \leftarrow \text{LeastCostPaths}(\mathcal{G}, \bar{x}^{so}), \text{LeastMarginalCostPaths}(\mathcal{G}, \bar{x}^{so})$ 
4:    $\hat{P} \leftarrow P^* \cup \tilde{P}$  ▷ Initialize the known path set
5:   while true do
6:      $f^u, f^c, \lambda \leftarrow \text{Solve Program 4.15, 4.14, 4.17, or, 4.18 using } \hat{P} \text{ as the path set } P$ 
7:      $P_k \leftarrow \text{LeastFleetMarginalPaths}(\mathcal{G}, \bar{x}, x^c)$ 
8:     if  $P_k \cap \tilde{P} \neq \emptyset$  then ▷ A new path was discovered.
9:        $\hat{P} \leftarrow \hat{P} \cup (P_k \cap \tilde{P})$ 
10:    else
11:      return  $f^u, f^c$ 
12:    end if
13:  end while
14: end procedure

```

4.5 Experiments

Real-world networks present two practical challenges to Algorithm 1, using either the MIP or LP version. First, the programs scale linearly in the number of paths and O-D pairs. Even moderately sized real-world networks will produce mixed integer programs too large for state-of-the-art solvers to solve in a reasonable amount of time.

The MIP may simply be run with a fixed computational budget or their LP relaxations may be used instead.

A second problem facing real-world networks has to do with the accuracy of the SO link flow. In both the heuristic LPs and the MIPs, if we were to fix the user path flow to zero, the program should recover an SO path flow as the fleet path flow variable. However, when we have access only to an approximation of the SO link flow, this guarantee no longer holds. In fact, the opposite is true: because the approximate link flow is not exactly SO, there must be flow on paths which are not least marginal cost. As a result, when directly implemented with an approximate SO link flow Programs 4.17 and 4.15 may be infeasible.

To overcome this second challenge, we use a modified version of Program 4.17 given by Program 4.19 below. In this modified program, the set of least marginal cost paths (and least cost paths for individual users) is constructed as the set of ϵ -least (marginal) cost paths: a path is least (marginal) cost if its (marginal) cost is less than or equal to $(1 + \epsilon)$ times the (marginal) cost of the least (marginal) cost path for its O-D pair. Correspondingly, instead of treating the link flow recovery as a constraint, it is treated as a penalty term in the objective weighted by $\beta > 0$. Larger penalty weights represent preference for accurately approximating the target aggregate link flow. This same modification is applied to CFS-SO LP and MIP, CFS-UE LP and MIP, and to the MCR.

$$\min_{\mathbf{f}^u \geq 0, \mathbf{f}^c \geq 0, \lambda} \sum_{p \in P^*} \mathbf{f}_p^c + \beta \left\| \mathbf{D}^* \mathbf{f}^u + \tilde{\mathbf{D}} \mathbf{f}^c - \bar{\mathbf{x}}^{\text{so}} \right\|_2 \quad (4.19a)$$

$$\text{s.t. } \tilde{\mathbf{c}}_r(\mathbf{f}^u) \leq \lambda_w(1 + \epsilon) \quad \forall r \in \tilde{P}_w \quad \forall w \in W \quad (4.19b)$$

$$\tilde{\mathbf{c}}_r(\mathbf{f}^u) \geq \lambda_w \quad \forall r \in P_w \quad \forall w \in W \quad (4.19c)$$

$$\mathbf{M}^* \mathbf{f}^u + \tilde{\mathbf{M}} \mathbf{f}^c = \bar{\mathbf{q}}^{\text{so}} \quad (4.19d)$$

Generally speaking, the more accurate the SO solution, the larger the penalty weight β and the smaller the path cost tolerance ϵ should be. Intuitively, a high quality SO solution should produce a link flow which is close to the true SO link flow and path costs that are close to the true path costs. Our mixed equilibrium then, should not be allowed to stray too far from the aggregate link flow, and the usable paths should be near least (marginal) cost. However, exactly which value is best and how they relate to common convergence criteria is not immediately clear. In order to best compare our results to existing work, we set our hyper parameters as the most stringent values (lowest ϵ , highest β) to closely reproduce the Sioux Falls minimum control ratio as reported in Z. Chen et al. (2020). This heuristic resulted in $\beta = 10$ and $\epsilon = 5\text{e}-4$. For the Pittsburgh network $\beta = 0.1$ and $\epsilon = 1\text{e}-8$; the least ϵ and largest β to bring its MCR close to the range reported for similar networks in Z. Chen et al. (2020). The choice

of hyperparameters also affects the solving time of the programs. Smaller values of ϵ result in smaller usable path sets reducing the number of variables and significantly reducing solve time.

Here too, network size presents some difficulty. From the perspective of Program 4.19, a high-quality SO approximation will assign a large fraction of flow to paths that are close to least marginal cost. Common equilibrium convergence criteria, such as relative gap or average excess cost (Boyles, Lownes, and Unnikrishnan 2021), do not capture this notion, making it difficult to know *a priori* whether an approximation is good enough, or how to set β and ϵ to compensate. This fact is not unique to our problem setting and is a challenge to any analysis which takes an equilibrium traffic assignment as input.

Taken together, these modifications allow us to use approximate SO (UE) link flow to compute an upper (lower) bound on CFS-SO (CFS-UE). In our experiments we employ the CVXPY modeling language (Diamond and Boyd 2016; A. Agrawal et al. 2018) with GUROBI v9.5.0 as the backend solver. All experiments were performed on an Ubuntu 20.04 Linux machine with 16GB RAM and an Intel i7-7700HQ processor.

4.5.1 Experimental setting and main results

Our experiments were performed on three networks. The Braess network served to validate our methodology and code. The Sioux Falls network (Stabler, Bar-Gera, and Sall 2021) with 76 links and 528 OD pairs serves as a simple small example and as a point of comparison to previous work. The Pittsburgh network with 5449 links and 4881 OD pairs, serves as a real-world example network. The Pittsburgh network contains the core urban areas of Pittsburgh, PA extracted from the Southwestern Pennsylvania Commission’s (SPC) road network (www.spcregion.org). We report the critical fleet size results below in Table 4.4 and devote the remainder of the section to our findings on these networks. The MIPs for CFS-SO and CFS-UE on the Sioux Falls network were run for 24 and 2 hours respectively; we also report the relative gap of the solution reported by the solver.

Table 4.4: Comparison of critical fleet size on experiment networks.

Network	MCR		CFS-SO		CFS-UE		
	LP	LP	MIP	rel. gap	LP	MIP	rel. gap
Braess	100.00%	100.00%	100.00%	0.00%	0.34%	66.67%	0.00%
Sioux Falls	14.12%	36.53%	33.03%	25.28%	46.12%	78.85%	15.48%
Pittsburgh	30.69%	83.04%	—	—	56.92%	—	—

From Z. Chen et al. (2020) we know that the minimum control ratio for the Braess Network is 1 so we know that the CFS-SO must also be 1; a finding reproduced by

our methodology. CFS-UE on the other hand is achieved by a fleet which uses two of the three paths exclusively, leaving the remaining path for the users. Because the CFS-UE LP, unlike the MIP, cannot abandon paths, it provides a very loose bound.

4.5.2 Link flow patterns at CFS-SO

The two immediate observations are that CFS-SO is highly variable across the two networks and is substantially higher than the MCR. In the Pittsburgh network at most 83% of vehicles must be in the fleet in order to achieve SO traffic flow, compared with 33% in the Sioux Falls network. However, that the presence of a self-interested fleet can give rise to SO flow on real-world networks without controlling all of the vehicles at all is remarkable in and of itself.




Figure 4.3: Fleet and user link flows at the CFS-SO mixed equilibrium on the Sioux Falls Network. Color and width indicate the amount of flow of each class on each link. Arrows show the direction of link flow, and annotations denote the percentage of flow belonging to each class.

Geographically, CFS-SO on both Sioux Falls and Pittsburgh produces fleet presence across the network. We show the user and fleet link flows at CFS-SO for Sioux Falls in Figure 4.3 and for Pittsburgh in Figure 4.4. At least part of the reason for this is that fleet presence on one O-D pair influences the fleet marginal cost on many others, particularly in dense networks such as these where paths for different O-D pairs are highly intertwined. As a result, fleet presence in one area of the network tends to induce fleet presence on at least some of the O-D pairs that share links with

it, leading to a cascading effect over the entire network. This effect may be contrasted with the minimum control ratio which treats O-D pairs in isolation resulting in areas of the network with no fleet presence whatsoever.

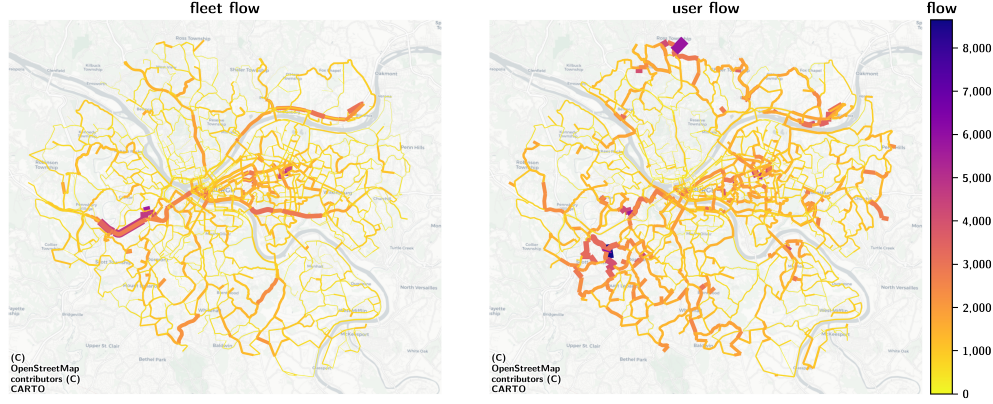


Figure 4.4: Fleet and user link flows at the CFS-SO mixed equilibrium on the Pittsburgh Network. Color and width indicate the amount of flow of each class on each link.

Nevertheless, fleet vehicles in Pittsburgh appear to accumulate along major highways and arterial roadways, suggesting that larger, more important roadways have an out-sized role to play in reducing total system travel cost with fleets. Fleets are most clearly present on the major highways leading into and out of Downtown Pittsburgh, near the center of Figure 4.4 and at the confluence of the three rivers. In contrast, user flow appears clustered in several relatively isolated geographic regions throughout the network, but notably along the periphery and within Pittsburgh’s densely populated East End, just east of Downtown.




Figure 4.5: CFS-SO and MCR both result in user and fleet exclusive O-D pairs. Here, the empirical cumulative distribution of the fraction of demand assigned to the fleet is shown for each solution concept and network.

4.5.3 Fleet O-D demand patterns at CFS-SO

Though fleets permeate the links of the network, they concentrate on paths and O-D pairs on the network. On both the Sioux Falls and Pittsburgh networks, the mixed equilibrium at CFS-SO divides most O-D pairs between being exclusively fleet and exclusively user. In analyzing the parallel network in Section 4.4.1 we noted that it is intuitive to imagine how CFS-SO would result in user or fleet exclusive O-D pairs, and less so to imagine O-D pairs with a mix of individual users and fleets. Indeed in both the Sioux Falls and Pittsburgh networks, about 95% and 85% of O-D pairs, respectively, are either exclusively user or exclusively fleet, as shown in Figure 4.5.

The requirement that fleet marginal cost be equalized across the paths used by the fleet between each OD pair, and specifically, that this equalizing requires fleet flow, is a key driver of the gap between MCR and CFS-SO. In the MCR setting, the fleet cost is simply the system marginal cost, which, along with link cost, is a constant. As a result, fleet flow can be assigned to those least marginal cost paths independently. By contrast, in the CFS setting, changes to fleet flow on one path must be balanced by fleet flow to equalize fleet marginal cost across used paths on each O-D pair.

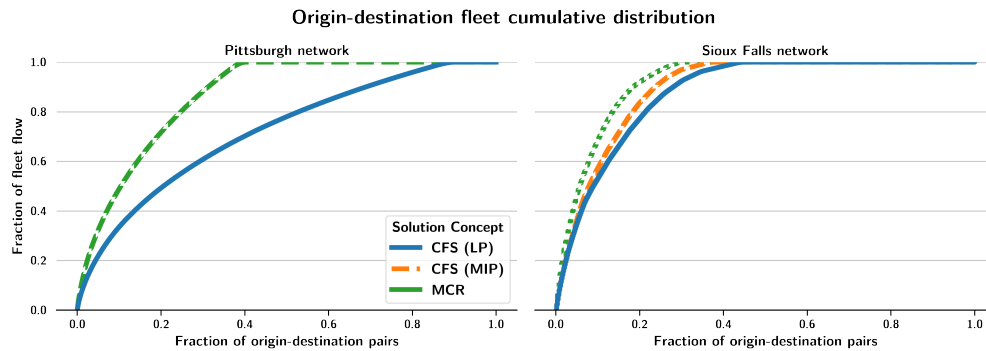


Figure 4.6: Cumulative fraction of fleet flow over O-D pairs. Fleet flow is concentrated with respect to O-D pairs: in Sioux Falls nearly 70% of the fleet comes from just 20% of O-D pairs. In the Pittsburgh network, roughly half of the fleet operates on the only 20% of O-D pairs.

If O-D pairs tend to be exclusive to either individual users or fleets, then transportation planners or fleet coordinators may focus their attention on a subset of O-D pairs accounting for an outsized proportion of the fleet. Figure 4.6 shows that for both the Sioux Falls and Pittsburgh Networks, a minority of the O-D pairs account for a majority of the fleet. In Sioux Falls just 10% of the O-D pairs house half of the fleet, for Pittsburgh, it takes 20% of O-D pairs. Remarkably, another 15% of O-D pairs in Pittsburgh do not require any fleet presence whatsoever.

4.5.4 Characteristics of paths and O-D pairs with fleet flow at CFS-SO

In both the Pittsburgh and Sioux Falls networks, O-D pairs with larger minimum marginal path cost generally have greater shares of fleet flow as shown in Figure 4.7. This conforms with the intuition that the fleet focuses on trips where congestion (and thus marginal cost) is greatest in order to achieve SO. In spite of this apparent trend, it is still difficult to distinguish between fleet-exclusive O-D pairs from user-exclusive ones *a priori*, or characterize what drives the fleet penetration at the O-D level. Simple criteria, such as those based on demand, O-D pair distance or (marginal) travel cost at SO, and number of paths fail to produce a useful understanding of the fleet penetration on O-D pairs. For example, despite the strong trend in Figure 4.7, marginal cost has little to no predictive power: among O-D pairs with the largest marginal cost only 60% have fleet flow on them. Instead we derive a metric, the path independence factor, from the optimality conditions of the CFS-SO LP capable of classifying fleet-exclusive paths with greater precision than simple criteria.

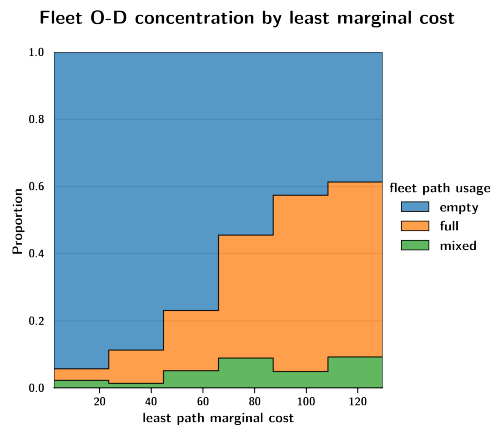


Figure 4.7: Proportion of fleet-exclusive (full), user-exclusive (empty), and mixed O-D pairs by the minimum marginal cost over paths for each O-D pair at CFS-SO (MIP) in the Sioux Falls network.

The path independence factor measures how sensitive fleet flow on a path could be to flow on other paths: how independent the cost of a path is from the flow on other paths. Fleet flow on a path that is highly intertwined with many other paths or has a large travel cost derivative has a greater capacity to increase fleet marginal cost than a path that is not. Following this intuition, we can compute path independence as the sum of the link cost derivatives on path r weighted by number of fleet usable paths

that do not include the link;

$$\text{PathIndependenceFactor}(r) = \sum_{p \in \tilde{P} \setminus \{r\}} \sum_{a \in A} \delta_{ap} \mathbf{t}'_a (1 - \delta_{ar}) = \sum_{a \in \text{links}(r)} \mathbf{t}'_a (|\tilde{P}| - N_a(\tilde{P})) \quad (4.20)$$

where $N_a(\tilde{P})$ is the number of paths in \tilde{P} containing link a . This term is derived from the part of the fleet marginal cost at SO flow that depends on fleet flow. A related form appears in the first-order optimality conditions for Program 4.17 (see Appendix A.9). The path independence factor is non-negative. It is zero for paths where every link is on every path in \tilde{P} . Paths with larger path independence have fewer high-cost interactions with paths in \tilde{P} . Empirically, paths with larger values of path independence tend to have more fleet flow on them, as shown in Figure 4.8. Paths that are completely unused tend to have the highest path independence factor; this is largely due to paths that have few interactions with fleet usable paths because they themselves are not fleet usable.

It is important to note that one may compute the path independence without having solved for CFS-SO; one only needs the set of least marginal cost paths at SO. As a first pass of testing the possibility that path independence might be used as a proxy for CFS-SO, we trained a Explainable Boosting Machine (EBM) (Lou et al. 2013; Nori et al. 2019) on the Sioux Falls path flows to predict whether the flow on a path that is usable by the fleet (i.e. it is least marginal cost) will be exclusively fleet or not at CFS-SO. Using just the path independence factor and whether the path may also be used by individual users (i.e. it is also least cost) the EBM was able to achieve an Area Under the Receiver Operating Characteristic Curve (ROC AUC) of 75%. In other words, the model will rank a random fleet exclusive path higher than a random non fleet exclusive path 75% of the time. As Figure 4.8 demonstrates, path independence is on wildly different scales on different networks so we cannot not expect any such model to generalize directly. Rather the model serves as a measure of the extent to which path independence is able to characterize this particular CFS-SO solution. The question of whether path independence can be generalized as a heuristic is left for future work.

4.5.5 Sensitivity of CFS-SO and CFS-UE to aggregate demand

For the Sioux Falls network, not only does the linear program offer a comparable value of CFS-SO, the solution itself tracks closely with the solution returned by the CFS-SO MIP. This is most evident in Figures 4.5 and 4.6 where the curves generated by the MIP tracks closely enough with the curves generated by the LP so that none of our conclusions would be altered by taking one solution instead of the other. To examine the quality of the linear program upper bound in a more general sense, we solve MCR, CFS-SO (exact and upper bound), and CFS-UE (exact and lower bound)




Figure 4.8: Paths with a higher independence factor (defined by (4.20)) tend to have more fleet, up to a point. The paths with the highest independence factor will tend to be unused completely. Here, ‘empty’ means at most 5% flow on the path is composed of fleet vehicles, ‘full’ means at least 95% flow on the path is fleet flow, and ‘mixed’ refers to the remaining paths.

on varied demand levels on the Sioux Falls network. Each mixed integer program was terminated at an absolute gap of 1% of the total travel demand or at a time limit of 2 hours, whichever came first. We compute an approximate lower bound on the CFS-SO, based on the best lower bound provided by the solver and on the assumption that the penalty term is an equal percentage of the objective at the incumbent solution and the best lower bound. Explicitly, the bound is computed using Equation (4.21), where CFS is the critical fleet size at the best integer solution returned by the solver. The results of the demand sensitivity analysis are summarized in Figure 4.9.

$$\text{CFS Lower Bound} = \text{CFS} \cdot \frac{\text{Objective Best Bound}}{\text{Objective Best Integer Solution}} \quad (4.21)$$

CFS-SO and MCR show similar trends with respect to the demand multiplier although CFS is more sensitive to changes in demand level. Both CFS-SO and MCR are small for very low demand levels, increasing with demand before decreasing at higher demand levels. In the very-low demand regime we expect CFS to be small or zero because we expect the network to be almost completely un-congested so that UE and SO are close. Encouragingly, past a certain point, both MCR and CFS-SO decrease, meaning that at high demands a smaller fraction of the total demand is needed for the system to benefit. Interestingly a similar pattern emerges for CFS-UE: past a certain demand level (roughly 80% of base demand) CFS-UE increases in demand level, implying that the network can handle high levels of optimized fleet demand without impacting UE traffic patterns at all.

In the Sioux Falls network, the linear program provides an upper bound within 6%

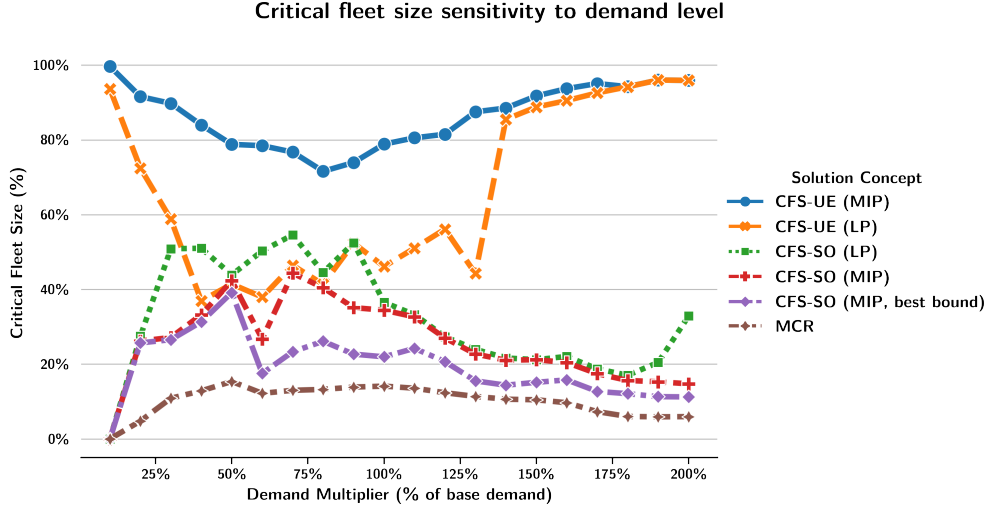


Figure 4.9: CFS-UE (solid blue line with circle marker), CFS-UE lower bound (orange line with 'x' marker), CFS-SO upper bound (dotted green line with square marker), CFS-SO MIP solution (dashed red line with '+' marker) with approximate lower bound (purple dot-dashed line with diamond marker), and MCR (Brown alternating dashed line with star marker) as a function of the demand multiplier which scales the demand volume on each O-D pair.

of the CFS-SO found by the mixed integer program for 14 of the 20 demand levels. For large demand levels, between 1.0x and 1.9x the base demand, the linear program provides a reasonable bound at a small fraction of the computational cost. It is in the low demand regime (0.2x-0.9x) where the upper bound performs particularly poorly, overestimating the CFS-SO found by the mixed integer program by up to 24%. However, the low demand regime is precisely where the linear program's computational advantages are least useful. For each demand level less than or equal to 50%, the MIP was solved to a relative gap of at most 8% within 2 hours. In these cases the upper bound is loose because the MIP was able to find high quality solutions. In contrast, no demand level above 50% was solved to a relative gap smaller than 22% within the 2 hour time limit. The LP lower bound for CFS-UE, however, is reasonably tight only for very low and very high demand regimes and at worst under-estimates CFS-UE by around 45%. The solutions returned by the CFS-UE MIP within the two hour time limit however, are of comparatively better quality than the solutions returned by the CFS-SO MIP within the same time limit. The CFS-UE solutions have an average relative gap of 3.6% compared to 24.5% for the CFS-SO solutions. Overall this indicates that the CFS-UE MIP is somehow less challenging for the solver than the CFS-SO MIP.

In general is it difficult to justify the tightness of the LP bounds for either CFS-UE or CFS-SO. Part of the difficulty is that the LP equalizes the fleet marginal cost over a potentially much larger set of paths than the set the true CFS-SO fleet would

use. It is not clear how to estimate *a priori* how much this affects the amount of fleet flow required. An option for practitioners who need both a bound and a measure of tightness is to run the MIP for a fixed computational budget.

4.6 Discussion

In both Pittsburgh and Sioux Falls networks, a non-trivial mixture of self-interested behavior induces SO traffic flow. Should CFS-SO be achieved, travelers would be on average are better off with the presence of the fleet than they would have been if they had all routed themselves independently. In contrast, the MCR, which also achieves SO traffic flow, may not make fleets better off. This important distinction highlights the advantage of explicitly modeling fleet behavior in traffic analysis and demonstrates the necessity of doing so. In this sense, of the many ways to partition SO as a mixed equilibrium, CFS-SO offers a particularly compelling case for its realism: no external subsidy is required.

Table 4.5: Coordination discount and average cost by flow class, solution concept, and network

network	solution concept	SO discount (%)	Coordination discount (%)		Average cost	
			fleet	user	fleet	user
Sioux Falls	CFS-SO (LP)	-3.82	-4.32	-3.30	27.64	15.52
	MCR		+3.44	-5.81	32.64	17.86
Pittsburgh	CFS-SO (LP)	-1.19	-1.19	-1.28	110,188.88	47,590.55
	MCR		-0.21	-2.34	135,722.77	83,179.67

Are fleets better off? Collectively, at both CFS-SO and MCR, individual users have lower average travel cost than the fleet simply because individual users use only least cost paths. One interesting effect of CFS-SO relative to MCR is that even though both achieve aggregate SO, they do so by partitioning demand between the fleet and individuals in very different ways as demonstrated by Table 4.5. First, the average travel cost, computed as the flow-weighted average path travel cost, is much higher for the fleet than for the user both at MCR and CFS-SO. This is partly due to the fact that individual users are only using least cost paths, but it is also due to the fact, demonstrated in Figure 4.7, that fleets accumulate on O-D pairs with large marginal cost. When the link cost function is monotonic increasing and convex, we should generally expect a positive correlation between cost and marginal cost ($t(x) > t(y) \implies t'(x) > t'(y)$). As a result, because fleet volume skews towards high marginal cost O-D pairs, and will not necessarily use least cost paths, the average fleet cost is much

higher than the average user cost.

