Exploring the Impact of Incorporating On-Demand Drivers into a Child Rideshare Platform

by

Katherine Dieppa