Second, the average travel cost of both the fleet *and* the individual users decreases from MCR to CFS-SO, even though the aggregate average cost is the same. This counter-intuitive result is possible only because the fleet and users have completely different O-D demand between MCR and CFS-SO. In shifting from MCR to CFS-SO, the fleet needs to acquire flow on lower cost paths, it achieves this by taking paths from users that are lower cost than the average fleet path cost, but higher cost than the average user path cost, resulting in a reduction in both averages.

Third, the travel cost discount that the fleet receives relative to UE is greater at CFS-SO than at MCR. Moreover, at MCR on the Sioux Falls network, the fleet travel time *increases* relative to UE. We can compare the fleet and individual average travel cost under CFS-SO and MCR with the average travel cost of a set of vehicles with the same O-D demand under aggregate UE using a measure we term the ‘coordination discount’:

$$\text{CoordinationDiscount} = \frac{\langle \mathbf{c}, \mathbf{f} \rangle}{\langle \lambda_{\text{ue}}, \mathbf{M}\mathbf{f} \rangle} - 1 \quad (4.22)$$

where \mathbf{f} is the path flow of interest (i.e. the fleet, user, or aggregate path flow) under the solution concept of interest (i.e. MCR, CFS-SO, or aggregate SO), \mathbf{c} is the path cost at \mathbf{f} , and λ_{ue} is the O-D cost under aggregate UE. When \mathbf{f} is an aggregate SO path flow the coordination discount is termed the SO discount and is related to the price of anarchy as $1/\text{PoA} - 1$. Positive values of the coordination discount indicate that the flow in question is *worse-off* under its current routing than they would have been at aggregate UE; negative values indicate that the flow in question is better off using its current routing. In Table 4.5 we can see that at MCR in Sioux Falls the fleet total cost *increases* 3.44% relative to UE even though the *system* total cost *decreases* by 3.82%. This is an example of the system optimality paradox discussed in Section 4.3.3. In contrast, a fleet at CFS-SO, *reduces* its total cost by 4.32% relative to UE. In Pittsburgh the fleet operator is only slightly better off under MCR than under UE, but could increase their savings to nearly match the system SO discount by optimizing their fleet.

4.6.1 CFS-SO as optimization-by-proxy

Viewed a different way, our results demonstrate that SO traffic flow is attainable by optimizing over a subset of vehicles, rather than over all vehicles. In optimizing the entire network, the traffic manager should know, at the very least, the total system travel time, a measurement over all vehicular flow in the network. In practice, this measurement is simply not feasible to take. What a nontrivial critical fleet size indicates is that the measurement need only be taken and the optimization need only be performed on a subset of the vehicles. If this subset of vehicles is easy to measure

(e.g. they are all using the same navigation software or information platform) then all that is additionally necessary are travel time estimates on the links of the network, which are substantially easier to measure than untracked vehicles.

But how large could we reasonably expect such a fleet to be? According to a 2015 study, over 90% of smartphone users use their smartphones for directions (Anderson 2016), roughly 73% of Americans in 2020 owned smartphones (O’Dea 2020). Finally, 84% of those who use smartphones for directions use a Google-owned navigation app (Ceci 2021). By a simple back-of-the-envelope calculation, nearly 66% of drivers use some form of smartphone navigation app and 55% of drivers use a Google-owned navigation app. By contrast, it is estimated that ride-sourcing vehicles make up at most 14% of vehicle miles traveled among the largest urban areas in the United States (Balding et al. 2019). While still less than the CFS-SO on the Pittsburgh network, an optimized fleet of Google Maps users is large enough to achieve SO on many networks, Sioux Falls included, if distributed in the right way.

It’s worth noting that CFS-SO finds the smallest total fleet size capable of achieving SO with no constraints on its O-D demand. Specifically in drawing an analogy to Google Maps, it may be of interest to find CFS-SO that balances the total fleet size with how reasonable the resulting O-D fleet demand pattern appears, for example how uniformly it is distributed across O-D pairs. We leave this as an area for future work.

In general however, the scale of the empirical CFS-SO results would seem to preclude the most obvious source of an optimized fleet; in the case of Sioux Falls for example it’s unlikely that a UPS or FedEx fleet would account for over a third of traffic volume, but connectivity enables the coordination of vehicles on a much larger scale. FedEx can optimize because it either owns vehicles and employs CDL (commercial driver’s license) drivers or contracts with freight operators. Both avenues are capital intensive and accordingly limit the size of the fleet. In contrast, Uber and Lyft have far lower driver acquisition costs because they do not own the vehicle, nor employ the driver (who needs only a standard driver’s license) directly. In effect, CFS-SO, and in many networks MCR as well, could *only* be achieved via the scale of vehicle coordination enabled by app-based services.

Finally, CFS-SO and CFS-UE capture two very different realities for fleet operators. When fleets on the network induce the same traffic flow as UE would, we are right to ask whether the fleet should bother optimizing itself at all. If mixed equilibrium is the same as UE, then the fleet has not gained anything by optimizing. In fact, if there is a cost to coordinating behavior, then the service is strictly worse off in coordinating the fleet. In contrast, when mixed equilibrium is SO, the potential benefits of optimization are clear. In fact, in both the Sioux Falls and Pittsburgh networks, the fleet vehicles are strictly better off optimizing themselves while at the same time driving the system to its optimal traffic flow. In contrast, they are worse off under the MCR.

4.6.2 Applications of CFS-SO and CFS-UE

There are three principal applications we see for critical fleet size analysis; it offers,

1. regulators a new way to inform congestion management policy,
2. transportation planners an important metric to evaluate road network improvements,
3. service providers, such as ride-sourcing companies, navigation services, and, car manufacturers increasingly interested in mobility more broadly, a better understanding of the network effects of their fleet coordination, specifically in identifying areas of the network where their efforts could align with system level goals, and,
4. the public a tool to understand how individual users might benefit from increased fleet presence on their roadways.

In the search for regulatory strategies for *ad hoc* fleets, including ride-sourcing services, CFS-SO may provide insights into corridors where fleets would be particularly helpful in reducing system-level congestion. Conversely, CFS-UE may provide useful insights into corridors where *ad hoc* fleets would not be expected to meaningfully change travel time or cost. Critical fleet size analysis can therefore inform more targeted congestion management schemes by recognizing where the interests of traffic managers and fleet coordinators align, where they do not align, and where fleets will not be expected to alter traffic patterns.

CFS-SO is also an important metric in evaluating a set of alternative road network improvements. A road network improvement that reduces CFS-SO could be leveraged by a fleet to improve system cost, potentially to SO levels. All else equal, a road improvement which offers a greater reduction in CFS-SO will pay dividends in the future as it makes it easier for a coordinated fleet to reach the threshold required to bring the system to optimality, without any external subsidy or intervention on the part of the transportation planner.

Services that are able to coordinate their fleets should ask themselves whether and how to optimize their fleet as it is often in their interest to guide the fleet toward a service-level objective. CFS-SO offers service providers a window into how well their service may align with network level objectives. In the case of green routes in Google Maps, this is an existential question. CFS-UE offers service providers insight into where optimizing their fleet may not be worth it at all.

4.7 Regularized critical fleet size

As formulated, our notion of critical fleet size simply seeks to find the OD demand pattern of the smallest possible fleet to induce system optimal traffic flow. In practice

it might be useful to prefer a different OD demand pattern that is larger but better conforms to some intuition. We discuss two specific cases in this section. In the first, termed demand-regularized CFS, we design a critical fleet size program which prefers OD demand patterns that are similar to some prior. For example, the prior could be the current OD demand pattern of a fleet of interest or the prior could be simply that the fleet should be as uniform as possible. In the second, termed the maximum likelihood CFS, we design a critical fleet size program which prefers fleet paths whose travel cost is close to the least cost path for each OD pair.

4.7.1 Demand-regularized critical fleet size

The demand regularized CFS-SO linear program is given below as Program 4.23 where $C > 0$ is a constant and g is the regularization function. Analogous versions can be constructed for the CFS-SO MIP, MCR, and CFS-UE (LP and MIP).

$$\min_{\mathbf{f}^u \geq 0, \mathbf{f}^c \geq 0, \lambda} \sum_{p \in \tilde{P}} \mathbf{f}_p^c + Cg(\mathbf{q}^c) \quad (4.23a)$$

$$\text{s.t. } \tilde{\mathbf{c}}_r(\mathbf{f}^u) = \lambda_w \forall r \in \tilde{P}_w \forall w \in W \quad (4.23b)$$

$$\tilde{\mathbf{c}}_r(\mathbf{f}^u) \geq \lambda_w \forall r \in P_w \forall w \in W \quad (4.23c)$$

$$\mathbf{D}^* \mathbf{f}^u + \tilde{\mathbf{D}} \mathbf{f}^c = \bar{\mathbf{x}}^{\text{so}} \quad (4.23d)$$

$$\mathbf{M}^* \mathbf{f}^u + \tilde{\mathbf{M}} \mathbf{f}^c = \bar{\mathbf{q}} \quad (4.23e)$$

When we have access to a demand pattern $\hat{\mathbf{q}}$ that serves as a prior on the distribution of the fleet demand, the regularized CFS program will balance total fleet size while remaining close to the given demand pattern. In this setting, we can simply take g to be the distance between the fleet demand variable and the prior demand as below;

$$g(\mathbf{q}) = \|\mathbf{q}^c - \hat{\mathbf{q}}\|^2 \quad (4.24)$$

The notion that demand should be uniform can be expressed with the following definition of the regularization function:

$$g(\mathbf{q}) = \sum_w \left(\frac{\mathbf{q}_w^c}{\mathbf{q}_w} - \frac{1}{|W|} \sum_w \frac{\mathbf{q}_w^c}{\mathbf{q}_w} \right)^2 \quad (4.25)$$

4.7.2 Maximum likelihood critical fleet size

The maximum likelihood critical fleet size LP is given below in Program 4.26 where $z_p \geq 0$ denotes the weight of the path.

$$\min_{\mathbf{f}^u \geq 0, \mathbf{f}^c \geq 0, \lambda} \sum_{p \in \tilde{P}} z_p \mathbf{f}_p^c \quad (4.26a)$$

$$\text{s.t. } \tilde{\mathbf{c}}_r(\mathbf{f}^u) = \lambda_w \forall r \in \tilde{P}_w \forall w \in W \quad (4.26b)$$

$$\tilde{\mathbf{c}}_r(\mathbf{f}^u) \geq \lambda_w \forall r \in P_w \forall w \in W \quad (4.26c)$$

$$\mathbf{D}^* \mathbf{f}^u + \tilde{\mathbf{D}} \mathbf{f}^c = \bar{\mathbf{x}}^{\text{so}} \quad (4.26d)$$

$$\mathbf{M}^* \mathbf{f}^u + \tilde{\mathbf{M}} \mathbf{f}^c = \bar{\mathbf{q}} \quad (4.26e)$$

We define Δ_r to be the **excess cost** of a path r which is computed as the difference between the travel cost on that path c_r and the least travel cost of all paths connected the same OD pair c_r^* :

$$\Delta_r = c_r - c_r^* \quad (4.27)$$

We define a family of weights, indexed by k , as follows:

$$z_p^{(k)} = \Delta_r^k \quad (4.28)$$

The first and second order weights ($z^{(1)}$ and $z^{(2)}$) are related to the likelihood of the fleet path flow \mathbf{f}^c under an exponential choice model and Gaussian choice model respectively.

We suppose that travelers prefer lesser travel costs. When the fleet mode is available between a given origin-destination pair it will necessarily offer a travel cost that is at least as large as the UE mode, which will always use the least cost path(s). Therefore, travelers who choose the fleet mode incur an excess cost. We define a probability density function g which maps path excess cost to the probability that the path is used. The likelihood of any path flow is then;

$$L(\mathbf{f}) = \prod_r g(\Delta_r)^{\mathbf{f}_r} \quad (4.29)$$

The negative log likelihood is then the weighted sum of paths flows;

$$-\ln L(\mathbf{f}) = -\sum_r \mathbf{f}_r \ln g(\Delta_r) \quad (4.30)$$

First we take g to be the exponential probability density function, $g(x) = \lambda \exp(-\lambda x)$, whose rate parameter λ determines how sharply the probability density decreases with x . When g is the exponential density function, the likelihood simplifies nicely:

$$-\sum_r \mathbf{f}_r \ln g(\Delta_r) = -\sum_r \mathbf{f}_r \ln(\lambda \exp(-\lambda \Delta_r)) \quad (4.31)$$

$$= -\sum_r \mathbf{f}_r (\ln(\lambda) - \lambda \Delta_r) \quad (4.32)$$

$$= \lambda \sum_r \Delta_r \mathbf{f}_r - \ln(\lambda) \sum_r \mathbf{f}_r \quad (4.33)$$

This second term is a constant since the aggregate demand is fixed. Further since Δ_r is zero for all UE paths by definition, the objective function for the ML-CFS under an exponential choice model is simply;

$$\sum_r \Delta_r \mathbf{f}_r^c = \sum_r z_r^{(1)} \mathbf{f}_r^c \quad (4.34)$$

When g is a half-normal (or normal) density function we can express the negatively log likelihood as follows, phrasing the density in terms of precision $\rho = \frac{1}{\sigma^2}$;

$$-\sum_r \mathbf{f}_r \ln g(\Delta_r) = -\sum_r \mathbf{f}_r \ln \left(\sqrt{\frac{2\rho^2}{\pi}} \exp \left(-\frac{1}{2}\rho^2 \Delta_r^2 \right) \right) \quad (4.35)$$

$$= -\sum_r \mathbf{f}_r \left(\ln \left(\sqrt{\frac{2\rho^2}{\pi}} \right) - \left(\frac{1}{2}\rho^2 \Delta_r^2 \right) \right) \quad (4.36)$$

$$= \frac{1}{2}\rho^2 \sum_r \Delta_r^2 \mathbf{f}_r - \frac{1}{2} \ln \left(\frac{2\rho^2}{\pi} \right) \sum_r \mathbf{f}_r \quad (4.37)$$

As before, the second term is a constant in our optimization so the objective simplifies:

$$\sum_r \Delta_r^2 \mathbf{f}_r^c = \sum_r z_r^{(2)} \mathbf{f}_r^c \quad (4.38)$$

In practice we take $\Delta_r = (1 + \epsilon)c_p - c_p^*$ for some small constant $\epsilon > 0$ (in this work we take $\epsilon = 0.01$) which ensures that $\Delta_r > 0$ for all paths r . This small constant can be taken to represent a cost associated with the fleet mode (e.g. a mileage based charge). The results are not very sensitive to the choice of ϵ , but when $\epsilon = 0$ the optimization may use least cost routes “for free”, which not only produces values of MCR and CFS which are artificially large, but produces a large set of optimal solutions.

Experiments

We compute the ML-CFS and ML-MCR on the Sioux Falls network for the exponential ($k = 1$) and normal ($k = 2$) choice models as well the no choice model scenario ($k = 0$). We report our results in Table 4.6. The fundamental trade-off the choice model navigates is between fleet and user average costs. The choice model implicitly penalizes higher cost paths and as a result, lowers the average traveler cost at the expense of a larger fleet.

4.8 Demographic analysis of critical fleet size

In discussing the system optimality paradox we noted that when flow on the network is partitioned, what is optimal for the system may not be desirable for every partition.

Table 4.6: Maximum Likelihood CFS Results

k	Fleet fraction (%)		Average cost	
	CFS-SO (LP)	MCR	CFS-SO (LP)	MCR
0	35.96	14.13	27.44	32.38
1	40.34	14.32	22.14	29.53
2	52.04	19.32	19.10	28.03

In this section we ask whether travel cost and fleet flow are different over partitions of flow based on race and poverty status. Concretely we study the intersection of race and poverty status at the tract level as reported by the US Census. We report the following findings for the Pittsburgh network.

1. The optimality gap for all race poverty subpopulations is negative, meaning that all race-poverty groups benefit under SO relative to UE.
2. Within each racial group, poverty status does not affect optimality gap.
3. Black populations benefit the most from SO; Asian populations benefit the least.
4. At CFS-SO there is no significant difference in fleet access between race-poverty groups.
5. At CFS-SO all race-poverty groups benefit whether they are using the fleet or not and the benefit is not significantly different between race-poverty groups.
6. At CFS-SO there is no significant difference between fleet volume per lane mile experienced across the race-poverty groups.

4.8.1 Data and methods

Our race and poverty data was sourced from 2020 American Community Survey (ACS) poverty status dataset (American Community Survey 2020). This dataset reports the total population and number of individuals living below the poverty line for each racial group at the tract level. The geometries for each tract were sourced from the 2020 census TIGER/Line shapefile for the state of Pennsylvania.

We employ two methods to associate traffic flow from the CFS-SO solution with census tracts. The first measures **access** whereby volume on a path is associated with a tract if the origin node of the path is contained in the tract. The second measures **exposure** whereby volume on a link is associated with a tract if the link intersects the tract. These two methods will differentiate a tract with volume from trips that pass through it from tracts with volume from trips originating from it. This is particularly important when considering the fleet vehicles in CFS-SO. A community with low access but high exposure to the fleet could argue that it is bearing the cost of fleet traffic

without enjoying any of its benefits. It should be noted that we use the term access to mean that the service is geographically accessible to the population of the tract. It is incumbent on the municipality to ensure that the service is economically and physically accessible. The exposure analysis focuses on the volume per lane-mile within the tract as a measure of congestion whereas the access analysis focuses jointly on fleet penetration within the tract and the average travel cost of all trips leaving the tract. We form an aggregate measure for the network by taking the average over tracts of the tract measurement (volume per lane-mile, penetration, and travel cost) weighted by the population group of interest within each tract. Concretely, the measure M for a statistic X for race r and poverty status p is given by,

$$M(X, r, p) = \frac{\sum_t N_t(r, p) X_t}{\sum_t N_t(r, p)} \quad (4.39)$$

where $N_t(r, p)$ is the number of individuals residing in tract t of race r and poverty status p and the summations are taken over tracts. This aggregate measure can be viewed as the expected value of the statistic X for a randomly selected individual of race r and poverty status p .

4.8.2 Results

Optimality gap by race and poverty

All race-poverty groups benefit from SO relative to UE as measured in percent travel time savings relative to UE, saving on average nearly 2% of travel time. The magnitude of the benefit is not significantly different (at the 95% confidence level) across race-poverty groups with one exception: black individuals above the poverty line benefit from SO more than Asian individuals living below the poverty line. However, this difference, while significant, is not robust and vanishes after removing either of the tracts with the largest population, about 2%, of each race-poverty group. Given the racial bias in highway construction, this is a somewhat surprising result.

Access

There is no significant difference in the ratio of fleet to user flow between the race-poverty groups. This uniformity is somewhat surprising given that fleet flow on the Pittsburgh network is highly concentrated with respect to OD pairs. What this means is that even though a full 15% of OD pairs in the Pittsburgh network have no fleet vehicles at all, an individual is equally likely to be nearby an OD pair with fleet service regardless of their race or poverty status. Moreover the fleet trips are of comparable quality across race-poverty groups. Compared with the UE cost of travel, individuals benefit from fleet presence equally over race-poverty groups.

Slight disparities emerge however when we examine the tracts with the least access. The top 5% of tracts with the least fleet flow account for roughly 2% of the population and 0.08% of the fleet flow but house nearly 4% of the population experiencing poverty. Moreover, those with the least access are 1.5 times more likely to live below the poverty line compared to those with greater access.

Exposure

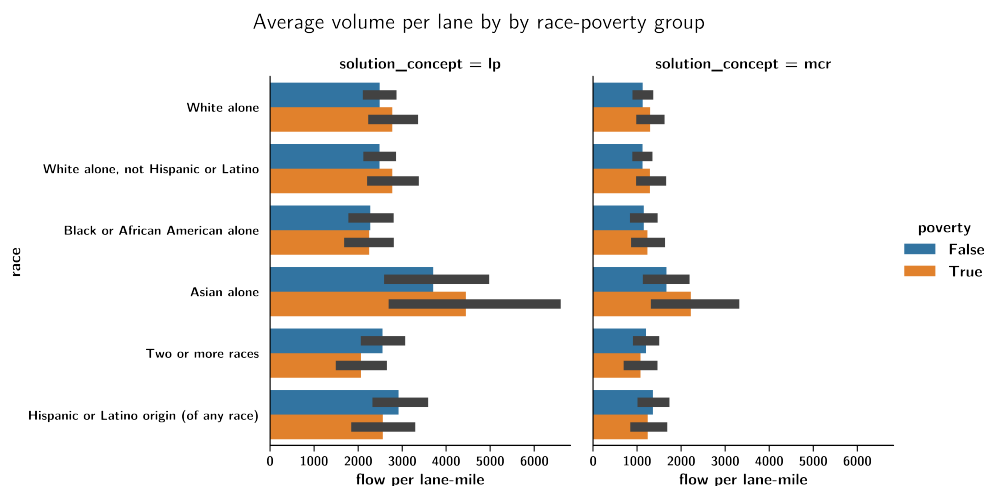


Figure 4.10: Fleet volume per lane mile by race-poverty group

Although there is no statistically significant difference in the fleet volume per lane-mile between the race-poverty groups, Asian individuals (particularly those under the poverty line) experience the highest fleet volume per lane-mile and Black or African-American individuals among the lowest, as shown in Figure 4.10. This is worth noting because Asian individuals benefit the least from SO but experience higher levels of fleet traffic in their home tracts relative to the other race-poverty groups who benefit more from SO but experience lower levels of fleet traffic. Although the inequity is not large enough to be called statistically significant at the 95% level, it may be large enough to cause concern if the presence of fleet vehicles in practice is enough of a burden.

Encouragingly there is a positive relationship between fleet volume per lane mile within a tract and the fleet flow servicing the tract implying that access and exposure are correlated. Notably, this relationship is preserved across poverty level and race: the top 10% of tracts with the highest share of people experiencing poverty have the roughly the same relationship as those in the bottom 90% and the top 10% of tracts with the largest share of non-white populations have roughly the same relationship as tracts with lower shares of non-white populations. In fact, tracts with higher shares

of non-white populations experience a higher ratio of access to exposure than tracts with lower shares.

4.9 Conclusions and future work

As technological advances in transportation continue to permeate our road networks, it becomes easier for services to coordinate vehicle behavior toward a service-level objective. In this work, we study if and how self-interested services that optimize their fleet can reduce system cost on a road network. We demonstrate that self-interested fleets are capable of both paradoxically increasing system cost over UE and reducing system cost up to achieving SO network flow. We provide a mathematical program with equilibrium constraints to solve for the smallest fleet that would induce SO, and the largest fleet that would induce UE, in mixed traffic with UE users. We present efficient solution methodologies and apply them to the Sioux Falls and Pittsburgh networks finding that 33% and at most 83% of vehicles respectively, must join the fleet in order for the network to reach SO. Moreover, these vehicles are better off than they would have been had they either routed themselves independently or participated in a MCR-like scheme, meaning that neither subsidies nor tolls would be required to compensate the fleet. At CFS-SO on the Pittsburgh network, fleet vehicles are most prevalent along highways and major arterials, while individual users tend to accumulate on shorter trips outside of the central business district. The ‘path impact’ metric, derived from the KKT conditions of the CFS linear program, is found to drive the concentration of fleet vehicles on paths. In other words, we have found two examples of networks on which a nontrivial mixture of self-interested behavior can induce minimum system cost on the network, without the need for external subsidy. Critical fleet size offers regulators a new way to inform congestion management policy, transportation planners a new metric to evaluate network improvements, and fleet operators/coordinators a better understanding of benefits of their fleet optimization efforts at the level of both their own platform and the entire network.

A key avenue of future work concerns the conditions under which a fleet will reduce system cost on a network. Although the condition given by Equation (4.3) yields some theoretical insight, it is not of practical use. A better practical condition would depend only on the network itself and, potentially, the UE or SO link flow. Whether such a condition exists and if it is easier to compute than mixed equilibrium is an area for future work. Additionally, the example of the fleet optimality paradox is contrived. Whether the paradox can be demonstrated in real world networks remains to be seen. If the paradox can be found in real world networks, these examples may inform better conditions or heuristics to characterize networks that may suffer from the presence of fleets.

Qualitatively, CFS is influenced both by aggregate demand patterns and also the

structure of the network itself. In one sense CFS and price of anarchy relate similarly to both network structure and demand levels, although the nature of this relationship is far from clear. At a high level this correspondence is intuitive: the price of anarchy measures the gap between congestion at UE and SO and CFS-SO measures the size of the fleet it would take to close it. As we have shown, CFS-SO is smallest when the network is either completely uncongested or highly congested, corresponding to low values of price of anarchy. Similarly, just as the Braess paradox shows us that network structure affects the price of anarchy, we should expect changes to the network structure to affect CFS-SO. The precise nature of this relationship is unknown and is left for future work. In particular, the question of how small CFS-SO can be made under a fixed budget for network intervention is a relevant and interesting network design problem. On the one hand such a network design problem generates promising road network improvements targeted at making it easier for coordinated fleets to optimize the network and on the other it would shed light on what kinds of network structure influence CFS-SO the most.

The use of critical fleet size in transportation planning exposes an interesting avenue of future work in characterizing the kinds of improvements which tend to reduce critical fleet size in real world networks. Such work would expand our understanding of what drives critical fleet size and more broadly, what drives system cost reduction when optimized fleets are present on the network. It may be of interest for public agencies to regulate a fleet coordinator toward a deployment pattern inline with CFS-SO. One would expect the fleet-optimal operation, in conjunction with incentives, may lead to SO flow, or considerably mitigate congestion.

As formulated, CFS enforces no constraints on the O-D fleet demand. This is intentional as in this work we are focused on finding strictly the smallest optimized fleet capable of inducing SO. There are larger fleets which would also achieve SO, but as our analysis shows, adding fleet to the CFS-SO demand pattern may break SO. As we discuss, a CFS O-D demand pattern that adheres to some prior notion of fleet demand might be more useful. Adding a penalty term to the CFS mathematical programs would be one way to achieve this and we leave such studies for future work.

In this chapter CFS is solved with the presence of only a single fleet. In reality, there could exist many such fleets. CFS-SO with multiple fleets finds the smallest total number of fleet vehicles such that the mixed equilibrium is SO. In Appendix A.7, our mathematical programs are expanded with additional constraints to ensure that *each* fleet meets FO criteria and that aggregate demand and link flows are conserved. Experiments on real networks and whether CFS increases or decreases when multiple fleets are present is left for future work. We hypothesize, however, that because multiple fleets will compete with each other, CFS-SO will be minimized when only one fleet is on the network. Refuting or proving this hypothesis is also left for future work.

In commenting on congestion, critical fleet size considers route-choice behavior

alone. Our framework elides many congestion-influencing behaviors that have been attributed to *ad hoc* fleets in the literature. A complete treatment of the congestion effects of such fleets should not only take into account the route choices, but also the behaviors associated with that particular service. Of course, not every service provider wishes to optimize total fleet travel cost, which this particular formulation of critical fleet size assumes. An important area of future research is then to adapt the methods in this work to other fleet-level objectives a service may have, such as total emissions in the case of Google Maps' green routes.

The relationship between the two notions of critical fleet size on networks with a non-unitary price of anarchy is not clear. Characterizing the relationship either empirically through further experiments on real-world networks or theoretically would expand our understanding of both critical fleet size as a quantity of interest and how networks in general should be expected to respond to the presence of optimized fleets.

We study CFS in the context of static equilibrium. This was a deliberate choice as we leveraged the simplicity of the static setting to isolate fleet behavior as the sole reason for the increase in total system cost in the FO paradox and as the sole reason that a mixture of fleets and individuals could achieve SO. CFS is also an interesting quantity in the dynamic setting and is an area of future work we are eager to explore.

Finally, solving critical fleet size, or its bound, in large networks is computationally challenging. Not only does the problem require a high-quality SO link flow solution, but the number of variables and constraints grows quickly with network size. There is no immediately obvious way of reducing the problem complexity, but a more clever row and column generating scheme, and tighter integration with solver software, may work well in practice to solve the smallest possible critical fleet size program. Finding ways to improve the computational efficiency of critical fleet size algorithms will make it a more easily applicable algorithm.

5

Learning from learners: behavioral inference from non-stationary policies

5.1 Introduction

When we observe a sequence of binary decision *outcomes* generated by agents involved in *active learning* can we deduce the reward signal they appear to be chasing? Such problems arise naturally when we are able to observe agent behavior but not the intent that generated the behavior. When the *policy* the agent is using does not change over time, existing methods such as inverse reinforcement learning and choice modeling are available to identify a reward function which best rationalizes the observed behavior. Existing methods, to the extent that they consider uncertainty, consider agents who have only *aleatory uncertainty* about the relationship between action, environment, and reward. As a result, these methods assume that successive observations of state, action, and reward carry no information about the environment that is not already known to the agent and will not affect their view of the world. In other words, when we use existing methods we assume that the agent whose behavior we are observing is no longer learning anything about their environment. However, it is often unreasonable to assume that an agent has no *epistemic uncertainty* about the relationship between state, action, and reward; particularly if the agent is interacting with a complex environment with a high-dimensional action space. In these cases, the observation of an outcome (and corresponding reward) will influence the agent's policy moving forward; in other words, the agent is learning successively better policies over time and we cannot apply existing methods to rationalize their behavior. Given a sequence of *outcomes* generated by an agent, we develop a sequentially-valid test statistic which reports when in the observation window, if ever, there is enough evidence to reject the notion that the agent is not seeking (avoiding) the outcome at a pre-specified confidence level. Notably, the test statistic does not require any knowledge of the action space. We demonstrate the test statistic on simulated interactions

with a multi-armed bandit and to the behavior of real ride-sourcing drivers before and during the COVID-19 lockdown in the city of Pittsburgh, Pennsylvania.

This methodology is most useful when we wish to rationalize observed decision making when it is likely that the decision-makers themselves are modifying the criteria they use to make decisions *in response to* the outcomes of their decisions over time. In other words, the agents refine their policies to reduce their *epistemic uncertainty* about the relationship between action and reward. Conventional notions of choice modeling assume a time-independent criteria and thus will mistake the reduction in epistemic uncertainty for a large degree of aleatory uncertainty. In contrast, our methodology aims to first identify *if* agents appear to be meaningfully improving their reward over the observation period and if so, to then identify the region of the decision outcome space the agents appear to be attracted to over the observation period.

The methodology we develop here may be directly applied to a more general class of problems, to answer the question: is an event becoming more (less) likely over time? In this chapter we are interested in detecting changes in behavior so it is useful to conceptualize the event as being generated by a learning agent interacting with an environment, but our process makes *no* assumptions on the data-generating process. Our only requirements are that the outcomes are observed sequentially in chronological order and that past outcomes do not depend on future outcomes. In this sense our method belongs to a wider class of *online changepoint detection* methods capable of identifying the moment when there is enough evidence to suggest that the probability of an event may be non-stationary in time or, if more specificity is required, that this probability may be increasing (decreasing) in time.

The remainder of this chapter is organized as follows. Section 5.2 introduces our methodologies. We first detail how sequences of events are extracted from OpenStreetMaps feature tag data and the ride-sourcing driver GPS trajectories. We then introduce our test statistic. Section 5.3 introduces and describes the datasets we use in this chapter. In Section 5.4 we then offer a descriptive analysis of the Ride-sourcing GPS trajectory data with a particular emphasis on comparisons between 2019 and 2020 trajectory data. In Section 5.5 we describe the set-up for our experiments. We first demonstrate the test statistic on two simulated datasets before describing how we apply the test statistic to our ride-sourcing GPS dataset. Section 5.6 reports the main results from the sequential hypothesis test on the ride-sourcing driver data. It is followed in Section 5.7 by a detailed quantitative and qualitative analysis of the results and their implications. In Section 5.8 we offer concluding remarks and directions of future work. We close this chapter in Section 5.9 with a discussion and preliminary results of the application of our procedure to empirically justify the use of a given travel cost function in a model of network equilibrium.

5.2 Methods

5.2.1 Inferring service-related stops from GPS trajectories

Suppose we have a sequence of n GPS pings that have been classified as coming from a stationary vehicle for the same driver on a shift ordered by timestamp. We are interested in grouping these pings into *service stops* representing a time and location for which a vehicle may have been involved in a service-related activity, for example a pickup or drop off. We will do so by estimating the following random variables:

- $S_i = 1$ if the i -th still ping was a service stop and 0 otherwise, and
- $C_i = 1$ if still pings i and $i + 1$ were associated with the same service stop.

Taken together these two indicators are used to transform a sequence of still GPS pings into a sequence of service stops using the procedure given by Algorithm 2. We describe each still ping by three features:

1. The speed measurement of the ping denoted by v_i ,
2. The time difference between pings i and $i + 1$ denoted by t_i , and
3. The distance between pings i and $i + 1$ denoted by d_i .

Algorithm 2 Stop inference procedure

```

procedure StopInference( $\{(v_i, t_i, d_i)\}_{i=1}^m$ )
     $n \leftarrow 0$  ▷ Initialize the service stop counter
    for  $i \in \{1, \dots, m\}$  do
         $S_i \leftarrow \arg \max_s \Pr[S_i = s \mid V = v_i]$  ▷ Infer  $S_i$ 
        if  $i > 1$  then
             $C_{i-1} \leftarrow \arg \max_c \Pr[C_{i-1} = c \mid T = t_{i-1}, D = d_{i-1}, S_{i-1}S_i = 1]$  ▷ Infer  $C_i$ 
            if  $C_i = 0$  then
                 $n \leftarrow n + 1$  ▷ Increment stop counter
            end if
        end if
        if  $S_i = 1$  then ▷ Assign a service stop label
             $N_i \leftarrow n$ 
        else
             $N_i \leftarrow -1$ 
        end if
    end for
    return  $\{N_i\}_{i=1}^n$ 
end procedure

```

First we consider S_i . Intuitively we are more confident that a still ping is in fact a stop if the speed measurement is small. We therefore model the probability that $S_i = 1$

given a speed measurement v by the logistic function whose mean \bar{v} we treat as a free parameter;

$$\Pr[S_i = 1 \mid V_i = v] = 1 - \text{LogisticCDF}(v; \bar{v}) \quad (5.1)$$

where,

$$\text{LogisticCDF}(v; \bar{v}) = \frac{1}{1 + \exp(-(v - \bar{v}))} \quad (5.2)$$

The use of the CDF of the logistic distribution captures our intuition that larger speed measurements imply that it is less likely that the ping was produced by a vehicle at a service stop. The parameter \bar{v} is the median of the logistic distribution and thus should be selected as the “break-even” speed: the speed measurement at which the vehicle is equally likely to be stationary or moving. Equivalently, \bar{v} is the largest (smallest) speed for which we would say that it is most likely that the vehicle was stationary (moving).

Next we consider C_i . If one or both of pings i and $i + 1$ are not service stops then clearly pings i and $i + 1$ cannot be associated with the same service stop. Conditioned on the event the $S_i \cdot S_{i+1} = 1$ then, the probability that $C_i = 1$ depends on the distance between the two pings, which we denote by D_i and the time difference between the two pings, denoted by T_i . The larger both D_i and T_i are the less likely it is that the two stops refer to the same location. As above, we encode this intuition via the CDF of the logistic distribution;

$$\Pr[C_i = 1 \mid T_i = t, D_i = d, S_i S_{i+1} = 1] = 1 - \text{LogisticCDF}(\rho t + d; \rho \bar{t} + \bar{d}) \quad (5.3)$$

where ρ serves to convert time into distance. As before \bar{t} and \bar{d} should be chosen to be our “break-even” time and distance: the largest time interval and distance we would be willing to accept between two consecutive pings associated with the same service stop. The conversion factor ρ has a ready interpretation as an average travel speed.

The location of a service stop is defined as the centroid of the GPS pings associated with the service stop. The timestamp of a service stop is defined as the earliest timestamp of the GPS pings associated with the service stop.

5.2.2 Non-stationary binary sequence testing

We observe a sequence of binary-valued random variables $\{X_i\}_{i=1, \dots, T}$. In this setting we imagine that they represent a reward or penalty experienced by an agent interacting with its environment, but this construction serves merely to motivate our particular application. Our objective is to determine the extent to which they appear

to be learning from experience: in short we ask, is the probability of $X_i = 1$ increasing (decreasing) over time?

In order to detect an increasing or decreasing probability of reward over time we will construct a generalization of a likelihood ratio known as an e-process against the null that the sequence of binary variables are i.i.d. Bernoulli random variables with a common unknown mean. Following Ramdas, Ruf, Larsson, and W. M. Koolen (2022), our process will take the form;

$$E_t := \frac{\prod_{s=1}^t (\hat{p}_s)^{X_s} (1 - \hat{p}_s)^{1-X_s}}{\binom{n_1^t}{t} \binom{n_0^t}{t}^{n_0^t}} \quad (5.4)$$

where n_1^t (n_0^t) are the number of times a 1 (0) has been observed up to and including time t and $\hat{p}_s \in [0, 1]$ is any estimate of $\Pr[X_s = 1 \mid X_1, \dots, X_{s-1}]$. This estimate may depend on the previously observed data X_1, \dots, X_{s-1} . The process by which the estimates \hat{p}_s are constructed will affect which alternatives the test has power under. In this work we will focus on three related estimators:

$$p_t^\pm = \frac{\sum_{s=1}^{t-1} (1 - \gamma)^{s-1} X_{t-(s-1)}}{\sum_{s=1}^{t-1} (1 - \gamma)^{s-1}} \quad (5.5)$$

$$p_t^+ = \max \left\{ \frac{n_1^{t-1}}{t-1}, p_t^\pm \right\} \quad (5.6)$$

$$p_t^- = \min \left\{ \frac{n_1^{t-1}}{t-1}, p_t^\pm \right\} \quad (5.7)$$

where $\gamma \in (0, 1)$ can be specified in advance or learned as data is collected. The first estimator, p_t^\pm , is simply an exponential weighted average of the data collected so far, representing the non-parametric alternative that the data-generating process is non-stationary but uses similar probabilities for data drawn nearby in time. A simple way to construct an estimator that makes Equation (5.4) sensitive in particular to non-stationary Bernoulli sequences for which the event probability is increasing (decreasing) is to leverage the following observations. When the probability of $X_t = 1$ is increasing (decreasing) in time, then the exponential weighted mean, p_t^\pm will be:

1. a better estimate of $\mathbb{E}_{t-1}[X_t]$ than the empirical mean, and
2. larger (smaller) than the empirical mean.

The second and third estimators encode these observations by restricting the range of the exponential weighted mean. The second estimator, p_t^+ , takes on values that are at least as large as the empirical mean and represents the alternative that the data-generating process has increasing probability of $X_t = 1$ over time. The third estimator, p_t^- , takes on values that are no larger than the empirical mean and represents the

alternative that the data-generating process is decreasing the probability of $X_t = 1$ over time. We note that the validity, that is the type-I error, of the test is unaffected by our choice of alternative likelihood.

Under the null, the likelihood using p_t^\pm will, in expectation, be at most the likelihood using the empirical mean precisely because the empirical mean is the maximum likelihood estimate for the mean of a sequence of i.i.d. Bernoulli variables. As a result, we have $\mathbb{E}_{t-1}[E_t] \leq 1$. However, under their respective alternatives and using the appropriate value of γ , the estimators p_t^\pm , p_t^+ , and, p_t^- provide a better estimate of $\mathbb{E}_{t-1}[X_t = 1]$ than the empirical mean does; as a result we should expect $\mathbb{E}_{t-1}[E_t] > 1$.

Although p_t^+ has power under increasing binary sequences, it does not have power under decreasing sequences. When we test a *decreasing* sequence using p_t^+ , the exponential weighted average will fall below the empirical mean and the max will cause $p_t^+ = n_1^{t-1}/t - 1$, and we have $\mathbb{E}_{t-1}[E_t]$ smaller than, but close to, 1. Conversely, p_t^- does not have power under increasing sequences. In this way we may construct a two-sided test via two applications of the test, the first against the increasing alternative and the second against the decreasing alternative.

In short, we are interested in testing the null H_0 against two non-parametric alternatives H^+ and H^- :

H_0 The sequence is i.i.d. Bernoulli with unknown mean.

H^+ The sequence is drawn from a Bernoulli with increasing mean.

H^- The sequence is drawn from a Bernoulli with decreasing mean.

In our context the observations are events that have occurred as an artifact of some agent's decision process. When we reject the null in favor of H^+ we say that the agent is *event-seeking*. When we reject the null in favor of H^- we say that the agent is *event-avoiding*.

Our question may be considered as a *changepoint detection problem* and our test as an *online changepoint detection* procedure. Given a sequence of observations $\{X_t\}_{t=1,2,\dots}$, changepoint detection aims to identify the time point, ν , at which the data-generating distribution changes from one given family of distributions, \mathcal{P} , to another given family of distributions \mathcal{Q} while controlling the probability of false detection (Shin, Ramdas, and Rinaldo 2022). Concretely we have,

$$X_0, \dots, X_\nu \sim P \in \mathcal{P} \tag{5.8}$$

$$X_{\nu+1}, \dots \sim Q \in \mathcal{Q} \tag{5.9}$$

It's important to note that in the most general setting \mathcal{P} and \mathcal{Q} are sets of distributions over *sequences*.

In the language of changepoint detection our question becomes: at what time does the sequence of observations start increasing? Our set of pre-change distributions \mathcal{P} is the set of i.i.d. Bernoulli sequences and our set of post-change distributions \mathcal{Q} is the set of distributions over sequences whose expectations are increasing (decreasing) in

time. Our statistic E_t will identify the point in time at which enough evidence has been gathered to say that the sequence is not i.i.d Bernoulli and may be increasing.

5.3 Data

5.3.1 Ride-sourcing driver GPS trajectories

We source our GPS data from the “rideshare and delivery assistant” app Gridwise¹. The Gridwise GPS data contains **183,548,522** individual GPS pings from **5352 drivers** in the Pittsburgh region for two **nine-month time periods: March to November of 2019 and March to November of 2020**.

Driver GPS trajectories are divided into shifts. A shift begins when a driver logs onto a service platform via the Gridwise app. A shift identifier is not recorded if the driver logged onto a service platform outside of the Gridwise app. We refer to pings labeled with a non-null shift identifier as *on-shift* pings. There are 23,982 distinct shifts in the data set which collectively account for two thirds of all pings. Although the median shift lasts approximately 5h30, a full 25% of shifts last at least 12 hours and 15% last more than 24 hours. We find it unlikely that drivers are truly driving over 12 hours at a time. We speculate that very long shifts result from drivers leaving the Gridwise app open on their device after they have finished using the TNC platform they initially logged into. Nevertheless we find short shifts to be a useful label for GPS pings as it provides some confidence that the driver was providing ride-hailing or delivery services for at least some of the time they were on-shift.

Table 5.1: Number and percent of GPS pings by activity type

activity type	n	%
unknown	5,697,293	3.10%
still	11,461,336	6.24%
on_foot	5,842,891	3.18%
walking	1,422,662	0.78%
running	93,625	0.05%
on_bicycle	286,545	0.16%
in_vehicle	158,744,169	86.49%

The GPS data is also labeled with an “activity type.” The activity type is a device-generated variable which reports the physical activity of the device inferred from its kinematic sensors. It is a categorical variable taking one of seven values summarized

¹<https://gridwise.io/>

in Table 5.1. Gridwise drivers spend most of their recorded time driving, a finding consistent with the use case for the Gridwise app. Of particular interest are the GPS pings classified as ‘still’. This set of pings contains the stops made by the driver and will form the basis of our subsequent analysis.

The GPS trajectories themselves are of high quality. Over 90% of pings are within 30 seconds of the most recent ping from the same driver and three-quarters are within 6 seconds. Despite the fact that drivers are logging on and off the platform over weeks if not months, the high temporal density of pings implies that the data contains long uninterrupted sequences of GPS pings. The geographic coordinates of half of all pings lie in a reported confidence radius of 5 meters (roughly 1 car length); 90% lie in a reported confidence radius of 16 meters.

The GPS trajectories are highly concentrated spatially. Of the US census tracts containing at least one GPS ping, the top 1% account for approximately 18% of all GPS pings. The tracts containing the Pittsburgh International Airport and Pittsburgh’s Downtown areas alone account for about 11% of all GPS pings.

5.3.2 OpenStreetMaps feature tags

OpenStreetMaps provides open source geographic data for the entire world, maintained by a community of users (OpenStreetMap contributors 2022). OpenStreetMaps represents physical entities in the world by associating one or more tags to a geographic shape. A tag is a ‘:’ delimited key-value pair of text data (e.g. `amenity:restaurant`). OpenStreetMaps uses a “free tagging system” meaning that map contributors are able to use any text they wish when tagging a geographic shape with data. However, informal standards have been adopted within the community around the most commonly used tags².

Our aim in using OpenStreetMap feature tags is to give context to the driver GPS trajectories and, ultimately, to find tags that were *relevant* to the driver’s decision-making process. To that end we focus on tags with the following keys:

- `aeroway`: airport-related tags.
- `amenity`: “useful and important facilities” for people.
- `building`: physical buildings and their associated use (e.g. residential, commercial, office).

In total, there are 31 tags of interest within the above tag keys.

We also leverage the `opening_hours` tag which reports the days and times when the entity is open. In particular this value is used to establish whether or not an entity was open at a particular moment (e.g. when a driver made a stop nearby).

²See https://wiki.openstreetmap.org/wiki/Map_features

5.4 TNC driver GPS analysis

There are several notable differences in the GPS data between 2019 and 2020 which we attribute to the onset of the COVID-19 pandemic. Our analysis suggests that in the months following the onset of the COVID-19 pandemic,

1. fewer, but more committed drivers took on an increased number of shorter shifts, and,
2. driver activity increased in the evening and decreased during the day.

An important caveat for the following analysis is that the sample of drivers represented on the Gridwise platform cannot be considered a random sample of TNC drivers. As a result, we cannot immediately use this analysis to conclude anything about the population of ride-hailing and delivery drivers. Drivers self-select to use Gridwise, presumably because they believe it will enable them to increase their earnings. Nevertheless, Gridwise data provides an unparalleled detailed look into the behavior of ride-hailing and delivery drivers.




Figure 5.1: Histogram of GPS pings by hour of day. Asterisks denote hours with statistically significant (Bonferroni corrected p-values at a 99% confidence level) difference between the histograms.

From March to November 2019, Gridwise saw a roughly 50% increase in monthly active drivers on the platform. After the onset of the COVID-19 pandemic, the platform steadily lost drivers. By November of 2020, the platform saw over 60% fewer drivers than it did in March of 2020. The decrease in monthly active drivers on the Gridwise platform is consistent with decreases in platform utilization of ride-hailing services

more broadly during this time period. In their second quarter filing with the U.S. Securities and Exchange Commission (SEC), Uber reported a 66% decrease in net revenue from mobility services and a 162% increase in net revenue from delivery services relative to the second quarter of 2019. Across all segments, Uber saw a 44% decline in monthly active customers and a 55% decrease in trips compared to the second quarter of 2019 (Uber Technologies 2020).

However, analysis of the monthly active drivers on Gridwise paints a different picture: instead of a sudden decline followed by slow recovery of drivers, Gridwise saw a steady decline in monthly active drivers over time. Drivers on Gridwise in 2020 had longer platform lifetimes than drivers in 2019 did, and spent more time on the platform on a per-driver basis. Measured as the number of days elapsed between the earliest and latest GPS pings associated with a driver id, average driver platform lifetime increased by over 30% from 2019 to 2020. Measured as the number of pings produced per driver, driver time on platform increased by nearly 30% from 2019 to 2020 and nearly 65% from July to November of 2020 compared to the same period in 2019. Taken together, we speculate that as the earning potential of being an TNC driver declined after the onset of COVID-19, the drivers that remained on the Gridwise platform were those who relied on TNCs for a large portion of their income. While drivers who had other options or only casually drove for TNCs simply left the platform, those that stayed doubled down, driving more shifts and longer and later hours.

Despite the decrease in drivers, the number of shifts on Gridwise remained largely the same, but shifts tended to be shorter with later start and end times. Drivers in 2020 tended to start their shifts later than their 2019 counterparts: comparatively more shifts started between 2pm and 10pm and fewer between 1am and 1pm. These shifts ended later too: the hours between 5pm and 3am saw an increase in the share of shifts ending from 2019 to 2020. In total, during the initial stages of COVID-19, drivers were more active in the afternoon and evening (11am-9pm) and less active in the late-night and mornings (9pm-11am) compared to 2020. Each day, defined as the 24 hour period starting at midnight local time, we computed a histogram of pings using 24 hour-aligned bins, each bin reporting the fraction of pings that occurred in each hour. We then applied a Bonferonni-corrected Welch’s t-test to determine the hours for which the average fraction of daily pings changed from 2019 to 2020. We report the average fraction of pings for 2019 and 2020 and significance at the 99% confidence level in Figure 5.1

5.4.1 OpenStreetMap analysis

What were drivers doing differently in 2020 compared to 2019? To address this question we join still GPS pings from the Gridwise data with OpenStreetMaps entity tag data. We define a ping “interaction” with a tag as the event that an entity with the

given tag was within 250 feet of the given ping. The interaction is a many-to-many relationship between entities and pings: each ping can have interactions with multiple entities (some having the same tag) and each entity may have interacted with multiple pings. We then sum the interactions for each tag over each hour and normalize by the total number of interactions in that hour, giving us a measure of collective driver interest in each tag by hour of day. In other words, the fraction of interactions answers the question: of all the entities nearby driver stops, what fraction of them had a given tag?




Figure 5.2: Average fraction of TNC driver interactions with tourism-related tags by month in 2019 and 2020. For comparison, Pittsburgh international airport passenger volumes (2019 volumes in solid black, 2020 volumes in dashed black) have been re-scaled and superimposed onto the airport interactions.

One of the more immediate impacts of COVID-19 on TNC ride-hailing operations was the reduction in airport trips due to the reduction in air travel. Airport trips represented 15% of Uber’s 2019 rides gross bookings (Uber Technologies 2019). In April 2020 Pittsburgh International Airport saw a 96% decrease in total passengers compared to April 2019 (Allegheny County Airport Authority 2020). We find that TNC drivers had dramatically fewer interactions with the airport and with hotels in 2020 as compared to the same month in 2019, both in absolute interactions and as a fraction of total interactions, as shown in Figure 5.2. Remarkably, the fraction of airport interactions tracks very closely with the passenger volumes (black dashed line in Figure 5.2) at the Pittsburgh airport from March to November of 2020.

A second noticeable effect of COVID-19 was that demand for TNC services shifted from passenger rides to delivery. In Uber’s 2020 second quarter (Q2) filing, passenger rides decreased substantially while trips on Uber’s food delivery platform soared

(Uber Technologies 2020). However, total activity, measured as both the number of trip (rides and delivery) and gross bookings, decreased. Many restaurants experienced a similar trend: an increase in delivery offset by a decrease in dine-in patrons. In their Q2 2020 filing, McDonald’s reported an 31% decrease in sales and a decrease in guest count compared to Q2 2019 even as delivery sales increased “significantly” (McDonald’s Corporation 2020). It stands to reason that any gains in TNC activity restaurants saw from Uber Eats were, on average, offset by even greater losses in TNC activity from passenger rides.

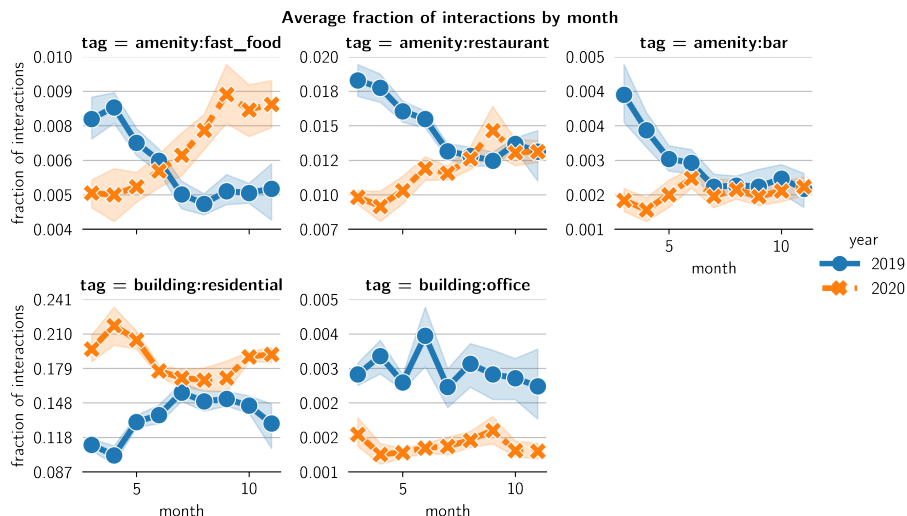


Figure 5.3: Average fraction of TNC driver interactions with passenger and delivery-related tags by month in 2019 and 2020

Can we demonstrate changes in tags that would be consistent with rides vs eats changes? Specifically we will focus on driver interactions with entities having the `amenity:fast_food`, `amenity:restaurant`, and `building:residential` tags. To ground our analysis we will also look at the `amenity:bar` and `building:office` tags. Under the ride-hailing prominent pre-COVID conditions we might expect bars and restaurants to have similar interactions with TNC drivers; likewise we might expect residences and offices to have similar interactions. Under post-covid conditions, when food delivery increased dramatically and passenger trips decreased dramatically, we should see diverging interaction trends between bars and restaurants and between residences and offices. Lastly, as Figure 5.1 demonstrates we expect to see relatively more activity in 2020 during the period from 11am until 9pm and relatively less activity during the period from 9pm until 11am. We term the periods the “late” and “early” periods respectively and will examine interaction in these two periods separately. Our findings are shown in Figure 5.3.

Overall, we find tag incidence data is consistent with the notion that COVID-19

produced a shift in consumer behavior away from travel and towards delivery. We find that the fraction of fast food and restaurant interactions increased from March to November 2020 despite a decreasing trend in those same months in 2019. During the late period, the fraction of fast food and restaurants interactions in 2020 met or exceeded the same-month fraction in 2019. In contrast, the fraction of bar interactions in 2020 was flat and below the 2019 same-month fraction. The fraction of residential interactions in 2020 comfortably exceeds the same-month fractions in 2019 for both the early and late periods. In contrast, the fraction of office interactions in 2020 remain decidedly below the same-month fractions in 2019 for both periods.

While it's tempting to read into these figures, it remains challenging to attribute these observations to a particular cause. As noted previously restaurant interactions are due a mix of delivery and passenger rides which we know to have changed over time. Similarly, residential interactions are a mix of home delivery and home rides. Since many more people were home for greater amounts of time, it is not clear how much of the difference in residential interactions is attributable to increased food delivery as opposed to more time spent at, and therefore traveling to or from, home.

5.5 Experiments

5.5.1 Simulated data

We demonstrate our e-process on simulated data. First, we consider sequences of i.i.d. Bernoulli data on which our e-process should fail to reject the null. Second, we consider sequences of rewards generated by agents interacting with a k -armed Bernoulli bandit where we expect our e-process should reject the null.

Stationary Bernoulli

We first apply our e-process to i.i.d Bernoulli data ($p = 0.8$). The values of the 10 e-processes are shown in Figure 5.4. As expected, none accumulates enough evidence for us to reject the null; here we use the increasing alternative to ease comparison with the following section. Though we fail to reject the null, the e-process does not stray too far from one; by comparison the likelihood ratio of the point alternative $p = 0.7$ against the (ground truth) point null of $p = 0.8$ yields a e-value around 10^{-8} after 1000 observations. In short, the exponential weighted mean is a poorer estimate of the true probability than the empirical mean, but not by a wide margin. This is helpful particularly in the changepoint application because if the data-generating distribution changed after $t = 1000$ it should not take long for the e-process to rebound and detect the change.




Figure 5.4: Our e-processes demonstrated on 10 different i.i.d Bernoulli sequences ($p = 0.8$)

Bernoulli bandit

We create a k -armed bandit ($k = 16$) with binary ($\{0, 1\}$) rewards whose arm means (p_i for $i = 1, \dots, 16$) are drawn i.i.d from a uniform distribution. Each time an agent pulls an arm it receives a reward of 1 with probability p_i . We examine the behavior of 10 agents each of which employs Thompson Sampling with a Beta(1, 1) prior (S. Agrawal and Goyal 2012). Each agent interacts with the same bandit instance for 1000 iterations. The behavior of the agents and their respective e-processes may be contrasted with the e-process for the i.i.d. Bernoulli sequences above. When an agent employs a fixed policy, the resulting reward sequence is i.i.d. Bernoulli and our test will fail to reject the null.

One e-process with an increasing alternative is generated for each of the reward sequences and shown in Figure 5.5. The legend lists the p-value derived from the test statistic next to its associated agent. For each of the agents, the e-process accumulates enough evidence to soundly reject the null. When we instead generate the e-process with the decreasing alternative, we fail to reject the null: the exponential weighted mean falls below the empirical mean and the alternative likelihood must use the empirical mean.

In short, this use case demonstrates the effectiveness of the test in establishing whether an agent *could possibly* have learned something about the environment they are interacting with, without knowledge of the action space or the environment.




Figure 5.5: Our e-process over bandit rounds, log scale. Black line indicates the 99.9% confidence level

5.5.2 Ride-hail driver GPS trajectories

We now apply our methodology to the GPS trajectories of ride-sourcing drivers. Specifically we analyze the on-shift GPS trajectories of drivers from March through November 2020. We apply the probabilistic model described in Section 5.2.1 to pings classified with the “still” activity type. We set the “break-even” speed, time, and distance ($\bar{v}, \bar{t}, \bar{d}$, respectively) to 8 miles per hour, 5 minutes, and 200 feet; ρ was set to the average measured speed over all trajectories: 14.48 meters per second (roughly 32 miles per hour). For the purposes of the test, we ignore all drivers who have fewer than 500 service stops; leaving 106 drivers. Each service stop is associated with a binary vector of length equal to the number of tags of interest as identified in Section 5.3.2. Non-zero entries in this vector represent the presence of at least one OpenStreetMaps entity with the given feature tag that is open at the time of the service stop and within 45 meters (roughly 150 feet) of the location of the service stop. For entities with no opening hours information, the entity was considered open if at least half of the entities having the same tag with opening hours information were open at that time. If no entity having the same tag was labeled with opening hours information, the entity was assumed to be open at all times.

We perform one two-sided test against the sequence of tag indicators over service stops for each tag and each driver, yielding 3,286 (106 drivers by 31 tags) individual tests. Within the exponential weighted mean function in the alternative likelihood we set the weight parameter γ as $1 - \exp(-\ln(2)/64)$, which ensures that the weight of each sample reduces by half every 64 service stops. This value was chosen as it seemed to provide a good balance of bias and variance in tests on the simulated data generated

in Section 5.5.1. As we previously noted, more complicated schemes could attempt to learn this value from data during the test.

The two-sided test combines two runs of the learning test, the first uses the increasing alternative and the second using the decreasing alternative. In effect, the first test supposes the agent wishes to maximize the tag incidence over time and the second supposes that the agent wishes to *minimize* the tag incidence over time. In our setting, when we reject the null using the increasing alternative we identify potential *tag-seeking* behavior and when we reject the null using the decreasing alternative we identify potential *tag-avoiding* behavior.

We also test the behavior of the drivers as a collective. To do so, we simply combine the service stops from all drivers into a single sequence of service stops in chronological order. To construct this sequence we use service stops from all 571 drivers resulting in 197,920 service stops. To account for the greater temporal density of service stops we set the weight parameter γ to produce a half-life of roughly one week. One two-sided test is applied to each tag from this single stream of service stops.

Bonferroni correction (Shaffer 1995) is applied to the p values returned by the test. A driver tag pair is deemed significant if its p value, multiplied by the number of driver-tag pairs tested, is less than $(1 - \alpha)/2$ for any confidence level $\alpha \in (0, 1)$. In this chapter we take $\alpha = 0.95$.

5.6 Results

5.6.1 Inferred service stops

Our inference process distills 2,552,460 still GPS pings from 648 drivers into 233,349 service stops from 571 drivers. On average service stops last about 12 minutes and contain 3 GPS pings. Service stops in our data range over 3235 square miles of Southwestern Pennsylvania intersecting six counties. The bulk of the activity, however, occurs in the Pittsburgh area and its containing county. A heatmap of all service stops is shown in Figure 5.6 whose extent has been narrowed for visual clarity. Service stops are heavily concentrated in the urban core of Pittsburgh, with substantial density across its inner suburbs.

5.6.2 Statistical hypothesis test results

At the 0.95 confidence level, after Bonferroni correction, we find 62 significant driver-tag pairs exhibiting tag-seeking behavior and 55 significant driver-tag pairs exhibiting tag-avoiding behavior out of the 3,286 driver-tag pairs tested.

Tables 5.2 and 5.3 report the Bonferroni-corrected p-values for tag-seeking and tag-avoiding behavior respectively. In both tables, the driver identifiers have been

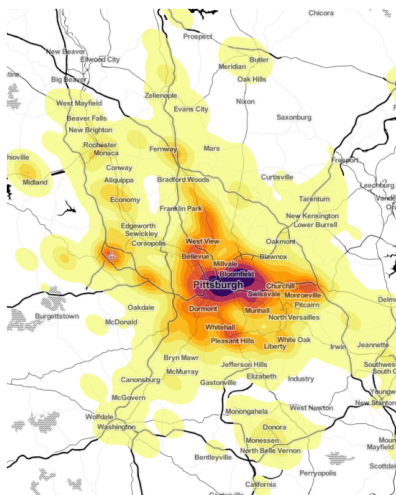


Figure 5.6: Kernel density map of service stops

elided and tags which appeared in more than one significant driver-tag pair are listed once alongside the minimum p-value for that tag over drivers and the number of significant driver-tag pairs the tag appears in (column n). An extended summary of test results is given in Table B.1 in B.1.

We now turn our attention to the tests performed on the sequence of service stops aggregated over all drivers. In this setting tag-seeking and tag-avoiding behavior should be taken to refer to the drivers as a collective. At the 0.95 confidence level, after Bonferroni correction, we find 13 significant tags exhibiting tag-seeking behavior and 8 significant tags exhibiting tag-avoiding behavior. Tables B.2 and B.3 in B.1 report the Bonferroni-corrected p-values for tag-seeking and tag-avoiding behavior respectively across the drivers as a collective as well as the timestamp at which the detection was made.

Notably, our procedure identifies the moment when enough evidence was gathered to conclude that the tag incidence probability was not constant in time and likely either increasing or decreasing. Conforming to our analysis of airport trips we find that airport trips were likely decreasing as early as late March and began slowly increasing again some time before early September. In contrast, food delivery related tags began to increase, starting with fast food and commercial areas in mid April followed by residences and restaurants in early May. It was comparatively longer before increases were seen in passenger oriented tags: theaters showed increases by June, casinos in August, bars in late October, and offices in November. Colleges and Universities showed increases by early September, which could be attributed to the beginning of the Fall semester, and highlights the difficulty in using this method to explain an increase in frequency rather than simply identify it.

Table 5.2: Bonferroni-corrected p-values for tags from driver-tag pairs exhibiting tag-seeking behavior. The final column (n) indicates the number of drivers for which the tag was found to be significant.

tag	p value	n
amenity:restaurant	9.43e-158	4
building:commercial	8.43e-127	8
amenity:college	5.32e-61	1
aeroway:aerodrome	1.30e-58	10
amenity:shelter	4.15e-38	5
building:residential	8.12e-38	17
amenity:place_of_worship	1.71e-15	1
amenity:grave_yard	5.06e-12	1
amenity:fast_food	8.07e-08	4
amenity:casino	1.75e-06	1
amenity:library	2.83e-06	1
amenity:university	4.29e-05	1
amenity:doctors	4.60e-03	1
amenity:dentist	7.41e-03	1

The p-values listed here are small due to the nature of the likelihood ratio under the alternative hypothesis. When the alternative hypothesis is true, the e process exhibits exponential growth in expectation as a result of the fact that it is an infimum over running products of terms each with expectation strictly greater than 1. Since the p-value is the inverse of the e-value, exponential growth in the e-value implies exponential shrinkage in the p-value resulting in very small p-values for some driver-tag pairs. It should be noted that not all p-values are small and moreover we report only the smallest p-values observed. The magnitude of the p-values can also be explained by the interpretation of this test as an online changepoint detection test. In online changepoint detection there is a fundamental tradeoff between accumulating evidence that a change has occurred and the time at which we declare that a change has occurred. If we want to detect a change as soon as the change is detectable, by necessity our p-values will be larger than if we had continued to collect evidence. In our setting, the p-values report the magnitude of the evidence gathered against the null hypothesis over the entirety of the study period, not just until we can detect a change. In many cases, the observations were drawn from the post-change distribution for most of the study period leading to very small p-values, for example, the airport tag under the alternative hypothesis that trip frequency is decreasing

The main benefit of this approach over heuristic approaches to changepoint detection is that by virtue of this being a statistical hypothesis test, the probability of a

Table 5.3: Bonferroni-corrected p-values for tags from driver-tag pairs exhibiting tag-avoiding behavior.

tag	p value	n
aeroway:aerodrome	2.94e-62	17
building:residential	1.11e-42	15
building:commercial	1.69e-32	5
amenity:restaurant	4.11e-17	4
amenity:shelter	6.46e-15	3
amenity:fast_food	1.44e-14	7
amenity:university	1.71e-10	1
amenity:casino	2.33e-08	1
amenity:cafe	5.40e-05	1
building:office	7.08e-03	1
amenity:hospital	8.10e-03	1

false null rejection is controlled at a pre-specified level. In this study for example, we have guaranteed that the probability of detecting at least one change when there was in fact no change is at most 5%. This property is particularly desirable in a practical setting where a detection of a change leads to some expensive intervention to the system. The desire to avoid wasting money on an intervention when no intervention is needed is exactly the error which statistical hypothesis testing controls for. However, the test does not control the probability of false retention, that is, not detecting a change when there in fact was a change. Controlling the probability of false retention is important when it is instead expensive to miss a change that has occurred. We can control for this type of error by swapping the null and alternative hypothesis as explained in Appendix B.2.

5.7 Discussion

When we reject the null for a driver tag pair, we are saying that the probability of observing a service stop near an open entity with that tag is unlikely to be constant in time. Further, by applying the two-sided test, we can identify driver-tag pairs in which the tag probability could possibly be increasing over time, and those for which the tag probability could possibly be decreasing over time.

Figures 5.7 and 5.8 demonstrate the running average tag incidence for the top 10 driver-tag pairs over time. The ‘reward signal’ in our case is a binary-valued variable representing whether or not the service stop was nearby an open entity with a given tag. As a result, the time-averaged reward is simply the frequency with which the tag




Figure 5.7: 60-day moving average tag incidence over time for the top 10 significant tag-seeking driver-tag pairs. Solid portions of the line begin when the null was rejected.

occurs nearby a service stop and is an empirical estimate of its probability. A driver-tag pair for which we reject the null in the tag-seeking test should exhibit increasing frequency in stops nearby the tag over time. Similarly, a driver-tag pair for which we reject the null in the tag-seeking test should exhibit decreasing frequency in stops nearby the tag over time. Accordingly, this behavior is observed for both the tag-seeking and tag-avoiding driver-tag pairs.

Our results suggest that our procedure is a highly effective behavioral trend detector by offering early trend detection with probabilistic guarantees. It can be observed in Figure 5.7 that the tag incidences (particularly `amenity:restaurant`, `building:commercial`, and `building:residential`) increase dramatically *after* we reject the null in favor of the tag-seeking alternative. In fact, most of the increase in tag incidence occurs after the null has been rejected. The `building:commercial` tag (orange in Figure 5.7) reaches a 60-day moving average low in mid-May at around 0.09 before rising over the next ten weeks to a high of 0.65 in early August. We first identify the increase in early June when we reject the null, meaning that the trend was identified within the first three weeks and first 10% of its eventual rise. Of course our test does not make any guarantees (probabilistic or otherwise) about what happens after we reject the null. Other tag incidences in Figure 5.7, for example `amenity:fast_food` and `amenity:shelter`, see only a modest increase in their frequency after we reject the null. However, the sensitivity of our procedure is tuned intuitively via the confidence level which may be directly interpreted as the probability of a false alarm. Performing a single test at the $1 - \alpha$ confidence level implies that a false rejection of the null (we identify a stationary sequence of tag-incidences as being tag-seeking) occurs with




Figure 5.8: 60-day moving average tag incidence over time for the top 10 significant tag-avoiding driver-tag pairs. Solid portions of the line begin when the null was rejected.

probability at most α . The Bonferroni correction applied to our experiments ($\alpha = 0.05$) further guarantees that the probability of any false rejections occur (over all tests) is at most 5%. For use cases where false alarms are costly, higher confidence levels may be used to require a greater level of evidence to reject the null. In these settings we should expect it to take longer and more of the trend to have been observed, for tag-seeking or tag-avoiding behavior to be identified. Conversely when early detection is more important and false alarms are tolerable, the confidence level may be lowered.

One caveat of the sequential nature of this hypothesis test is that it reports if *at any point* the sequence of tag frequency *up to that point in time* appears to have increased (decreased). A sequence of average rewards which initially rises and then falls may have risen enough initially to reject the null in the tag-seeking test, and if it falls enough, reject the null in the tag-avoiding test. In fact, this is evident in our test results: the `amenity:restaurant` is significantly increasing (blue in Figure 5.7) and also significantly decreasing (red in Figure 5.8). This same phenomenon can be observed with the `building:commercial` and `building:residential` tags.

When we inspect the e-processes for the `building:residential` driver tag pair in Figures 5.7 and 5.8 we gain some intuition about how the e-process accumulates evidence. In Figure 5.9 both the increasing alternative (seek) and decreasing alternative (avoid) e-processes are plotted in the lower panel. In the upper panel we show a moving average of the tag incidence which serves as a local estimate of the probability of stopping at a residential building. When the likelihood ratio favors the alternative over the null, the e-process will rise, indicating that it is accumulating evidence. The rate at which it rises is indicative of how much it favors the alternative. Conversely,

when the likelihood ratio favors the null over the alternative the e-process will fall. Our e-process will increase much faster under the alternative than it will decrease under the null and this is evident in the figure. Under the null, the alternative likelihood will fall back to the likelihood under the empirical mean in hindsight, which will result in likelihood ratios near one.




Figure 5.9: How the e-processes for increasing and decreasing sequences respectively accumulate evidence.

Statistically speaking, the rise and fall of e-processes is irrelevant; once the e-process crosses our significance threshold we are justified in rejecting the null, no matter what it does after. However, qualitatively it is instructive to relate the accumulation of evidence in the e-process with the increasing and decreasing probability in the sequence. In Figure 5.9 we illustrate this by qualitatively (via smoothing and finite differences) defining regions where the e-process is accumulating evidence. In the upper panel, timespans during which the e-process is accumulating evidence that the driver is avoiding residential buildings are shaded in red, while timespans during which the e-process is accumulating evidence that the driver is seeking out residential buildings are shaded in green. We emphasize that we are no longer using any statistical properties of the e-process in defining these regions so there is no statistical claim to be made regarding these regions. What this does show is that there is an intuitive alignment of the e-process and the observations: when the frequency with which we observe an event increases, the increasing alternative e-process rises and the decreasing alternative e-process stagnates or declines.

A second caveat is that the test itself cannot explain *why* the time-average rewards increase. When we are comfortable with the assumption that the environment (i.e. the action-conditioned reward distribution) is stationary, then, by process of elimination,

we must conclude that it was the agent, rather than the environment that changed behavior over time. In our setting, this assumption may not hold, as COVID-related restrictions shifted considerably over the lockdown period. So what we can say in interpreting the test results is that we can identify certain drivers who over the course of the lockdown changed their behavior to either seek out services near tags that they previously had not frequented or avoid services near other tags that they previously had. It might be the case that these drivers sought or avoided these entities because of a change in government restrictions, a change in some other aspect of the environment, or because of a change in how they valued these entities.

Perhaps the most conservative interpretation of these results is that statistical significance represents a shift in the underlying process generating the driver's behavior. There are a myriad of factors involved in producing an observation that a driver made a service stop at, say, a residential building: not only the drivers own decisions but the demand for TNC services, decisions of stores on the delivery platform, the platform decisions, government regulations, seasonal trends, and even traffic incidents or road work among others. Our results can be interpreted as gathering evidence that the process is now different, without specifying which component we believe to have been driving the change.

If we had more specific hypotheses (and additional data) however, this methodology gives us the tools to evaluate them. First, if we had more precise time series information about the specific entities in question, and in particular what COVID restrictions were applied to them at that time, we could use a conditional incidence as our sequence of observations, that is, for example, the event that a driver stopped nearby a restaurant *and* it was open for takeout. Another example would be the event that a driver stopped nearby a venue *and* there was an event ongoing at that time. At a technical level this is not different from how we used the opening hours information to condition our entity incidences.

Second, if we were to hypothesize that another binary metric was driving tag incidence, then if we applied our test to both, we should expect first that they are both significant with respect to the same alternative, and that the time at which the metric became significant precedes the time that the tag incident became significant. This application does not test the hypothesis directly and is most useful when there may be an unknown temporal lag in the relationship between the two. For example, we could hypothesize that the event that COVID hospitalization or case count were above a certain threshold would make take-out ordering more likely. In combination with the previous method, we would test the tag incidence of service stops nearby restaurants open for take out and also test the sequence of observations of COVID hospitalizations or case counts, hoping to see evidence for the latter accumulate to significance before the former.

Finally, it may be possible to directly test these hypotheses by modifying the e-

process appropriately. In this setting we wish to test against the null that the two metrics are independent where the alternative is that they are dependent, potentially constraining the estimates to specifying the direction of the relationship (e.g. for binary Y , $\mathbb{E}_{t-1}[X_t \mid Y_t = 1] \geq \mathbb{E}_{t-1}[X_t \mid Y_t = 0]$). In this setting we treat one of the two random variables as outside information, the likelihood under the null ignores this information while the alternative likelihood uses it:

$$E_t := \frac{\prod_{s=1}^t \left(\frac{n_{1|Y_s}^t}{n_{Y_s}} \right)^{X_t} \left(\frac{n_{0|Y_s}^t}{n_{Y_s}} \right)^{1-X_t}}{\binom{n_1^t}{t} \binom{n_0^t}{t}} \quad (5.10)$$

where $n_{1|Y_s}$ and $n_{0|Y_s}$ represent the number of times the variable X took on values of 1 and 0 respectively conditioned on the value of Y and $n_{Y_s} = n_{1|Y_s} + n_{0|Y_s}$. In the above e-process Y is assumed to be discrete. If it is not, then we can replace the empirical probabilities in the numerator with any estimator of $\mathbb{E}_{s-1}[X_s \mid Y_s]$.




Figure 5.10: Moving average tag incidence for significant seekers (green) and avoiders (red). Lines are dotted until they are found to be significant, then solid thereafter.

The results of our hypothesis test suggest that there was a great deal of heterogeneity in the driver behavior during the initial months of the COVID-19 pandemic. Several of the same tags are found to be both significantly increasing and also decreasing in frequency. Some of these significant sequences of observations belonged to the same driver, as in Figure 5.9. In these cases the sequence both increased and decreased at different points in time. However others, for example the residential building tag as shown in Figure 5.10, were found to be increasing for some drivers

while simultaneously decreasing for others. This suggests that the behavior changed differently across drivers. As before, this test does not explain *why* those differences exist: it could be that different drivers had different strategies or different responses to a changing environment or that a changing environment affected their different strategies in different ways, or even that they operate in different environments entirely.

We characterize the heterogeneity across drivers by fitting a Latent Dirichlet Allocation (LDA) model (Hoffman, Bach, and Blei 2010). LDA is often referred to as a topic model because it was developed to group a set of text documents into categories based on the words in the document. LDA simultaneously learns the categories or topics as a distribution over words (topic-word distribution), and the topic of the documents expressed as a distribution over topics (document-topic distribution). The documents are represented by a “document-term” matrix where the element at row i and column j is the number of times the word j appears in document i . In this setting we use LDA as a clustering method to identify behavioral profiles from the e-values of the drivers. We construct our document-term X as follows, for each driver i and tag i :

$$X_{i,j} = \ln \left(1 + \max_t E_t^{(i,j)} \right) \quad (5.11)$$

where $E_t^{(i,j)}$ is the value of our e-process for driver i and tag j at time t . We construct one document term matrix using the tag-seeking e-processes and another using the tag-avoiding e-processes and apply LDA to each separately. We take the maximum e-value to associate with each driver and tag the *largest* amount of evidence gathered in favor of the alternative over the observation period. Since e-values are multiplicative measures of evidence and the (pseudo-)count evidence that LDA considers are additive measures of evidence we apply a log transform to re-scale the e-values to conform with the assumptions of LDA. We set the document topic prior to 0.1 to encourage sparsity in the topic distribution assigned to each document, in our case this means that around 85% of drivers were assigned a topic distribution that placed at least 95% of its probability mass on a single topic. We assign component labels to each driver by taking the arg max of the document-topic distribution. The number of components (topics) was selected via cross-validation. We present the LDA results for tag-seeking behavior only.

The profiles extracted by LDA represent tags that are often sought together, but not necessarily at the same time. Component 2, which counts restaurants, fast food, commercial buildings, and residences as its top tags, appears to represent drivers who sought out food delivery. Component 0, on the other hand, which counts residences, the airport, commercial districts and colleges as its most likely tags, appears to represent drivers who sought passenger rides. Component 1 captures the majority of drivers and represents drivers who never accumulate much evidence against the null.

ever, is stronger in the aggregate analysis as we are able to pool stop data across drivers.

One particularly useful application of this test is to detect when an event begins to occur more frequently than it has in the past, without the need to specify how frequently we should expect this event to occur nor the process by which this event is generated. When applied to TNC operations in particular, observers can detect with confidence that driver behavior has changed soon after it changed. The value of this information depends on the observer and what they are trying to do. We imagine that a data scientist at a TNC is trying to establish when, for example, an effort to increase the share of drivers involved in delivery (measured as, for example, the frequency with which a completed trip was a delivery trip) has started to be effective. We imagine a regulator attempting to assess, for example, if and when a recent intervention at the curb is decreasing congestion or safety on the adjacent street. We imagine a data scientist at Gridwise who wants to detect when certain advice being offered to drivers is no longer as effective as it once was. Finally, we imagine that a TNC or government has implemented a policy under a set of assumptions about TNC driver behavior, perhaps for example, that drivers are more involved in rides than delivery. This methodology can be used to detect when these assumptions are no longer met.

In general we see two broad categories of use cases for this methodology. First, to establish with confidence and speed when some intervention begins to take effect in the desired way. Second, to establish with confidence and speed when an assumption that underlies a particular policy is no longer valid. Not only are both of these use cases powerful for regulators and organizations, but they are not methods to determine truth in hindsight, rather they are designed specifically to be monitored in real-time as data is being collected and to report with confidence as soon as the change is detected. Within choice modeling in particular, this methodology can be used to detect when the factors considered by a given choice model have a non-stationary relationship with the decisions observed over time.

In TNC operations, quickly detecting which kinds of entities drivers are increasingly seeking and avoiding can help identify congestion problems and policy interventions both within government and within TNCs themselves. From the government perspective, this information can be used to assess TNC operations in the city without delving too deeply into operational metrics. Despite the fact that the data we use here neither was provided by a TNC nor contains any TNC-specific information, we are able to learn quite a bit about TNC driver behavior. When TNCs are reluctant to provide detailed operational data, governments may yet find ways to leverage our procedure to understand TNC operations from proxy data, as we have done here. As it concerns our particular work, policymakers and regulators can use this method to understand how TNC trends are changing in their cities, potentially in response to government

regulation or changes in TNC policies or operations.

To illustrate how broadly this methodology may be applied we close this section by detailing an application to an important, but unrelated area of transportation: pedestrian safety at intersections (e.g. Vision Zero (Tingvall and Haworth 1999)). One important goal for planners and policymakers in this effort is to accurately estimate the probability of a pedestrian or bicycle collision at a particular intersection. For the sake of the example, say that every weekday during dusk hours we measure whether or not at least one pedestrian collision has occurred at a given intersection. To estimate the probability we can simply take the average of our observations, but this assumes that the underlying process generating the observations is stationary: the probability of a collision is not changing. As a result, we are also interested in whether or not the probability of a collision is increasing. When the probability is increasing our average is an *underestimate* of tomorrow's probability and we risk not acting out of a false sense of safety. If we can quickly detect that an intersection's collision probability has begun increasing, we can attempt to identify a cause and provide potentially life-saving intervention.

5.8 Conclusion

In this work we apply recent advances in sequential statistical hypothesis testing to extract behavioral changes in TNC driver behavior during the initial stages of the COVID-19 pandemic in the greater Pittsburgh area. We construct a composite null and composite non-parametric alternative sequential hypothesis test via an e-process capable of detecting when the probability of observing an event has increased or has decreased over a sequence of observations. Intuitively this e-process increases (accumulates evidence for the alternative) when the event starts occurring more frequently than it has in the past. The test is administered to the sequence of OpenStreetMap feature tags nearby service stops of 106 drivers over the first 9 months of the COVID-19 pandemic in Pittsburgh. The test found 112 driver tag pairs for which enough evidence existed to reject the null suggesting that the tag incidence probability was increasing (or decreasing) over time.

When a test finds that the sequence of tag incidences for a driver was becoming either significantly more or less frequent, it cannot tell us why. However, because the test places no assumptions on either the decision process in use by the driver or the environment, it can show us where to look for examples of changing behavior. Further analysis of these drivers could then yield insight into why certain changes were observed. Moreover, we find that certain tags became more frequent for some drivers while simultaneously becoming less frequent for others. This suggests heterogeneity in either or both of the driver decision processes or the environment that may be difficult to capture in a unified model. We describe this heterogeneity by fitting a Latent

Dirichlet Allocation model to the driver-tag e-values to cluster drivers into behavioral trend profiles. We recover food delivery and passenger rides profiles consistent with our earlier analysis. Lastly, the e-process itself can be leveraged to test more specific hypotheses, ones that, in particular, attempt to address why certain tags occur more or less frequently for certain drivers.

In this work we restrict our focus to binary metrics. Although the e-process can be directly extended to accommodate any discrete metric, a key avenue of future work is to extend the testing methodology to bounded metrics. Processes that are martingales under our null that the sequence has constant (conditional) mean are simple to construct. In the language of Shafer and Vovk (2019) (see also Waudby-Smith and Ramdas 2020) the challenge is to construct strategies to “bet against” the null so that under the alternative the e-process increases quickly.

A second key area of future work is a continued investigation into what was influencing the observed changes in behavior among these drivers. That the test does not explain the changes it identifies is on the one hand a strength of the methodology: because it does not rely on *any* assumptions on the behavior-generating process, failing to reject the null should be taken to mean that under *no model* did the *net* frequency of the event under study change significantly over time. However, from the practitioner’s perspective, it may be more interesting to focus on the drivers and tags for which we reject the null and for which the test offers no detailed explanation. In these cases the test identifies *what* needs to be explained and *when* the changes were first detected. Further investigations can thus be guided by these results.

5.9 Testing network equilibrium

We close this chapter by discussing in detail one of the more theoretically interesting applications of our procedure: assessing the possibility that aggregate flow on a network is approaching an equilibrium. The draw of equilibrium theory to practitioners and theorists alike lies in its ability to answer counter-factual questions about the future state of traffic on the network in a parsimonious and computationally efficient way. In Section 2.2 we noted that its validity is very difficult to assess and its underlying assumptions have been questioned by empirical studies. We also noted that since equilibrium models describe the limiting behavior, the most we should hope to establish is that the behavior of the network is tending toward it. In much the same way that our procedure may be applied to bandit agents to establish whether they appear to be earning over time but not that they *will necessarily* achieve the optimal policy, our test can be applied to transportation networks to establish that they appear to be progressing toward an equilibrium, but not that they *will necessarily* progress to any particular equilibrium.

Concretely, we aim to test the notion that a network may be approaching a user

equilibrium where the travel cost function is travel time. We assume we have access to a GPS trajectory dataset \mathcal{D} describing trips as well as segment level travel time estimates. We define a binary stochastic process parameterized by $\epsilon > 0$ and indexed by the time $t(r)$ at which a trip $r \in \mathcal{D}$ was completed.

$$X_{t(r)}^{(\epsilon)} = \mathbb{I}[T(r) < (1 + \epsilon)T^*(r)] \quad (5.12)$$

where $T(r)$ is the observed travel time of the trip r and $T^*(r)$ is the *best travel time in hindsight* for trip r and is defined as the travel time of quickest (time-dependent) route the traveler *could have* taken based on the road segment travel time estimates. Without loss of generality we assume that the values of $t(r)$ are distinct across all trips.

Each random variable in the binary stochastic process $X_{t(r)}^{(\epsilon)}$ represents the event that the traveler's realized travel time was within a small margin of the best travel time available to them; we refer to such trips as ϵ -*optimal*. At a perfect user equilibrium where the users cost function is travel time, the travel time of every trip will be equal and minimal; in other words, we have $\Pr[X_{t(r)}^{(\epsilon)} = 1] = 1$ for any trip r and for all $\epsilon > 0$. Conversely, for any ϵ , $\Pr[X_{t(r)}^{(\epsilon)} = 1] < 1$ implies that the network is not at equilibrium. If a network at disequilibrium, that is, a network for which $\Pr[X_{t(r)}^{(\epsilon)} = 1] < 1$ for some ϵ , is ever to achieve equilibrium then the probability that a trip is ϵ -optimal **must** increase over time. We can use our e-process (Equation (5.4)) with the increasing alternative as given by Equation (5.6), to directly test our sequence of observations against the null that ϵ -optimal trips are not becoming more frequent over time. When we fail to reject the null in favor we can say that there is not sufficient evidence to suggest that the network is trending toward equilibrium. When we instead reject the null we can say that the system may possibly be progressing toward an equilibrium. When we do reject the null we do not gain the authority to say that the system will eventually reach a given equilibrium state, only that it is behaving in a way that is consistent with a network that is progressing toward an equilibrium. As a result, our procedure is best used to eliminate cost functions from consideration for use in equilibrium models. In other words, to say that a cost function should be used to represent behavior in a network equilibrium model it should *at the very least* be deemed statistically significant by our procedure.

We demonstrate the application of our procedure using synthetic data generated by an equilibrium solution algorithm on the Braess Network. It is important to note that solution algorithms for equilibrium should not be interpreted as being representative of any real behavioral adjustment process that may be occurring on the road network. However, they do provide examples of processes that do converge to an equilibrium so it is suitable for use as a proof of concept. We solve for UE on the Braess network using the method of successive averages (MSA) (Boyles, Lownes, and Unnikrishnan 2021, Section 6.2.1) with a fixed step size as to extend the period during which the network


Algorithm 3 Observing an equilibrium adjustment process

```
1: Require: feasible link flow  $x$ ,  $\lambda > 0$ ,  $\epsilon > 0$ 
2:  $t \leftarrow \text{LinkCost}(x)$ 
3: for  $k \in (0, 1, \dots)$  do
4:    $x^* \leftarrow \text{AllOrNothing}(t)$ 
5:    $x \leftarrow \lambda x^* + (1 - \lambda)x$ 
6:    $t \leftarrow \text{LinkCost}(x)$ 
7:   Sample a path  $r$  with probability proportional to path volume
8:   yield  $\mathbb{I}[T(r) < (1 + \epsilon) \min_s T(s)]$ 
9: end for
```

is approaching equilibrium. Although this choice of step size does not converge to equilibrium, a small enough step size, like the one we will use, causes the algorithm to approach an equilibrium before oscillating near it. On each iteration of MSA, we sample a traveler at random and compute the value of Equation (5.12) taking $\epsilon = 0.01$. The full process is given in Algorithm 3 where `AllOrNothing` assigns all flow to shortest paths. In our demonstration the procedure is run for 1000 iterations and we take $\lambda = 1e - 3$.

As our experiments demonstrate, and as shown in Figure 5.12a, the frequency with which travelers experience near minimal travel time increases with time and as the relative gap decreases. Correspondingly, our e-process is able to reject the null in favor of the increasing alternative as shown in Figure 5.12b. What is clear from the true probability is that at no point do the travelers appear to choose quickest routes. It's not until the relative gap is within 3%, that is the total system travel time is within 3% of the best system travel time, that more than half of the travelers use ϵ -quickest routes; at a relative gap of 0.005% only 63% of travelers are using least cost paths. In short, even though our simulated travelers are genuinely trying to minimize their own travel time, this fact is not apparent from the data at any one point in time. In fact, it will *only* be apparent once the network reaches very near to its equilibrium state. In practice, however, we do not expect this traffic state to be achieved. In contrast, by considering the evolution of the system our procedure is able to demonstrate that the travelers as a whole are possibly improving their ability to minimize their travel time, even if they are as a collective failing to do so at any given moment.

In practice, testing equilibrium travel cost requires two ingredients: the evaluation of the travel cost function on a real observed trip trajectory through the network and the shortest travel cost available to the traveler. In most cases the former is simple to obtain since it is a function of our observation, and in this setting we may use any function of our data as a travel cost function. The latter, however, is usually not observable directly, but estimates of upper bounds may be available. An upper bound that is equally tight over time is valid in this application because it will simply



(a) Binary observations (blue), the true probability (orange) and exponential weighted average (green) over iterations of the modified MSA algorithm listed in Algorithm 3. The relative gap (red) demonstrates the progress to equilibrium.

(b) E-process using the increasing alternative applied to our observations of a system tending toward equilibrium.

Figure 5.12: Equilibrium testing procedure results.

shift the probability that a trip was completed within a small percentage of the least cost, but not alter the trend in the probability. In contrast, an upper bound that gets progressively tighter over time will make it appear, all else equal, that travelers are avoiding near least cost paths more frequently over time. The assumption that the estimation error is stationary over time is difficult to evaluate empirically, but depending on the data set, travel cost function, and the method used to construct the estimate the assumption may be reasonable.

We discuss a few practical ways to overcome the fact that we cannot observe the least travel cost directly. First, if the travel cost is link-separable and the dataset of trajectories is dense enough, it may be possible to estimate the travel cost on each link for each time window. These link level travel cost estimates can then be used to directly compute the least cost path through the graph to establish whether a lesser cost path might have existed at the time of the departure. However, if the coverage or density of the trip data on the network changes over time, it would be difficult to justify the assumption that the estimation error of the least cost path is stationary. Another possibility is to compute time-dependent link costs based on external data sources (e.g. traffic speeds from a third party provider) either instead of or in addition to the computations of travel cost from the trajectories themselves.

6

Conclusion

Technology in transportation mediates the relationships between the physical transportation systems and the social systems that use them in complex and opaque ways. This mediation is complex because of the large amount of data collected and the vast computing power leveraged by these technologies enables intelligent supply coordination in real-time and without human intervention. This mediation is opaque because the methodology, software, and data on which the automated decisions are based are not generally available to be analyzed directly. To understand technology in transportation purely in traditional terms at the intersection of social and physical systems is overly reductionist. If we, the planners, the engineers, and society at large, fail to understand the impacts of these technologies on their own terms and as a distinct system within transportation, we risk not only missing opportunities to enhance our transportation systems but risk exacerbating their negative impacts. However, precisely because technology in transportation participates in a “very close coupling” of a cyber-physical-social system, we are able to understand and hopefully influence the behavior of these technologies by examining how they relate to social-physical systems we can better measure. In this thesis we begin to bridge the gap between the analysis of social-physical and cyber systems operating within our transportation systems in three distinct but interrelated efforts that together are aimed at addressing three questions:

- What is the relationship between the cyber and social-physical components of a transportation system?
- What are the goals of decision makers who utilize a virtual layer to traverse the physical world?, and
- In light of these goals, what effect will our infrastructure decisions today have on the network of tomorrow?

Through these efforts we have learned that the automated decision-making in the

application of technology to transportation systems both affects the physical network and is affected by social-physical systems in a consistent way—one that can be learned from observations of the social-physical systems. In particular, the short-term evolution of surge-pricing in TNC systems can be understood as the result of social-physical factors interacting with both the TNCs algorithm and its vehicle supply.

That coordinating signals *affect* physical transportation systems like road networks, led us to our second effort. We demonstrated that when a platform coordinates transportation supply it can leverage its market power to reduce its total transportation cost. In some cases, this coordination can worsen congestion on the network while in others the coordination serves to ease congestion for all road users. The diverging potential impacts of these systems presents an opportunity to reshape our social-physical systems to mitigate harm and, in fact, promote equitable benefits of the application of technology to transportation.

That coordinating signals are *affected* by our social-physical systems led us to our third effort. We demonstrated that we may reason about the potential goals of highly complex decision making within technologically-driven transportation systems by statistical analysis of observed behavior, without relying on a particular model of the decision-maker. Moreover, by relying on sequential statistical hypothesis testing we can establish in real-time when the decision maker appears to begin changing their behavior. In particular we examined the surroundings of ride-hailing drivers when they made service related stops and were able to identify the drivers whose service stop behavior changed as well as the moment it changed. Finally, we demonstrate that we can leverage our procedure to justify the use of a particular cost function in a network equilibrium model based on empirical measurements of the network over time, thus allowing the practitioner to propose a notion of travel cost for an emerging transportation technology, empirically evaluate it, and then use it within a network equilibrium model to forecast the long-run impact of this technology on the road network.

In all, our treatment of the role of technology in transportation reveals that these systems have a profound impact on our transportation systems and warrant an understanding that acknowledges their idiosyncrasies. At the same time, these systems are profoundly affected by the social-physical elements of our transportation systems and we can leverage this interrelation to expose and influence their behavior in our societies.

6.1 Future work

Perhaps the thread that unites our efforts here is best defined by the work that remains. This thesis lays the foundation for justifying and evaluating concrete policy interventions aimed at equitably and effectively regulating cyber-physical-social sys-

tems in transportation. In broad terms, the aim of the future work is to build on this foundation in discovering, designing, and critiquing specific and realizable policy interventions, particularly as they relate to transportation network companies or autonomous vehicle fleets on public roadways.

The major theoretical contribution of this thesis is to demonstrate that social-physical models of transportation systems, particularly network equilibrium models, should acknowledge the new capabilities unlocked by the addition of cyber components. In Chapter 4 we demonstrate a notion of equilibrium which considers one such capability: vehicle coordination in service of total fleet travel time. Failing to account for this capability in planning models can result in *underestimates* of congestion; explicitly modeling it can lead to interventions that benefit the entire system. Other fleet level goals and behaviors are possible; some are more difficult to model than the one we investigated, some may more realistically capture a particular fleet of interest. When we are able to measure a virtual transportation system’s internal coordinating signal directly, we may leverage our efforts in Chapter 3 to understand what social-physical features are most tightly coupled with the coordinating signal. When machine learning methods are applied in this context we learn how the cyber component of the cyber-physical-social transportation system understands its relationship to the social-physical system with which it interacts. When internal measurements of the cyber component are not available the task ahead is to leverage our efforts in Chapter 5 to infer decision rules from data *as they are being learned*. Further, these two approaches may be combined to study cyber components whose coordinating signals are only partially measurable. Subsequently, our efforts in Chapter 4 serve a template for first examining the mixed equilibria these decision rules induce, and second highlight interventions capable of bending the mixed equilibria toward equitable and efficient network states. Moreover, applying our efforts Chapter 5 to network level metrics themselves we can ask what set of equilibria does the network appear to be tending towards.

Together, this thesis may be viewed as a roadmap for transportation planning in the face of complex technologies capable of influencing transportation network utilization in distinct and powerful ways. We close this thesis by imagining specific and important questions we view as the intellectual successors to the work presented here.

1. How can TNC fleets be leveraged to improve network efficiency?
 - (a) Can we infer the goals of TNC drivers?
 - (b) What social-physical features are TNC drivers sensitive to?
 - (c) How would government intervention alter TNC driver behavior?
 - (d) What is the mixed equilibrium that would be achieved if TNC drivers continued to follow their policies?
 - (e) What infrastructure changes would best align TNC behavior with trans-

portation planning goals?

2. How can we discover heterogeneous goals and their equilibria?
 - (a) Can we leverage our sequential hypothesis test to refute certain notions of equilibrium?
 - (b) Can we demonstrate potential travel goals or utility functions from observed traffic behavior, even as the system and the individuals are exploring and exploiting the network?
3. How does real-time navigation affect long-term traffic flow?
 - (a) How does sample bias in the measurements used to offer real-time navigation affect long-term network utilization?
 - (b) Are there platform-level goals real-time navigation providers have that differ from individual traveler goals?
4. What role does micro-mobility play in urban transportation?
 - (a) How do travelers learn to value micro-mobility as a mode of transportation?
 - (b) What social-physical features best predict micro-mobility utilization?
5. How can we effectively monitor and manage transportation infrastructure?
 - (a) Is infrastructure utilization changing?
 - (b) How long does it take an infrastructure intervention to produce changes in behavior?

A

On the impact of optimized fleets Appendices

A.1 Proof of Equation (4.3)

Proof. Let $T(\mathbf{x}) = \langle \mathbf{x}, t(\mathbf{x}) \rangle$ represent total delay at link flow \mathbf{x} , let \mathbf{y} represent the aggregate link flow at user equilibrium and let \mathbf{x} represent the aggregate link flow at mixed equilibrium as described by Harker (1988).

By convexity of T we have

$$\langle \nabla T(\mathbf{x}), (\mathbf{y} - \mathbf{x}) \rangle + T(\mathbf{x}) \leq T(\mathbf{y}) \quad (\text{A.1})$$

Therefore to show that $T(\mathbf{x}) \leq T(\mathbf{y})$, it suffices to show that $\langle \nabla T(\mathbf{x}), (\mathbf{y} - \mathbf{x}) \rangle \geq 0$.

$$\langle \nabla T(\mathbf{x}), (\mathbf{y} - \mathbf{x}) \rangle = \sum_{a \in A} (t_a(\mathbf{x}_a) + \mathbf{x}_a t'_a(\mathbf{x}_a)) (\mathbf{y}_a - \mathbf{x}_a) \quad (\text{A.2})$$

The aggregate link flows can be decomposed into feasible fleet link flows: $\mathbf{x}_a = \mathbf{x}_a^u + \sum_{i \in \mathcal{F}} \mathbf{x}_a^{c_i}$ and $\mathbf{y}_a = \mathbf{y}_a^u + \sum_{i \in \mathcal{F}} \mathbf{y}_a^{c_i}$ and then factored,

$$= \sum_{a \in A} (t_a(\mathbf{x}_a) + \mathbf{x}_a t'_a(\mathbf{x}_a)) \left((\mathbf{y}_a^u - \mathbf{x}_a^u) + \sum_{i \in \mathcal{F}} (\mathbf{y}_a^{c_i} - \mathbf{x}_a^{c_i}) \right) \quad (\text{A.3})$$

$$\begin{aligned} &= \sum_{a \in A} t_a(\mathbf{x}_a) (\mathbf{y}_a^u - \mathbf{x}_a^u) + \sum_{i \in \mathcal{F}} [(t_a(\mathbf{x}_a) + \mathbf{x}_a^{c_i} t'_a(\mathbf{x}_a)) (\mathbf{y}_a^{c_i} - \mathbf{x}_a^{c_i})] \\ &\quad + \mathbf{x}_a t'_a(\mathbf{x}_a) (\mathbf{y}_a - \mathbf{x}_a) - \sum_{i \in \mathcal{F}} \mathbf{x}_a^{c_i} t'_a(\mathbf{x}_a) (\mathbf{y}_a^{c_i} - \mathbf{x}_a^{c_i}) \end{aligned} \quad (\text{A.4})$$

Since $\mathbf{x}^u, \{\mathbf{x}^{c_i}\}_{i \in \mathcal{F}}$ are the mixed equilibrium traffic flows we have the following inequalities (Harker 1988)

$$0 \leq \sum_{a \in A} t_a(\mathbf{x}_a)(\mathbf{y}_a^u - \mathbf{x}_a^u) \quad (\text{A.5})$$

$$0 \leq \sum_{a \in A} (t_a(\mathbf{x}_a) + \mathbf{x}_a^{c_i} t'_a(\mathbf{x}_a))(\mathbf{y}_a^{c_i} - \mathbf{x}_a^{c_i}) \quad \forall i \in \mathcal{F} \quad (\text{A.6})$$

$$0 \leq \sum_{a \in A} t_a(\mathbf{x}_a)(\mathbf{y}_a^u - \mathbf{x}_a^u) + \sum_{i \in \mathcal{F}} [(t_a(\mathbf{x}_a) + \mathbf{x}_a^{c_i} t'_a(\mathbf{x}_a))(\mathbf{y}_a^{c_i} - \mathbf{x}_a^{c_i})] \quad (\text{A.7})$$

In particular, this means that if the following holds then the total delay under user equilibrium will exceed the total delay under mixed equilibrium.

$$\sum_{a \in A} \mathbf{x}_a^u t'_a(\mathbf{x}_a)(\mathbf{y}_a - \mathbf{x}_a) \geq \sum_{a \in A} \sum_{i \in \mathcal{F}} \mathbf{x}_a^{c_i} t'_a(\mathbf{x}_a)(\mathbf{y}_a^{c_i} - \mathbf{x}_a^{c_i}) \quad (\text{A.8})$$

When there is only one fleet, we can simplify

$$\sum_{a \in A} \mathbf{x}_a^u t'_a(\mathbf{x}_a)(\mathbf{y}_a - \mathbf{x}_a) \geq \sum_{a \in A} \mathbf{x}_a^c t'_a(\mathbf{x}_a)(\mathbf{x}_a^u - \mathbf{y}_a^u) \quad (\text{A.9})$$

□

A.2 System optimality paradox

We consider a two-node network with two parallel links and 5 units of demand. We define link cost functions for the two links as follows;

$$t_1(x_1) = 5 + x_1 \quad (\text{A.10})$$

$$t_2(x_2) = 1 + x_2 \quad (\text{A.11})$$

For completeness, the marginal costs are

$$\dot{t}_1(x_1) = 5 + 2x_1 \quad (\text{A.12})$$

$$\dot{t}_2(x_2) = 1 + 2x_2 \quad (\text{A.13})$$

We report the UE and SO solutions in Table A.1. First, we see that the UE solution equalizes travel cost (t) across the two links and the SO solution equalizes the marginal cost across the two links, thus verifying that the link flows are UE and SO solutions respectively. Second, we see that at SO, the 1.5 units of flow assigned to link 1 are worse off compared to the 3.5 units of flow assigned to link 2. We also see that at SO, the 1.5 units of flow assigned to link 1 are worse off compared to the cost they would have incurred at the UE flow. In total, users of link 1 are each absolutely worse off by 1 unit of travel cost, adding 1.5 to total cost, and users of link 2 are each absolutely better off by 1 unit of travel cost, reducing total cost by 3.5, for a net reduction in total cost of 2 relative to UE.

Table A.1: Demonstration of the SO paradox

	UE			SO		
	x	t	\dot{t}	x	t	\dot{t}
link 1	0.5	5.5	6.0	1.5	6.5	8.0
link 2	4.5	5.5	10.0	3.5	4.5	8.0
total	5.0	27.5		5.0	25.5	

A.3 Proofs of critical fleet size on a parallel network

A.3.1 Proof of Proposition 1

Proof. The reverse implication is true by construction.

To show the forward implication, let A_+ denote the set of links on which $\bar{x}_a^{\text{ue}} > 0$ and A_0 denote the set of links on which $\bar{x}_a^{\text{ue}} = 0$. By the user equilibrium principle we have that there exists a $t^* > 0$ such that $t_a(\bar{x}^{\text{ue}}) = t^*$ for all $a \in A_+$ and $t_a(\bar{x}^{\text{ue}}) \geq t^*$ for all $a \in A_0$.

At mixed equilibrium we also have that the fleet marginal cost \tilde{t}_a is equalized on all used paths. That is, there exists a $\tilde{t}^* > 0$ such that $\tilde{t}^* = t_a(\bar{x}^{\text{ue}}) + \mathbf{x}_a^c t'_a(\bar{x}^{\text{ue}})$ for all links a used by the fleet. Since we have assumed that the mixed equilibrium is non-trivial, we know that there is fleet flow on at least one link.

Suppose that there exists a link $a^* \in A_+$ with $\mathbf{x}_{a^*}^c = 0$. Then $\tilde{t}^* \leq t_{a^*}(\bar{x}^{\text{ue}}) + 0 \cdot t'_{a^*}(\bar{x}^{\text{ue}}) = t_{a^*}(\bar{x}^{\text{ue}})$. However, since $a^* \in A_+$ we have that $t_{a^*}(\bar{x}^{\text{ue}}) = t^*$. In short, fleet marginal cost is less than or equal to the user cost: $\tilde{t}^* \leq t^*$.

Now take any link on which $\mathbf{x}_a^c > 0$. It also must be the case that $\bar{x}_a^{\text{ue}} > 0$. Therefore both the link cost and the fleet marginal link cost must be minimal. So, we have $\tilde{t}^* = t^* + \mathbf{x}_a^c t'_a(\bar{x}^{\text{ue}})$. Since we have assumed the link cost function is increasing we have $t'_a > 0$ which implies that $\tilde{t}^* > t^*$. We have a contradiction. Therefore if $\bar{x}_a^{\text{ue}} > 0$ then $\mathbf{x}_a^c > 0$. \square

A.3.2 Proof of Proposition 2

Proof. Since we are at mixed equilibrium we have that the fleet marginal cost \tilde{t}_a is equal and minimal on all used links. By the previous proposition, we have that the set of links used by the fleet must be A_+ . So we have, $\tilde{t}^* = t^* + \mathbf{x}_a^c t'_a$ for all $a \in A_+$. By the flow conservation constraint we know that $\sum_{a \in A_+} \mathbf{x}_a^c = q_c$. Combining these equations

we have,

$$\frac{\tilde{t}^* - t^*}{\mathbf{t}'_a} = \mathbf{x}_a^c \quad \forall a \in A_+ \quad (\text{A.14})$$

$$\sum_{a \in A_+} \frac{\tilde{t}^* - t^*}{\mathbf{t}'_a} = \sum_{a \in A_+} \mathbf{x}_a^c = \mathbf{q}^c \quad (\text{A.15})$$

Rearranging we have,

$$\tilde{t}^* = \frac{\mathbf{q}^c}{\sum_{a \in A_+} \frac{1}{\mathbf{t}'_a}} + t^* \quad (\text{A.16})$$

However, our implicit constraint that $\mathbf{x}_a^c \leq \bar{\mathbf{x}}_a^{\text{ue}}$ can be violated by this solution in one of two ways. First, for all unused links we require that the marginal link cost, \tilde{t}^* , is less than or equal to the fleet marginal cost on those links. Since these links are also unused by the aggregate link flow, the (fleet) marginal cost is simply the link cost;

$$\tilde{t}^* \leq \mathbf{t}_a \quad \forall a \in A_0 \quad (\text{A.17})$$

Second, we could over-fill a link so that $\mathbf{x}_a^c > \bar{\mathbf{x}}_a^{\text{ue}}$. In which case the resulting fleet marginal cost would be greater than the marginal cost evaluated at the aggregate link flow, that is, we must have that

$$\tilde{t}^* \leq t^* + \bar{\mathbf{x}}_a^{\text{ue}} \mathbf{t}'_a \quad \forall a \in A_+ \quad (\text{A.18})$$

These conditions can be concisely expressed as,

$$\tilde{t}^* \leq \min_{a \in A} \mathbf{t}_a + \bar{\mathbf{x}}_a^{\text{ue}} \mathbf{t}'_a \quad (\text{A.19})$$

□

A.3.3 Proof of Proposition 3

Proof. By Proposition 2, we see that the fleet marginal link cost of each link at equilibrium is monotonic increasing in the fleet demand. Therefore we can maximize the fleet demand by binding the upper bound on the fleet marginal link cost. □

A.3.4 Proof of Proposition 4

We begin by first proving the following Lemma;

Lemma 1. *On the parallel network where mixed equilibrium is SO, exactly one of the following is true of the links that have positive aggregate link flow:*

1. *all links are used exclusively by the fleet;*

2. all links are used exclusively by the individual users; or
3. all links have a strictly positive user and fleet flow.

Proof. With \bar{x}^{so} representing the SO link flow and x_a^u, x_a^c representing the user and fleet flow on link a , we can identify four possible cases for each link summarized in Table A.2.

Table A.2: Cases for aggregate, user, and fleet link flow

case	\bar{x}_a^{so}	x_a^u	x_a^c
1	> 0	> 0	> 0
2	> 0	$= 0$	$= \bar{x}_a^{\text{so}}$
3	> 0	$= \bar{x}_a^{\text{so}}$	$= 0$
4	$= 0$	$= 0$	$= 0$

Since we are at mixed equilibrium there exist $t^* > 0$ such that $t^* = t_a$ for all $x_a^u > 0$ and $t^* \leq t_a$ otherwise. For the fleet there exists $\tilde{t}^* > 0$ such that $\tilde{t}^* = t_a + x_a^c t'_a$ if $x_a^c > 0$ and $\tilde{t}^* \leq t_a$ otherwise. The aggregate link flow is system optimal so there must exist $\pi > 0$ such that $\pi = t_a + \bar{x}_a^{\text{so}} t'_a$ if $\bar{x}_a^{\text{so}} > 0$ and $\pi \leq t_a + \bar{x}_a^{\text{so}} t'_a$ otherwise.

Consider a representative link a_1 which satisfies case 1. Since the link is used by both users and the fleet we have that $t^* = t_a$ and $\tilde{t}^* = t_a + x_a^c t'_a$. Since it is used at all we also have that $\pi = t_a + \bar{x}_a^{\text{so}} t'_a$. Together with the assumption that the link cost function is strictly increasing we have that $t^* < \tilde{t}^* < \pi$.

Now consider a representative a_2 from case 2. Since this link is used by the fleet exclusively we have $\tilde{t}^* = t_a + x_a^c t'_a = t_a + \bar{x}_a^{\text{so}} t'_a = \pi$. Immediately we see that cases 1 and 2 are incompatible: if there is a link used exclusively by the fleet then all links will be exclusively used by the fleet and if there is one link used by a mix of users and the fleet no other link is used exclusively by the fleet.

Now consider a representative a_3 from case 3. Since this link is used by users exclusively we have $t^* = t_a$ and $\tilde{t}^* \leq t_a$. Since it is used at all we also have that $\pi = t_a + \bar{x}_a^{\text{so}} t'_a$. Taken together we see that $\tilde{t}^* \leq t^* < \pi$ which implies that case 3 is incompatible with case 1 and 2: if one link is unused by the fleet then no links may be used by the fleet, and if any link is used by the fleet then no link is unused by the fleet.

Finally, links in case 4 are compatible with each of the other three cases because they are unused by users and fleets, which implies they are neither minimal cost, minimal fleet marginal cost, nor minimum marginal cost paths. As a result, they may be safely ignored. \square

We conclude with the proof of Proposition 4;

Proof. Suppose user equilibrium flow on the network is the same as the system optimal flow. Then the minimum marginal cost paths are also the minimum cost paths and users alone are sufficient to induce system optimal flow.

Alternatively suppose user equilibrium flow is not the same as system optimal flow; that is the set of used paths at system optimal is different from the set of used paths at user equilibrium. In particular, this implies that at system optimum some used paths are not least cost paths. At mixed equilibrium then, no users may be assigned to these paths. As a result, these paths will be used exclusively by the fleet. However, by Lemma 1, if one path is exclusively used by the fleet then all must be exclusively used by the fleet and therefore the fleet demand will be equal to the total demand.

It also may be the case that at system optimum some unused paths are least cost paths. In this case, the only way to ensure the path remains unused is when there are no users on the network, that is, all flow is fleet flow. \square

A.4 Critical fleet size heuristic solution algorithm via sensitivity analysis

In this section we develop a heuristic solution algorithm for the fleet-optimal minimum control ratio based on sensitivity analysis (Tobin and Friesz 1988). Recall the fleet equilibrium operator, FO defined in Program 4.12. Its first parameter is the user link flow \mathbf{x}^u which appears in the argument to the link cost function in the objective. Equivalently, it may be viewed as a parameter to the link cost function which, as long as $t_a(\cdot)$ is continuously differentiable in its argument, is also continuously differentiable in \mathbf{x}_a^u , a key requirement of sensitivity analysis. The second parameter of FO is simply the fleet demand .

We can reduce Program 4.13 by eliminating several variables and re-writing the constraints.

$$\min_{\mathbf{f}^u} \sum_{w \in W} \left(\bar{q}_w - \sum_{p \in P_w^*} \mathbf{f}_p^u \right) \quad (\text{A.20a})$$

$$\text{s.t. } \mathbf{f}^u \geq 0 \quad (\text{A.20b})$$

$$\mathbf{M}^* \mathbf{f}^u \leq \bar{\mathbf{q}} \quad (\text{A.20c})$$

$$\mathbf{D}^* \mathbf{f}^u + \mathbf{D} \text{FO}(\mathbf{D}^* \mathbf{f}^u, \bar{\mathbf{q}} - \mathbf{M}^* \mathbf{f}^u) = \bar{\mathbf{x}}^{\text{so}} \quad (\text{A.20d})$$

Where \mathbf{D}^* represents the incidence matrix between shortest paths and links at SO link flow and similarly, \mathbf{M}^* represents the incidence matrix between O-D pairs and shortest paths at SO link flow. In short, the reduced program removes definitional constraints and re-parameterizes everything in terms of the user path flows. This is not only more succinct, but also highlights that the user path flows, restricted to the

set of shortest paths at SO are all that need to be specified in order to determine all of the other values in the system. We also note that in this formulation the fleet equilibrium operator FO is now parameterized by the user path flow alone. In what follows, we define $\widetilde{\text{FO}}(\mathbf{f}^u) = \text{FO}(\mathbf{D}^*\mathbf{f}^u, \bar{\mathbf{q}} - \mathbf{M}^*\mathbf{f}^u)$.

Typically, sensitivity analysis is applied when the objective function is a function of the link flow at equilibrium in which case the approximate gradient found by sensitivity analysis is used directly in the computation of the descent (ascent) direction. In this case however, the objective is purely a function of the so-called perturbation parameters, not of the equilibrium these parameters induce. Instead, the induced equilibrium must obey link-wise constraints. As a result, we use an Augmented Lagrangian method (ALM) (Hestenes 1969) in combination with sensitivity analysis to develop a solution algorithm. ALM solves a constrained optimization problem by replacing the objective function with its Lagrangian plus an additional penalty term. The method then alternates updating the primal and dual variables, as well as the penalty parameter until convergence.

Here we are interested in treating the aggregate link flow constraint (A.20d) as a penalty in the objective. Because we have access to a local approximation of the gradient of the fleet optimal link flow via sensitivity analysis, we cannot enforce the equality constraint on the fleet link flow during descent: what appears to be feasible direction may not remain feasible after a step in that direction. When the constraint is instead penalized in the objective function such steps may be easily corrected in subsequent iterations.

Below we introduce the ALM formulation of Program A.20. First we define the function h such that $h(\mathbf{f}^u) = 0$ is exactly equivalent to Constraint A.20d: $h(\mathbf{f}^u) = \mathbf{D}^*\mathbf{f}^u + \mathbf{D}\widetilde{\text{FO}}(\mathbf{f}^u) - \bar{\mathbf{x}}^{\text{so}}$. Next we define the Augmented Lagrangian function: $\mathcal{L}_{\lambda,\rho}(\mathbf{f}^u) = -\mathbf{1}^T\mathbf{f}^u + \lambda^T h(\mathbf{f}^u) + \frac{\rho}{2}\|h(\mathbf{f}^u)\|^2$. The ALM is performed by the following iterative updates: In the k -th iteration, given a Lagrange multiplier estimate λ_{k-1} and a penalty weight ρ_{k-1} :

$$\mathbf{f}_k^u = \arg \min_{\mathbf{f}^u \in \mathcal{F}^*} \mathcal{L}_{\lambda_{k-1}, \rho_{k-1}}(\mathbf{f}^u) \quad (\text{A.21})$$

$$\lambda_k = \lambda_{k-1} + \rho_{k-1} h(\mathbf{f}_k^u) \quad (\text{A.22})$$

$$\rho_k = \alpha \rho_{k-1} \quad (\text{A.23})$$

Where $\alpha > 1$ and the feasible set \mathcal{F}^* is given by the remaining constraints from Program A.20, which form a convex polytope: $\mathcal{F}^* = \{\mathbf{f}^u \in \mathbb{R}^{|P^*|} \mid \mathbf{f}^u \geq 0, \mathbf{M}^*\mathbf{f}^u \leq \bar{\mathbf{q}}\}$.

The updates given by (A.22) and (A.23) are straightforward. However, the update given by (A.21) requires the minimization over the convex polytope \mathcal{F}^* of the augmented Lagrangian function $\mathcal{L}_{\lambda,\rho}$. Since we only have a locally-valid approximate gradient, the best we can do is to repeatedly take small steps in the direction of the local approximation of the negative gradient.

If we have a feasible user path flow at iteration k , denoted by $\mathbf{f}_k^u \in \mathcal{F}^*$ then each iteration we aim to find a vector $\mathbf{v}_k \in \mathbb{R}^{|P^*|}$ such that $\mathbf{f}_{k+1}^u = \mathbf{f}_k^u + \mathbf{v}_k$ is feasible and decreases the objective function $\mathcal{L}_{\lambda_k, \rho_k}$. Plugging \mathbf{f}_{k+1}^u into the constraints we have,

$$\mathbf{f}_k^u + \mathbf{v}_k \geq \mathbf{0} \quad (\text{A.24})$$

$$\mathbf{M}^* \mathbf{v}_k \leq \bar{\mathbf{q}} - \mathbf{M}^* \mathbf{f}_k^u \quad (\text{A.25})$$

A \mathbf{v}_k satisfying (A.24) and (A.25) will ensure that \mathbf{f}_{k+1}^u is in a feasible direction relative to \mathbf{f}_k^u . It is also a descent direction if $\mathcal{L}_{\lambda_k, \rho_k}(\mathbf{f}_k^u + \mathbf{v}_k) \leq \mathcal{L}_{\lambda_k, \rho_k}(\mathbf{f}_k^u)$. Because we neither have a closed form for $\mathcal{L}_{\lambda_k, \rho_k}$ nor its gradient, we cannot solve directly for \mathbf{v}_k , we have to rely on the linear approximation centered at \mathbf{f}_k^u $\mathcal{L}_{\lambda_k, \rho_k}(\mathbf{f}_k^u + \mathbf{v}_k) \approx \mathcal{L}_{\lambda_k, \rho_k}(\mathbf{f}_k^u) + \nabla \mathcal{L}_{\lambda_k, \rho_k}(\mathbf{f}_k^u)^T \mathbf{v}_k$. Using this approximation, \mathbf{v}_k is a descent direction if, $\nabla \mathcal{L}_{\lambda_k, \rho_k}(\mathbf{f}_k^u)^T \mathbf{v}_k \leq 0$. The gradient of the augmented Lagrangian may be expressed in terms of the Jacobian of the equilibrium fleet link flow found by sensitivity analysis, $\mathbf{J}[\mathbf{DF}\tilde{\mathbf{O}}(\mathbf{f}_k^u)]$; $\nabla \mathcal{L}_{\lambda_k, \rho_k}(\mathbf{f}_k^u) \approx -\mathbf{1}^T + (\lambda_k^T + \rho_k \|h(\mathbf{f}_k^u)\|)(\mathbf{D}^* + \mathbf{J}[\mathbf{DF}\tilde{\mathbf{O}}(\mathbf{f}_k^u)])$. We may then find the best feasible direction at iteration k by solving the following linear program,

$$\min_{\mathbf{v}} \nabla \mathcal{L}_{\lambda_k, \rho_k}(\mathbf{f}_k^u)^T \mathbf{v} \quad (\text{A.26a})$$

$$\text{s.t. } \mathbf{f}_k^u + \mathbf{v} \geq \mathbf{0} \quad (\text{A.26b})$$

$$\mathbf{M}^* \mathbf{v} \leq \bar{\mathbf{q}} - \mathbf{M}^* \mathbf{f}_k^u \quad (\text{A.26c})$$

A step size $\gamma \in [0, 1]$ is then selected and an update $\mathbf{f}_{k+1} = \mathbf{f}_k + \gamma \mathbf{v}_k$ can be performed. In practice γ should be small since sensitivity analysis provides only a local estimate of the gradient. Because Program A.26 is linear, the approach described above is exactly a Sequential Linear Programming (SLP) (Wright, Nocedal, et al. 1999) solution method. We now present the algorithm to complete the primal update given by (A.21) in Algorithm 4. We introduce the full ALM solution method in Algorithm 5.

Algorithm 5 has a few limitations. First, the main theoretical limitation is that since the fleet optimal link flow is not a convex function of the user path flow, the algorithm is guaranteed only to return a local optimum, not necessarily the globally optimal solution. The main practical limitation is that the algorithm requires the computation of equilibrium and sensitivity analysis every iteration of the primal update which itself runs many times within the solution algorithm. Moreover, the aggregate link flow constraint imposes a specific structure on the equilibrium condition which the equilibrium computation within each inner iteration has no knowledge of. We may instead leverage this structure to propose a heuristic that will return an upper bound on the control ratio, and a mixed integer program that is guaranteed to return the globally optimal solution.

Algorithm 4 Primal Update via Sequential Linear Programming

```
1: procedure PrimalUpdate( $\lambda, \rho$ )
2:    $\mathbf{f}^u \leftarrow \mathbf{0} \in \mathbb{R}^{|P^*|}$   $\triangleright \mathbf{f}^u = \mathbf{0}$  is always a feasible solution.
3:   repeat
4:     Compute  $\widetilde{\text{FO}}(\mathbf{f}^u)$ 
5:     Estimate  $J[\text{D}\widetilde{\text{FO}}(\mathbf{f}^u)]$  via sensitivity analysis
6:      $\mathbf{v} \leftarrow \text{Solve (A.26)}$ 
7:     Pick  $\gamma \in [0, 1]$ 
8:      $\mathbf{f}^u \leftarrow \mathbf{f}^u - \gamma \mathbf{v}$ 
9:   until converged
10:   $\mathbf{f}^c \leftarrow \widetilde{\text{FO}}(\mathbf{f}^u)$ 
11:  return  $\mathbf{f}^u, \mathbf{f}^c$ 
12: end procedure
```

Algorithm 5 Critical Fleet Size via ALM

```
1: procedure CFS( $\alpha$ )
2:   Initialize  $\lambda, \rho$ 
3:   repeat
4:      $\mathbf{f}^u, \mathbf{f}^c \leftarrow \text{PrimalUpdate}(\lambda, \rho)$ 
5:      $\lambda \leftarrow \lambda + \rho h(\mathbf{f}^u)$ 
6:      $\rho \leftarrow \alpha \rho$ 
7:   until converged
8:   return  $\mathbf{f}^u, \mathbf{f}^c$ 
9: end procedure
```

A.5 Proof of the MCR lower bound of CFS-SO

Proof. The Minimum Control Ratio (MCR) is defined by Z. Chen et al. (2020) and Sharon et al. (2018) to be the smallest fleet of system-optimizing vehicles (where all other vehicles are user-optimizing) which achieves system optimal traffic assignment on the network. Both works phrase the MCR as a linear program, reproduced below using our notation;

$$\min_{\mathbf{f}^u \geq 0, \mathbf{f}^c \geq 0} \sum_{p \in \tilde{P}} \mathbf{f}_p^c \quad (\text{A.27a})$$

$$\text{s.t. } \mathbf{D}^* \mathbf{f}^u + \tilde{\mathbf{D}} \mathbf{f}^c = \bar{\mathbf{x}}^{\text{so}} \quad (\text{A.27b})$$

$$\mathbf{M}^* \mathbf{f}^u + \tilde{\mathbf{M}} \mathbf{f}^c = \bar{\mathbf{q}}^{\text{so}} \quad (\text{A.27c})$$

where \mathbf{D}^* , $\tilde{\mathbf{D}}$, \mathbf{M}^* , and $\tilde{\mathbf{M}}$ are defined as above. The constraints of this program ensure that;

1. users, whose path flow is denoted by \mathbf{f}^u , only use least cost paths,
2. fleets, whose path flow is denoted by \mathbf{f}^c , only use least marginal cost paths,
3. the aggregate link flow is system optimal, and,
4. the original origin-destination demand is split between the two classes.

Immediately, we see that the constraint set of Program A.27 contains the constraint set of Program 4.17. As proved later, Program 4.17 represents an upper bound on Critical Fleet Size when we take \tilde{P} (the set of fleet-usable paths) to be the set of minimum marginal cost paths in Constraints (4.17b) and (4.17c). This implies that the *upper bound* on Critical Fleet Size is at least as large as the Minimum Control Ratio.

However, because fleets may only use minimum marginal cost paths, there must exist a path set $P' \subseteq \tilde{P}$ which represents the set of fleet paths at the mixed equilibrium at aggregate system optimal link flow. Although this path set cannot be known in advance, if we were to use it instead of \tilde{P} in Program 4.17, we would solve not for an upper bound on Critical Fleet Size, but for Critical Fleet Size itself. The replacement of \tilde{P} with a subset of it, P' is realized by adding yet more constraints to the MCR formulation. As a result, the Critical Fleet Size is at least as large as the Minimum Control Ratio. \square

A.6 Proof of Proposition 5

We first consider CFS-SO. Because the aggregate link flow is SO, it follows that any path r with a strictly positive aggregate flow $((\mathbf{f}^u + \mathbf{f}^c)_r > 0)$ must be a least marginal cost for the O-D pair it connects. Conversely, if r is not a least marginal cost path

then $(f^u + f^c)_r = 0$ must hold. This implies that at a mixed equilibrium that is also system optimum, fleets may only use least marginal cost paths, that is if P' represents the set of paths in use by the fleet at CFS-SO, then it must be the case that $P' \subseteq \tilde{P}$. Constraints 4.17b and 4.17c ensure that fleet marginal cost of any path that *could* be used by the fleet is equal and minimal. However, if we knew in advance the true path set P' , we could compute the true CFS-SO by replacing \tilde{P} in 4.17b with P' . Because $P' \subseteq \tilde{P}$, we see that any solution to the original program is also a solution to the modified linear program. Since the program is a minimization, the solution to the original linear program is therefore at least as large as the modified program, in other words the heuristic linear program 4.17 provides an upper bound on CFS-SO.

For CFS-UE, the aggregate link flow is UE. It follows that any path r with a strictly positive aggregate flow $((f^u + f^c)_r > 0)$ must be least cost for the O-D pair it traverses. Conversely, if r is not a least marginal cost path then $(f^u + f^c)_r = 0$ must hold. Analogously to the above, at a mixed equilibrium that is also UE, fleets may only use least cost paths, that is if P' represents the set of paths in use by the fleet at CFS-UE, then it must be the case that $P' \subseteq P^*$. Again, if we knew in advance the true path set P' we could compute the true CFS-SO by replacing P^* in 4.18b with P' . Because $P' \subseteq P^*$, we see that any solution to the original program is also a solution to the modified linear program. Since the program is a maximization, the solution to the original linear program is no larger than the modified program, in other words, the heuristic linear program 4.18 is a lower bound on CFS-UE.

This completes the proof.

At an intuitive level, solving our linear program can be thought of in analogy to the following. We start with $f^u = 0$ so all vehicles are fleet vehicles and λ_w is simply the minimum marginal cost over the paths between O-D pair w . As flow is drained to individual users, λ_w will decrease. When at least one of the paths is completely drained, that is $f_r^c = 0$ (with $r \in \tilde{P}$), the minimum fleet marginal fleet cost may not be further reduced, because path r , despite having zero flow, is still contained in the set of used paths. Hence, the linear program returns an upper bound on the minimum fleet size to achieve SO link flow. Conversely, the linear program for CFS-UE returns a lower bound on CFS-UE.

A.7 Critical fleet size with multiple fleets

Here we identify in Program A.28 the formulation of critical fleet size for system optimum in a network with individual users and $m > 1$ fleets. An analogous program can be constructed for CFS-UE in the case of multiple fleets, and by extension both admit an LP which offers an upper and lower bound for CFS-SO and CFS-UE with multiple fleets respectively. We arrive at Program A.28 by adding flow and cost variables (and associated constraints) for each fleet to the MIP given by Program 4.15. The link flow

and demand conservation constraints are also modified accordingly. Let $\{\mathbf{f}^{\mathbf{c}_1}, \dots, \mathbf{f}^{\mathbf{c}_m}\}$ denote the path flow vector for each of the m fleets in the system, let λ_w^i denote the least fleet marginal path cost between O-D pair w for fleet \mathbf{c}_i , and let $z_k^i \in \{0, 1\}$ denote whether path k is in use by fleet \mathbf{c}_i .

$$\min_{\mathbf{f}^{\mathbf{u}} \geq 0, \mathbf{f}^{\mathbf{c}_i} \geq 0, \lambda^{\mathbf{c}_i}, z^{\mathbf{c}_i}} \sum_{i=1}^m \sum_{p \in \tilde{P}} \mathbf{f}_p^{\mathbf{c}_i} \quad (\text{A.28a})$$

$$\text{s.t. } \tilde{\mathbf{c}}_r(\mathbf{f}^{\mathbf{c}_i}) \geq \lambda_w^i \quad \forall r \in P_w \forall w \in W \forall i \in \{1, \dots, m\} \quad (\text{A.28b})$$

$$\tilde{\mathbf{c}}_r(\mathbf{f}^{\mathbf{c}_i}) \leq \lambda_w^i + m_1 \cdot (1 - z_r^i) \quad \forall r \in \tilde{P}_w \forall w \in W \forall i \in \{1, \dots, m\} \quad (\text{A.28c})$$

$$\mathbf{f}_r^{\mathbf{c}_i} \leq m_2 \cdot z_r^i \quad \forall r \in \tilde{P}_w \forall w \in W \forall i \in \{1, \dots, m\} \quad (\text{A.28d})$$

$$\mathbf{D}^* \mathbf{f}^{\mathbf{u}} + \tilde{\mathbf{D}} \sum_{i=1}^m \mathbf{f}^{\mathbf{c}_i} = \bar{\mathbf{x}}^{\text{so}} \quad (\text{A.28e})$$

$$\mathbf{M}^* \mathbf{f}^{\mathbf{u}} + \tilde{\mathbf{M}} \sum_{i=1}^m \mathbf{f}^{\mathbf{c}_i} = \bar{\mathbf{q}}^{\text{so}} \quad (\text{A.28f})$$

$$z_k^i \in \{0, 1\} \quad \forall k \in \tilde{P}_w \forall w \in W \forall i \in \{1, \dots, m\} \quad (\text{A.28g})$$

where $\tilde{\mathbf{c}}_r$ is the fleet marginal cost for path r defined as $\tilde{\mathbf{c}}(\mathbf{f}^{\mathbf{c}_i}) = \mathbf{c} + \mathbf{D}^T \text{diag}(\mathbf{t}') \mathbf{D} \mathbf{f}^{\mathbf{c}_i}$. We note that any feasible solution to Program 4.15 is trivially a feasible solution to Program A.28 so the single fleet notion of CFS-SO is an upper bound of CFS-SO for multiple fleets. Similarly, the single fleet notion of CFS-UE is a lower bound of CFS-UE for multiple fleets.

A.8 Path set generation

In this section we describe one way, used in our numerical experiments, to generate the path sets used in the CFS-SO and CFS-UE programs.

Path set generation contains two steps: first, each path generated during the network equilibrium solution algorithm is stored in a database, second, before the CFS program is formed, the database is queried to find all least cost and least marginal cost paths at the given link flow solution.

The Frank-Wolfe (FW) (Boyles, Lownes, and Unnikrishnan 2021) network equilibrium solution algorithm was used in Chapter 4 to compute both the user equilibrium and system optimal solutions on each of our test networks. Each iteration of FW performs two steps, first it computes the link cost based on the link flow produced by the previous iteration and second it re-distributes link flow based on link costs. In order to update link flow, FW performs a weighted shortest path search between each OD pair using the link costs computed in the first step and loads all of the flow onto these paths, producing what is termed the “all-or-nothing” (AON) link flow assignment. The

link flow produced by the iteration is a convex combination of the link flow from the previous step and the AON link flow. As a result, the link flow can be equivalently expressed as a linear combination of path flows over the set of paths that have been seen by FW. Each iteration of FW, the set of paths computed by the weighted shortest path algorithm are stored in a LevelDB database (Ghemawat and Dean 2014), a fast file-based key-value store. As a key-value store, LevelDB is optimized for fast key retrieval; by using the sequence of link ids as the key (and the OD pair id as the value) we can store exactly one copy of each path found for a low computational cost.

When we want to build a path set for a CFS mathematical program, we iterate through the paths in the database, identifying those that are either least cost or least marginal cost. Because the database contains the link ids and OD pair id for each path, the (marginal) cost of each path can be determined from a vector of link costs and the link-path incidence matrix \mathbf{D} and the OD-path incidence matrix \mathbf{M} can be formed.

A.9 First order optimality conditions for the critical fleet size upper bound linear program

In this section we examine the Karush-Kuhn-Tucker (KKT) optimality conditions of the linear program for the upper bound on critical fleet size given by Program 4.17. To make the KKT conditions more clear we re-write Program 4.17 to express the fleet marginal cost in terms of the fleet path flow and expand the constraints.

$$\min_{\mathbf{f}^u, \mathbf{f}^c, \lambda} \sum_{p \in P} \mathbf{f}_p^c \quad (\text{A.29a})$$

$$\text{s.t. } \lambda_w - \mathbf{c}_r - \sum_{a \in A} \sum_{s \in P} \delta_{ar} \delta_{as} \mathbf{f}_s^c \mathbf{t}'_a = 0 \quad \forall r \in \tilde{P}_w \forall w \in W \quad (\text{A.29b})$$

$$\lambda_w - \mathbf{c}_r - \sum_{a \in A} \sum_{s \in P} \delta_{ar} \delta_{as} \mathbf{f}_s^c \mathbf{t}'_a \leq 0 \quad \forall r \in P_w \setminus \tilde{P}_w \forall w \in W \quad (\text{A.29c})$$

$$\sum_{r \in P^*} \delta_{ar} \mathbf{f}_r^u + \sum_{r \in \tilde{P}} \delta_{ar} \mathbf{f}_r^c - \bar{\mathbf{x}}_a^{\text{so}} = 0 \quad \forall a \in A \quad (\text{A.29d})$$

$$\sum_{r \in P_w^*} \mathbf{f}_r^u + \sum_{r \in \tilde{P}_w} \mathbf{f}_r^c - \bar{\mathbf{q}}_w^{\text{so}} = 0 \quad \forall w \in W \quad (\text{A.29e})$$

$$-\mathbf{f}_r^u \leq 0 \quad \forall r \in P \quad (\text{A.29f})$$

$$-\mathbf{f}_r^c \leq 0 \quad \forall r \in P \quad (\text{A.29g})$$

We use α to denote the KKT multipliers associated with the constraints. Superscripts will be used to identify the multiplier with the associated constraint group,

e.g. $\alpha^{(b)}$ for constraint (A.29b); subscripts will be used to index the multiplier within the constraint, e.g. $\alpha_r^{(b)}$ for the multiplier for path r within constraint group b .

Stationarity conditions yield,

$$1 - \sum_{r \in \tilde{P}} \alpha_r^{(b)} \sum_{a \in A} \delta_{ar} \delta_{ap} \mathbf{t}'_a - \sum_{r \notin \tilde{P}} \alpha_r^{(c)} \sum_{a \in A} \delta_{ar} \delta_{ap} \mathbf{t}'_a + \sum_{a \in A} \alpha_a^{(d)} \delta_{ap} + \alpha_{w(p)}^{(e)} - \alpha_p^{(g)} = 0 \quad \forall p \in \tilde{P} \quad (\text{A.30})$$

$$\sum_{a \in A} \alpha_a^{(d)} \delta_{ap} + \alpha_{w(p)}^{(e)} - \alpha_p^{(f)} = 0 \quad \forall p \in P^* \quad (\text{A.31})$$

$$\sum_{r \in \tilde{P}_w} \alpha_r^{(b)} + \sum_{r \in P_w \setminus \tilde{P}_w} \alpha_r^{(c)} = 0 \quad \forall w \in W \quad (\text{A.32})$$

By dual feasibility we have,

$$\alpha_r^{(c)} \geq 0 \quad \forall r \in P \quad (\text{A.33})$$

$$\alpha_r^{(f)} \geq 0 \quad \forall r \in P^* \quad (\text{A.34})$$

$$\alpha_r^{(g)} \geq 0 \quad \forall r \in \tilde{P} \quad (\text{A.35})$$

By complementary slackness,

$$\alpha_r^{(c)} \left(\lambda_w - \mathbf{c}_r - \sum_{a \in A} \sum_{s \in P} \delta_{ar} \delta_{as} \mathbf{f}_s^{\mathbf{c}} \mathbf{t}'_a \right) = 0 \quad \forall r \notin \tilde{P} \quad (\text{A.36})$$

$$-\alpha_r^{(f)} \mathbf{f}_r^{\mathbf{u}} = 0 \quad \forall r \in P^* \quad (\text{A.37})$$

$$-\alpha_r^{(g)} \mathbf{f}_r^{\mathbf{c}} = 0 \quad \forall r \in \tilde{P} \quad (\text{A.38})$$

We define a *degenerate* solution as one in which at least one path outside of \tilde{P} is a least fleet marginal cost path. Non-degenerate solutions are therefore characterized by $\alpha_r^{(c)} = 0$ for all paths $r \notin \tilde{P}$. When a path p in a non-degenerate solution is used by both the fleet and users, $\alpha_p^{(g)} = \alpha_p^{(f)} = 0$ and the stationarity conditions simplify considerably,

$$1 - \sum_{r \in \tilde{P}} \alpha_r^{(b)} \sum_{a \in A} \delta_{ar} \delta_{ap} \mathbf{t}'_a = 0 \quad (\text{A.39})$$

$$\sum_{r \in \tilde{P}} \alpha_r^{(b)} = 0 \quad (\text{A.40})$$

Together, these equations imply,

$$\alpha_p^{(b)} = - \sum_{r \in \tilde{P} \setminus \{p\}} \alpha_r^{(b)} \quad (\text{A.41})$$

$$1 - \left(\alpha_p^{(b)} \sum_{a \in A} \delta_{ap} \mathbf{t}'_a + \sum_{r \in \tilde{P} \setminus \{p\}} \alpha_r^{(b)} \sum_{a \in A} \delta_{ar} \delta_{ap} \mathbf{t}'_a \right) = 0 \quad (\text{A.42})$$

$$1 + \sum_{r \in \tilde{P} \setminus \{p\}} \alpha_r^{(b)} \sum_{a \in A} \delta_{ap} \mathbf{t}'_a - \sum_{r \in \tilde{P} \setminus \{p\}} \alpha_r^{(b)} \sum_{a \in A} \delta_{ar} \delta_{ap} \mathbf{t}'_a = 0 \quad (\text{A.43})$$

$$1 + \sum_{r \in \tilde{P} \setminus \{p\}} \alpha_r^{(b)} \sum_{a \in A} \delta_{ap} \mathbf{t}'_a (1 - \delta_{ar}) = 0 \quad (\text{A.44})$$

The inner sum in the final equation represents, for each path $r \neq p$, the sum of the link cost derivatives for all links on path p *except* those also on path r . The link cost derivative measures the increase in travel cost incurred by each user of the link should an additional unit of flow be added to the link. This inner term is then the potential travel cost impact on users of path p *not* attributable to increases of flow on path r , that is the impact to travelers on path p of additional flow on path p minus the impact of additional flow to travelers on path p of additional flow on path r weighted by the coefficients $\alpha_r^{(b)}$. Alternatively, when a path $p \in \tilde{P} \cap P^*$ is used by the fleet but not by users, $a_p^{(f)} \geq a_p^{(g)} = 0$, yielding,

$$1 + \sum_{r \in \tilde{P} \setminus \{p\}} \alpha_r^{(b)} \sum_{a \in A} \delta_{ap} \mathbf{t}'_a (1 - \delta_{ar}) \leq 0 \quad (\text{A.45})$$

Finally, when a path $p \in \tilde{P} \cap P^*$ is used by users but not the fleet, $a_p^{(g)} \geq a_p^{(f)} = 0$, yielding,

$$1 + \sum_{r \in \tilde{P} \setminus \{p\}} \alpha_r^{(b)} \sum_{a \in A} \delta_{ap} \mathbf{t}'_a (1 - \delta_{ar}) \geq 0 \quad (\text{A.46})$$

B

Learning from learners Appendices

B.1 Extended summary of test results

Each tested tag is listed in Table B.1. For each tested tag, we report the smallest p value (before Bonferoni correction) over the set of drivers tested as well as the number of drivers for which the tag was found to be significant for each of the two sides of the two-sided test, “seeking” and “avoiding” respectively.

Tables B.2 and B.3 list the tags found to be significant, as well as when significance was first detected, by the procedure applied to the aggregated sequence of stops.

B.2 Increasing probability as the null hypothesis

In this work we treat the hypothesis that the sequence is i.i.d. Bernoulli with unknown mean as the null hypothesis and the hypothesis that the mean is increasing (decreasing) as the alternative. We choose this assignment of null and alternative for two reasons: one conceptual and one technical. Conceptually, the null should be chosen to represent the uninteresting case: the intervention has *no effect*, there is *no relationship* between two measurements, or, in the case of the very first null hypothesis, Muriel Bristol is no better than chance at distinguishing which of tea or milk was first added to a cup (Fisher 1936). In our setting we argue that the uninteresting case is when the probability of a service stop being nearby an entity with a given tag is constant in time. The technical reason is that in order for our statistic to be an e-process we must be able to compute the supremum in the denominator Equation (5.4). This may be done in closed form when the null is i.i.d. Bernoulli, but when the null is instead the set of Bernoulli sequences with non-decreasing mean we may use the following convex optimization to identify the maximum likelihood under the null;

$$\ell_+^*(X) = \max_{\mathbf{p} \in \mathbb{R}^n} \sum_{i=1}^n X_i \log(p_i) + (1 - X_i) \log(1 - p_i) \quad (\text{B.1a})$$

$$\text{s.t. } 0 \leq p_1 \leq p_2 \leq \cdots \leq p_n \leq 1 \quad (\text{B.1b})$$

An analogous program may be used for the hypothesis that the probabilities are non-increasing.

$$\ell_-^*(X) = \max_{\mathbf{p} \in \mathbb{R}^n} \sum_{i=1}^n X_i \log(p_i) + (1 - X_i) \log(1 - p_i) \quad (\text{B.2a})$$

$$\text{s.t. } 1 \geq p_1 \geq p_2 \geq \cdots \geq p_n \geq 0 \quad (\text{B.2b})$$

It should be noted that the i.i.d. Bernoulli product distribution is a member of both of the above constraints sets. As a result, i.i.d. Bernoulli makes a poor alternative hypothesis when either the non-decreasing or non-increasing hypothesis is used as the null. Instead, the non-decreasing and non-increasing hypotheses can be used interchangeably as null and alternative to one another, or the non-stationary alternative p^\pm defined in Equation (5.5) may be used.

Table B.1: Summary of test results

tag	Seeking		Avoiding	
	best p-value	n	best p-value	n
aeroway:aerodrome	3.95e-62	10	8.94e-66	17
amenity:arts_centre	1.48e+00		6.91e-01	
amenity:bank	1.30e-03		6.43e-04	
amenity:bar	2.53e-05		3.00e-04	
amenity:cafe	1.17e-01		1.64e-08	1
amenity:casino	5.31e-10	1	7.10e-12	1
amenity:childcare	1.85e+00		1.80e-01	
amenity:cinema	3.57e-04		5.54e-01	
amenity:college	1.62e-64	1	1.41e-02	
amenity:concert_hall	2.00e+00		2.00e+00	
amenity:courthouse	1.25e+00		5.16e-04	
amenity:crematorium	2.00e+00		2.00e+00	
amenity:dentist	2.25e-06	1	6.34e-02	
amenity:doctors	1.40e-06	1	2.08e-04	
amenity:events_centre	2.00e+00		2.00e+00	
amenity:events_venue	2.00e+00		1.27e-02	
amenity:fast_food	2.46e-11	4	4.37e-18	7
amenity:funeral_home	2.00e+00		2.00e+00	
amenity:grave_yard	1.54e-15	1	5.11e-01	
amenity:hospital	8.64e-04		2.46e-06	1
amenity:library	8.63e-10	1	5.96e-01	
amenity:nightclub	1.77e+00		1.76e-01	
amenity:place_of_worship	5.20e-19	1	9.03e-04	
amenity:pub	3.04e-01		3.42e-04	
amenity:restaurant	2.87e-161	4	1.25e-20	4
amenity:shelter	1.26e-41	5	1.97e-18	3
amenity:theatre	1.25e+00		2.97e-02	
amenity:university	1.30e-08	1	5.22e-14	1
building:commercial	2.57e-130	8	5.13e-36	5
building:residential	2.47e-41	17	3.39e-46	15
building:office	1.25e+00		2.15e-06	1

Table B.2: Bonferroni-corrected p-values for significant tag-seeking behavior.

tag	p value	first detected
amenity:fast_food	3.43e-104	2020-04-21 19:46:07.441
amenity:restaurant	1.22e-76	2020-05-05 18:40:57.001
amenity:college	1.58e-69	2020-09-09 14:57:19.000
building:commercial	2.05e-60	2020-04-14 02:30:59.047
amenity:grave_yard	9.46e-39	2020-07-12 22:41:35.114
amenity:shelter	1.24e-29	2020-05-30 19:49:09.731
aeroway:aerodrome	3.21e-29	2020-09-10 22:53:09.000
amenity:university	1.46e-18	2020-09-10 07:16:22.163
amenity:casino	1.33e-17	2020-08-08 18:34:05.000
building:residential	3.92e-16	2020-05-01 17:06:29.999
amenity:bar	2.93e-12	2020-10-31 18:52:31.048
amenity:cafe	6.54e-11	2020-05-23 12:52:53.143
amenity:hospital	1.41e-10	2020-09-18 07:56:58.000
building:office	1.81e-07	2020-11-02 15:11:41.000
amenity:theatre	5.94e-07	2020-06-08 20:47:15.070

Table B.3: Bonferroni-corrected p-values for significant tag-avoiding behavior.

tag	p value	first detected
aeroway:aerodrome	0.00e+00	2020-03-27 21:25:11.512
amenity:casino	2.63e-31	2020-03-22 20:09:48.494
amenity:dentist	3.30e-26	2020-05-26 09:22:41.397
building:residential	4.02e-24	2020-08-12 19:34:19.999
amenity:arts_centre	1.71e-15	2020-06-11 09:51:00.001
amenity:university	1.92e-08	2020-06-26 23:11:47.998

References

- Agrawal, Akshay et al. (2018). “A rewriting system for convex optimization problems”. In: *Journal of Control and Decision* 5.1, pp. 42–60.
- Agrawal, Shipra and Navin Goyal (2012). “Analysis of thompson sampling for the multi-armed bandit problem”. In: *Conference on learning theory*. JMLR Workshop and Conference Proceedings, pp. 39–1.
- Alcántara, Ann-Marie (Apr. 23, 2021). “Google Maps to Add a Greenest Route to Its Driving Directions”. In: *The Wall Street Journal*. url: <https://www.wsj.com/articles/google-maps-to-add-a-greenest-route-to-its-driving-directions-11619197255> (visited on 09/17/2021).
- Alemi, Farzad et al. (2018). “What influences travelers to use Uber? Exploring the factors affecting the adoption of on-demand ride services in California”. In: *Travel Behaviour and Society* 13, pp. 88–104.
- Allegheny County Airport Authority (Apr. 1, 2020). *Pittsburgh International Airport Summary of Airline Traffic April 2020*. url: <https://flypittsburgh.com/acaa-corporate/about/airport-statistics/>.
- Alonso-Mora, Javier et al. (Jan. 2017). “On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment”. In: *Proceedings of the National Academy of Sciences* 114.3, pp. 462–467. issn: 0027-8424. doi: 10.1073/pnas.1611675114. arXiv: arXiv:1507.06011.
- American Community Survey (2020). *S1701: Poverty Status in the Last 12 Months*. url: <https://data.census.gov/cedsci/table?q=S1701&tid=ACSST5Y2020.S1701>.
- Anderson, Monica (2016). “More Americans using smartphones for getting directions, streaming TV”. In: *Pew Research Center* 29.
- AP Stylebook (June 30, 2015). *Ride-hailing services such as Uber and Lyft may also be called ride-booking services. Do not use ride-sharing*. url: <https://twitter.com/APStylebook/status/615957399233499136> (visited on 03/10/2022).
- Avedian, Sergio (Nov. 8, 2018). *Penny Surge Tried to Defeat Me – But Here’s How I’m Fighting Back!* url: <https://therideshareguy.com/ubers-new-flat-rate-surge/> (visited on 03/10/2020).

- Balding, Melissa et al. (2019). *Estimated Percent of Total Driving by Lyft and Uber*. Tech. rep. Fehr & Peers. url: <https://drive.google.com/file/d/1FIUskVkj9lsAnWJQ6kLhAhNoVLjfFdx3/view>.
- Bansal, Prateek et al. (2020). "Eliciting preferences of TNC users and drivers: evidence from the United States". In: *Travel Behaviour and Society* 20, pp. 225–236.
- Baraniuk, Richard G (2007). "Compressive sensing [lecture notes]". In: *IEEE signal processing magazine* 24.4, pp. 118–121.
- Battifarano, Matthew and Zhen Sean Qian (2019). "Predicting real-time surge pricing of ride-sourcing companies". In: *Transportation Research Part C: Emerging Technologies* 107, pp. 444–462.
- Beckmann, Martin, Charles B McGuire, and Christopher B Winsten (1956). *Studies in the Economics of Transportation*. New Haven, Connecticut: Yale University Press.
- Ben-Akiva, Moshe E. (1985). *Discrete choice analysis theory and application to travel demand*. eng. MIT Press series in transportation studies ; 9. Cambridge, Mass: MIT Press. isbn: 0585138265.
- Ben-Elia, Eran et al. (2013). "The impact of travel information's accuracy on route-choice". In: *Transportation Research Part C: Emerging Technologies* 26, pp. 146–159.
- Berliner, R and Gil Tal (2018). "What drives Your Drivers: An In-depth look at Lyft and Uber Drivers". In: *UC Davis Institute of Transportation Studies*. <https://steps.ucdavis.edu/wpcontent/uploads/2018/02/BERLINER-TAL-What-Drives-Your-Drivers.pdf>.
- Bhat, Chandra R (1997). "Recent methodological advances relevant to activity and travel behavior analysis". In: *International Association of Travel Behavior Research Conference, Austin, Texas*. Citeseer.
- Bimpikis, Kostas, Ozan Candogan, and Daniela Saban (2019). "Spatial pricing in ride-sharing networks". In: *Operations Research*.
- Boyce, David (2007). "Forecasting travel on congested urban transportation networks: review and prospects for network equilibrium models". In: *Networks and Spatial Economics* 7.2, pp. 99–128.
- Boyles, Stephen D., Nicholas E. Lownes, and Avinash Unnikrishnan (2021). *Transportation Network Analysis*. 0.89. Vol. 1. url: <https://sboyles.github.io/blubook.html>.
- Cain, Alasdair, Mark Burris, and Ram Pendyala (2001). "Impact of Variable Pricing on Temporal Distribution of Travel Demand". In: *Transportation Research Record: Journal of the Transportation Research Board* 1747.01, pp. 36–43. issn: 0361-1981. doi: 10.3141/1747-05.
- California Public Utilities Commission (Sept. 13, 2013). *CPUC Establishes rules for Transportation Network Companies*. url: <https://docs.cpuc.ca.gov/PublishedDocs/Published/G000/M077/K132/77132276.PDF> (visited on 03/10/2022).

- Ceci, L. (2021). *Most popular mapping apps in the United States as of April 2018, by reach*. Tech. rep. Statista. url: <https://www.statista.com/statistics/865419/most-popular-us-mapping-apps-ranked-by-reach/>.
- Chen, Le, Alan Mislove, and Christo Wilson (2015). "Peeking Beneath the Hood of Uber". In: *Proceedings of the 2015 Internet Measurement Conference*. IMC '15. New York, NY, USA: ACM, pp. 495–508. isbn: 978-1-4503-3848-6. doi: 10.1145/2815675.2815681.
- Chen, M Keith and Michael Sheldon (2016). "Dynamic Pricing in a Labor Market: Surge Pricing and Flexible Work on the Uber Platform". In: *Proceedings of the 2016 Acm Conference on Economics and Computation*, pp. 1–19. issn: 9781450339360. doi: 10.1145/2940716.2940798.
- Chen, Zhibin et al. (2020). "Path controlling of automated vehicles for system optimum on transportation networks with heterogeneous traffic stream". In: *Transportation Research Part C: Emerging Technologies* 110, pp. 312–329.
- Clewlöw, Regina R and Gouri S Mishra (2017). *Disruptive transportation: The adoption, utilization, and impacts of ride-hailing in the United States*. Tech. rep. Institute of Transportation Studies, University of California, Davis.
- Cochran, Abigail L and Daniel G Chatman (2021). "Use of app-based ridehailing services and conventional taxicabs by adults with disabilities". In: *Travel Behaviour and Society* 24, pp. 124–131.
- Cohen, Peter et al. (2016). *Using Big Data to Estimate Consumer Surplus: The Case of Uber*. Tech. rep. National Bureau of Economic Research. doi: 10.3386/w22627.
- Cominetti, Roberto, Jose R Correa, and Nicolás E Stier-Moses (2009). "The impact of oligopolistic competition in networks". In: *Operations Research* 57.6, pp. 1421–1437.
- Dafermos, Stella C (1972). "The traffic assignment problem for multiclass-user transportation networks". In: *Transportation science* 6.1, pp. 73–87.
- Daganzo, Carlos F and Yosef Sheffi (1977). "On stochastic models of traffic assignment". In: *Transportation science* 11.3, pp. 253–274.
- Delle Site, Paolo (2021). "Pricing of connected and autonomous vehicles in mixed-traffic networks". In: *Transportation Research Record*, p. 0361198120985850.
- Diamond, Steven and Stephen Boyd (2016). "CVXPY: A Python-embedded modeling language for convex optimization". In: *Journal of Machine Learning Research* 17.83, pp. 1–5.
- Dias, Felipe F et al. (2017). "A behavioral choice model of the use of car-sharing and ride-sourcing services". In: *Transportation* 44.6, pp. 1307–1323.
- Erhardt, Gregory D et al. (2019). "Do transportation network companies decrease or increase congestion?" In: *Science advances* 5.5, eaau2670.

- Fagnant, Daniel J and Kara M Kockelman (2014). "The travel and environmental implications of shared autonomous vehicles, using agent-based model scenarios". In: *Transportation Research Part C: Emerging Technologies* 40, pp. 1–13.
- Fei, Xiang, Chung-cheng Lu, and Ke Liu (2011). "A bayesian dynamic linear model approach for real-time short-term freeway travel time prediction". In: *Transportation Research Part C* 19.6, pp. 1306–1318. issn: 0968-090X. doi: 10.1016/j.trc.2010.10.005.
- Fisher, Ronald Aylmer (1936). "Design of experiments". In: *British Medical Journal* 1.3923, p. 554.
- Friesz, Terry L (1985). "Transportation network equilibrium, design and aggregation: key developments and research opportunities". In: *Transportation Research Part A: General* 19.5-6, pp. 413–427.
- Friesz, Terry L et al. (1994). "Day-to-day dynamic network disequilibria and idealized traveler information systems". In: *Operations Research* 42.6, pp. 1120–1136.
- Ghemawat, Sanjay and Jeff Dean (2014). *LevelDB*. url: <https://github.com/google/leveldb>.
- Grahn, Rick et al. (2020). "Socioeconomic and usage characteristics of transportation network company (TNC) riders". In: *Transportation* 47.6, pp. 3047–3067.
- Grünwald, Peter, Rianne de Heide, and Wouter M Koolen (2020). "Safe testing". In: *2020 Information Theory and Applications Workshop (ITA)*. IEEE, pp. 1–54.
- Guda, Harish and Upender Subramanian (2017). "Your Uber is Arriving : Managing On-Demand Workers through Surge Pricing , Forecast Communication and Worker Incentives". In: *Management Science*.
- Guha, Samayita, Emre M Demirezen, and Subodha Kumar (2018). "Dynamics of Competition in On-Demand Economy: A Differential Games Approach". In: *Available at SSRN 3263152*.
- Gurley, Bill (2014). *A Deeper Look at Uber's Dynamic Pricing Model*. UBER Newsroom. url: <https://www.uber.com/newsroom/guest-post-a-deeper-look-at-ubers-dynamic-pricing-model/>.
- Hall, Jonathan, Cory Kendrick, and Chris Nosko (2015). "The Effects of Uber's Surge Pricing: A Case Study". In: *Uber Under The Hood*, pp. 1–8.
- Hall, Jonathan D, Craig Palsson, and Joseph Price (2018). "Is Uber a substitute or complement for public transit?" In: *Journal of Urban Economics* 108, pp. 36–50.
- Hall, Jonathan V and Alan B Krueger (2018). "An analysis of the labor market for Uber's driver-partners in the United States". In: *Ilr Review* 71.3, pp. 705–732.
- Harker, Patrick T (1988). "Multiple equilibrium behaviors on networks". In: *Transportation science* 22.1, pp. 39–46.
- Hazelton, Martin L (2022). "The emergence of stochastic user equilibria in day-to-day traffic models". In: *Transportation Research Part B: Methodological* 158, pp. 102–112.

- He, Fang and Zuo Jun Max Shen (Sept. 2015). "Modeling taxi services with smartphone-based e-hailing applications". In: *Transportation Research Part C: Emerging Technologies* 58, pp. 93–106. issn: 0968090X. doi: 10.1016/j.trc.2015.06.023.
- Head, Megan L et al. (2015). "The extent and consequences of p-hacking in science". In: *PLoS biology* 13.3, e1002106.
- Hensley, Russell, Asutosh Padhi, and Jeff Salazar (July 17, 2017). *Cracks in the ridesharing market—and how to fill them*. McKinsey & Company. url: <https://www.mckinsey.com/industries/automotive-and-assembly/our-insights/cracks-in-the-ridesharing-market-and-how-to-fill-them> (visited on 03/15/2022).
- Hestenes, Magnus R (1969). "Multiplier and gradient methods". In: *Journal of optimization theory and applications* 4.5, pp. 303–320.
- Hoffman, Matthew, Francis Bach, and David Blei (2010). "Online learning for latent dirichlet allocation". In: *advances in neural information processing systems* 23.
- Howard, Steven R and Aaditya Ramdas (2019). "Sequential estimation of quantiles with applications to A/B-testing and best-arm identification". In: *arXiv preprint arXiv:1906.09712*.
- Howard, Steven R, Aaditya Ramdas, et al. (2020). "Time-uniform Chernoff bounds via nonnegative supermartingales". In: *Probability Surveys* 17, pp. 257–317.
- Jan, Oliver, Alan J Horowitz, and Zhong-Ren Peng (2000). "Using global positioning system data to understand variations in path choice". In: *Transportation Research Record* 1725.1, pp. 37–44.
- Ke, Jintao, Siyuan Feng, et al. (2021). "Joint predictions of multi-modal ride-hailing demands: A deep multi-task multi-graph learning-based approach". In: *Transportation Research Part C: Emerging Technologies* 127, p. 103063.
- Ke, Jintao, Xiaoran Qin, et al. (2021). "Predicting origin-destination ride-sourcing demand with a spatio-temporal encoder-decoder residual multi-graph convolutional network". In: *Transportation Research Part C: Emerging Technologies* 122, p. 102858.
- Ke, Jintao, Hai Yang, et al. (2019). "Hexagon-Based Convolutional Neural Network for Supply-Demand Forecasting of Ride-Sourcing Services". In: *IEEE Transactions on Intelligent Transportation Systems* 20.11, pp. 4160–4173. doi: 10.1109/TITS.2018.2882861.
- Ke, Jintao, Hongyu Zheng, et al. (2017). "Short-term forecasting of passenger demand under on-demand ride services: A spatio-temporal deep learning approach". In: *Transportation Research Part C: Emerging Technologies* 85.October, pp. 591–608. issn: 0968090X. doi: 10.1016/j.trc.2017.10.016. arXiv: 1706.06279.
- Kontou, Eleftheria, Venu Garikapati, and Yi Hou (2020). "Reducing ridesourcing empty vehicle travel with future travel demand prediction". In: *Transportation Research Part C: Emerging Technologies* 121, p. 102826.
- Koster, Raph (2017). *Still Logged In: What AR and VR Can Learn from MMOs*. Game Developers Conference 2017. url: <https://vimeo.com/208372546> (visited on 03/11/2022).

- Lapowsky, Issie (July 2015). *Uber Wins Its Battle Against NYC's Mayor—For Now*. url: <https://www.wired.com/2015/07/uber-wins-battle-nyc-mayor-now/>.
- Lau, Johann (Sept. 3, 2020). *Google Maps 101: How AI helps predict traffic and determine routes*. url: <https://blog.google/products/maps/google-maps-101-how-ai-helps-predict-traffic-and-determine-routes/> (visited on 05/04/2022).
- Lavieri, Patrícia S and Chandra R Bhat (2019). "Investigating objective and subjective factors influencing the adoption, frequency, and characteristics of ride-hailing trips". In: *Transportation Research Part C: Emerging Technologies* 105, pp. 100–125.
- Lee, Jason D et al. (2016). "Exact post-selection inference, with application to the lasso". In: *The Annals of Statistics* 44.3, pp. 907–927.
- Lee, Min Kyung et al. (2015). "Working with Machines: The Impact of Algorithmic and Data-Driven Management on Human Workers". In: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. CHI '15. Seoul, Republic of Korea: ACM, pp. 1603–1612. isbn: 978-1-4503-3145-6. doi: 10.1145/2702123.2702548.
- Leiner, Barry M. et al. (1997). *Brief History of the Internet*. Internet Society. url: <https://www.internetsociety.org/internet/history-internet/brief-history-internet/> (visited on 03/17/2022).
- Leonard, Robert J (1994). "Reading Cournot, reading Nash: The creation and stabilisation of the Nash equilibrium". In: *The Economic Journal* 104.424, pp. 492–511.
- Levinson, David (2003). "The value of advanced traveler information systems for route choice". In: *Transportation Research Part C: Emerging Technologies* 11.1, pp. 75–87.
- Li, Aoyong and Kay W Axhausen (2020). "Short-term traffic demand prediction using graph convolutional neural networks". In: *AGILE: GIScience Series* 1, pp. 1–14.
- Li, Yang et al. (2017). "Forecasting short-term subway passenger flow under special events scenarios using multiscale radial basis function networks". In: *Transportation Research Part C: Emerging Technologies* 77, pp. 306–328.
- Licklider, J. C. R. (1960). "Man-Computer Symbiosis". In: *IRE Transactions on Human Factors in Electronics* HFE-1.1, pp. 4–11. doi: 10.1109/THFE2.1960.4503259.
- Licklider, Joseph CR and Robert W Taylor (1968). "The computer as a communication device". In: *Science and technology* 76.2, pp. 1–3.
- Liu, Jielun et al. (2019). "Spatial-temporal inference of urban traffic emissions based on taxi trajectories and multi-source urban data". In: *Transportation Research Part C: Emerging Technologies* 106, pp. 145–165.
- Lockhart, Richard et al. (2014). "A significance test for the lasso". In: *Annals of statistics* 42.2, p. 413.
- Loftus, Joshua R and Jonathan E Taylor (2014). "A significance test for forward stepwise model selection". In: *arXiv preprint arXiv:1405.3920*.

- Lou, Yin et al. (2013). "Accurate intelligible models with pairwise interactions". In: *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 623–631.
- Ma, Wei and Sean Qian (2019). "Measuring and reducing the disequilibrium levels of dynamic networks through ride-sourcing vehicle data". In: *arXiv preprint arXiv:1905.05386*.
- Macfarlane, Jane (2019). "When apps rule the road: The proliferation of navigation apps is causing traffic chaos. It's time to restore order". In: *IEEE Spectrum* 56.10, pp. 22–27. doi: 10.1109/MSPEC.2019.8847586.
- Mansourianfar, Mohammad Hadi et al. (2021). "Joint routing and pricing control in congested mixed autonomy networks". In: *Transportation Research Part C: Emerging Technologies* 131, p. 103338.
- McDonald's Corporation (June 30, 2020). *Form 10-Q*. url: <https://www.sec.gov/ix?doc=/Archives/edgar/data/0000063908/000006390820000063/mcd-6302020x10q.htm>.
- McFadden, Daniel et al. (1973). "Conditional logit analysis of qualitative choice behavior". In: .
- McFadden, Daniel (1974). "The measurement of urban travel demand". In: *Journal of public economics* 3.4, pp. 303–328.
- Mehr, Negar and Roberto Horowitz (2019). "How will the presence of autonomous vehicles affect the equilibrium state of traffic networks?" In: *IEEE Transactions on Control of Network Systems* 7.1, pp. 96–105.
- Miller, John, Yu Nie, and Amanda Stathopoulos (2017). "Crowdsourced urban package delivery: Modeling traveler willingness to work as crowdshippers". In: *Transportation Research Record* 2610.1, pp. 67–75.
- Min, Wanli and Laura Wynter (2011). "Real-time road traffic prediction with spatio-temporal correlations". In: *Transportation Research Part C* 19.4, pp. 606–616. issn: 0968-090X. doi: 10.1016/j.trc.2010.10.002.
- Mohammad M. Hamed, Hashem R. Al-Masaeid, and Ahi M. Bani Said (1995). "Short-Term Prediction of Traffic Volume in Urban Arterials". In: *Journal of Transportation Engineering* 121.3, pp. 249–254.
- Moran, Maarit and Philip Lasley (2017). "Legislating transportation network companies". In: *Transportation Research Record* 2650.1, pp. 163–171.
- Moulin, Herve (1986). *Game theory for the social sciences*. eng. 2nd and rev. ed. Studies in game theory and mathematical economics. New York: New York University Press. isbn: 0814754309.
- National Association of City Transportation Officials (2020). *Shared Micromobility in the U.S.: 2019*. url: <https://nacto.org/shared-micromobility-2019/> (visited on 03/15/2022).
- Nori, Harsha et al. (2019). "InterpretML: A Unified Framework for Machine Learning Interpretability". In: *arXiv preprint arXiv:1909.09223*.

- Noursalehi, Peyman, Haris N Koutsopoulos, and Jinhua Zhao (2018). "Real time transit demand prediction capturing station interactions and impact of special events". In: *Transportation Research Part C: Emerging Technologies* 97, pp. 277–300.
- O'Dea, S. (2020). *Smartphone penetration rate as share of the population in the United States from 2010 to 2021*. Tech. rep. Statista. url: <https://www.statista.com/statistics/201183/forecast-of-smartphone-penetration-in-the-us/>.
- OpenStreetMap contributors (2022). *OpenStreetMaps Data*. <https://www.openstreetmap.org>.
- Park, Mee Young and Trevor Hastie (2007). "L1-regularization path algorithm for generalized linear models". In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 69.4, pp. 659–677. doi: 10.1111/j.1467-9868.2007.00607.x.
- Pedregosa, F. et al. (2011). "Scikit-learn: Machine Learning in Python". In: *Journal of Machine Learning Research* 12, pp. 2825–2830.
- Peeta, Srinivas and Athanasios K Ziliaskopoulos (2001). "Foundations of dynamic traffic assignment: The past, the present and the future". In: *Networks and spatial economics* 1.3, pp. 233–265.
- Phillips, Bob (2017). *Balancing Supply and Demand in a Two-Sided Marketplace*. Plenary Talk.
- Punel, Aymeric and Amanda Stathopoulos (2017). "Modeling the acceptability of crowd-sourced goods deliveries: Role of context and experience effects". In: *Transportation Research Part E: Logistics and Transportation Review* 105, pp. 18–38.
- Ramdas, Aaditya, Johannes Ruf, Martin Larsson, and Wouter Koolen (2020). "Admissible anytime-valid sequential inference must rely on nonnegative martingales". In: *arXiv preprint arXiv:2009.03167*.
- Ramdas, Aaditya, Johannes Ruf, Martin Larsson, and Wouter M Koolen (2022). "Testing exchangeability: Fork-convexity, supermartingales and e-processes". In: *International Journal of Approximate Reasoning* 141, pp. 83–109.
- Rayle, Lisa et al. (2014). "App-based, on-demand ride services: Comparing taxi and ridesourcing trips and user characteristics in san francisco university of california transportation center (uctc)". In: *University of California: Berkeley, CA, USA*.
- Rosenblat, Alex and Luke Stark (2016). "Algorithmic Labor and Information Asymmetries: A Case Study of Uber's Drivers". In: *International Journal of Communication* 10.0, p. 27. issn: 1932-8036. doi: 10.2139/ssrn.2686227.
- Roughgarden, Tim (2005). *Selfish routing and the price of anarchy*. MIT press.
- Ruijter, Arjan de et al. (2022). "Evolution of labour supply in ridesourcing". In: *Transportmetrica B: Transport Dynamics* 10.1, pp. 599–626.
- San Francisco County Transportation Authority (2017). "TNCs Today: A Profile of San Francisco Transportation Network Company Activity". In.

- Satariano, Adam (Sept. 2017). *Uber Losing Battle in London After Regulator Revokes License*. url: <https://www.bloomberg.com/news/articles/2017-09-22/london-authority-revokes-uber-s-private-hire-license>.
- Shafer, Glenn (2019). "The language of betting as a strategy for statistical and scientific communication". In: *arXiv preprint arXiv:1903.06991*.
- Shafer, Glenn, Alexander Shen, et al. (2011). "Test martingales, Bayes factors and p-values". In: *Statistical Science* 26.1, pp. 84–101.
- Shafer, Glenn and Vladimir Vovk (2019). *Game-Theoretic Foundations for Probability and Finance*. Vol. 455. John Wiley & Sons.
- Shaffer, Juliet Popper (1995). "Multiple hypothesis testing". In: *Annual review of psychology* 46.1, pp. 561–584.
- Sharon, Guni et al. (2018). "Traffic optimization for a mixture of self-interested and compliant agents". In: *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Sheffi, Yosef (1985). *Urban transportation networks*. Vol. 6. Prentice-Hall, Englewood Cliffs, NJ.
- Sheth, Amit, Pramod Anantharam, and Cory Henson (2013). "Physical-cyber-social computing: An early 21st century approach". In: *IEEE Intelligent Systems* 28.1, pp. 78–82.
- Shin, Jaehyeok, Aaditya Ramdas, and Alessandro Rinaldo (2022). "E-detectors: a non-parametric framework for online changepoint detection". In: doi: 10.48550/ARXIV.2203.03532. url: <https://arxiv.org/abs/2203.03532>.
- Smart, John, Jamais Cascio, and Jerry Paffendorf (2007). *The Metaverse Roadmap*. url: <https://www.metaverseroadmap.org/MetaverseRoadmapOverview.pdf> (visited on 03/11/2022).
- Smith, Michael J (1979). "The existence, uniqueness and stability of traffic equilibria". In: *Transportation Research Part B: Methodological* 13.4, pp. 295–304.
- SRI International (Dec. 10, 1968). *1968 Demo - FJCC Conference Presentation Reel #1*. url: https://archive.org/details/XD300-23_68HighlightsAResearchCntAugHumanIntellect (visited on 03/17/2022).
- Stabler, Ben, Hillel Bar-Gera, and Elizabeth Sall (2021). *Transportation Networks for Research*. Accessed 10 November 2021. url: <https://github.com/bstabler/TransportationNetworks>.
- Stephenson, Neal (1992). *Snow Crash*. Bantam Books. isbn: 0-553-08853-X.
- Sun, Hao, Hai Wang, and Zhixi Wan (2019). "Model and analysis of labor supply for ride-sharing platforms in the presence of sample self-selection and endogeneity". In: *Transportation Research Part B: Methodological* 125, pp. 76–93.
- Sutton, Richard S, Andrew G Barto, et al. (1998). "Introduction to reinforcement learning". In.
- Terdiman, Daniel (May 26, 2010). *SF AppShow gives developers leg up on competition*. CNET. url: <https://www.cnet.com/culture/sf-appshow-gives-developers-leg-up-on-competition/> (visited on 03/11/2022).

- Tingvall, Claes and Narelle Haworth (1999). "Vision Zero-An ethical approach to safety and mobility". In: *6th ITE International Conference Road Safety & Traffic Enforcement: Beyond 2000*.
- Tobin, Roger L and Terry L Friesz (1988). "Sensitivity Analysis for Equilibrium Network Flow". In: *Transportation Science* 22.4, pp. 242–250.
- Uber Blog (Jan. 8, 2020). *To California drivers: keeping you in the driver's seat, Part 2*. url: <https://www.uber.com/blog/california/keeping-you-in-the-drivers-seat-2/> (visited on 03/10/2022).
- Uber Technologies (2018). *Riding with Uber: Upfront Pricing*. url: <https://www.uber.com/ride/how-uber-works/upfront-pricing/>.
- (Dec. 31, 2019). *Form 10-K*. url: <https://www.sec.gov/ix?doc=/Archives/edgar/data/0001543151/000154315120000010/fy2019q410kfinancialst.htm>.
- (June 30, 2020). *Form 10-Q*. url: <https://www.sec.gov/ix?doc=/Archives/edgar/data/0001543151/000154315120000029/uber-20200630.htm>.
- (2021). *Uber Technologies, Form 10-K*. url: <https://www.sec.gov/ix?doc=/Archives/edgar/data/1543151/000154315122000008/uber-20211231.htm>.
- (2022). *How does Uber match riders with drivers?* url: <https://www.uber.com/us/en/marketplace/matching/> (visited on 05/04/2022).
- Vanajakshi, Lelitha and Laurence R. Rilett (2004). "A Comparison Of The Performance Of Artificial. Neural Networks And Support Vector Machines For The Prediction Of Traffic Speed". In: *IEEE Intelligent Vehicles Symposium*. Parma, Italy, pp. 194–199.
- Ville, Jean (1939). "Etude critique de la notion de collectif". In: *Gauthier-Villars, Paris*.
- Vovk, Vladimir and Ruodu Wang (2021). "E-values: Calibration, combination and applications". In: *The Annals of Statistics* 49.3, pp. 1736–1754.
- Wald, A (1945). "Sequential tests of statistical hypotheses". In: *Annual Mathematica Statistics*.
- Wang, Dong et al. (2017). "DeepSD: Supply-Demand Prediction for Online Car-Hailing Services Using Deep Neural Networks". In: *2017 IEEE 33rd International Conference on Data Engineering (ICDE)*, pp. 243–254. doi: 10.1109/ICDE.2017.83.
- Wang, Jiawei et al. (2020). "Controllability analysis and optimal control of mixed traffic flow with human-driven and autonomous vehicles". In: *IEEE Transactions on Intelligent Transportation Systems*.
- Wang, Xiaolei et al. (Sept. 2016). "Pricing strategies for a taxi-hailing platform". In: *Transportation Research Part E: Logistics and Transportation Review* 93, pp. 212–231. issn: 13665545. doi: 10.1016/j.tre.2016.05.011.
- Ward, Jacob W et al. (2019). "Effects of on-demand ridesourcing on vehicle ownership, fuel consumption, vehicle miles traveled, and emissions per capita in US States". In: *Transportation Research Part C: Emerging Technologies* 108, pp. 289–301.
- Wardrop, John Glen (1952). "Some Theoretical Aspects of Road Traffic Research". In: *Proceedings of the institution of civil engineers* 1.3, pp. 325–362.

- Wasserman, Larry, Aaditya Ramdas, and Sivaraman Balakrishnan (2020). "Universal inference". In: *Proceedings of the National Academy of Sciences* 117.29, pp. 16880–16890.
- Watling, David and Martin L Hazelton (2003). "The dynamics and equilibria of day-to-day assignment models". In: *Networks and Spatial Economics* 3.3, pp. 349–370.
- Waudby-Smith, Ian and Aaditya Ramdas (2020). "Estimating means of bounded random variables by betting". In: *arXiv preprint arXiv:2010.09686*.
- Wei, Yu and Mu Chen Chen (2012). "Forecasting the short-term metro passenger flow with empirical mode decomposition and neural networks". In: *Transportation Research Part C: Emerging Technologies* 21.1, pp. 148–162. issn: 0968090X. doi: 10.1016/j.trc.2011.06.009.
- Wright, Stephen, Jorge Nocedal, et al. (1999). "Numerical optimization". In: *Springer Science* 35.67-68, p. 7.
- Xiong, Gang et al. (2015). "Cyber-physical-social system in intelligent transportation". In: *IEEE/CAA Journal of Automatica Sinica* 2.3, pp. 320–333.
- Yan, Xiang, Xinyu Liu, and Xilei Zhao (2020). "Using machine learning for direct demand modeling of ridesourcing services in Chicago". In: *Journal of Transport Geography* 83, p. 102661.
- Yang, Hai and Teng Yang (May 2011). "Equilibrium properties of taxi markets with search frictions". In: *Transportation Research Part B: Methodological* 45.4, pp. 696–713. issn: 01912615. doi: 10.1016/j.trb.2011.01.002.
- Yang, Hai, Xiaoning Zhang, and Qiang Meng (Oct. 2007). "Stackelberg games and multiple equilibrium behaviors on networks". In: *Transportation Research Part B: Methodological* 41.8, pp. 841–861. issn: 0191-2615.
- Zha, Liteng, Yafeng Yin, and Zhengtian Xu (2018). "Geometric matching and spatial pricing in ride-sourcing markets". In: *Transportation Research Part C: Emerging Technologies* 92, pp. 58–75.
- Zha, Liteng, Yafeng Yin, and Hai Yang (Oct. 2016). "Economic analysis of ride-sourcing markets". In: *Transportation Research Part C: Emerging Technologies* 71, pp. 249–266. issn: 0968090X. doi: 10.1016/j.trc.2016.07.010.
- Zhang, Kenan and Yu Marco Nie (2018). "Mitigating the impact of selfish routing: An optimal-ratio control scheme (ORCS) inspired by autonomous driving". In: *Transportation Research Part C: Emerging Technologies* 87, pp. 75–90.
- Zhang, Pinchao, Zhen, and Qian (2017). "User-centric interdependent urban systems: using time-of-day electricity usage data to predict morning roadway congestion". In: *Transportation Research Part C: Emerging Technologies*. arXiv: 1710.08437.
- Zheng, Weizhong et al. (2006). "Short-Term Freeway Traffic Flow Prediction : Bayesian Combined Neural Network Approach". In: *Journal of Transportation Engineering* 132.February, pp. 114–121.

- Zheng, Zimu et al. (2016). "Urban Traffic Prediction through the Second Use of Inexpensive Big Data from Buildings". In: *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management - CIKM '16*, pp. 1363–1372. isbn: 9781450340731. doi: 10.1145/2983323.2983357.
- Zhu, Shanjiang and David Levinson (2015). "Do people use the shortest path? An empirical test of Wardrop's first principle". In: *PloS one* 10.8, e0134322.
- Ziebart, Brian D et al. (2008). "Maximum entropy inverse reinforcement learning." In: *Aaai*. Vol. 8. Chicago, IL, USA, pp. 1433–1438.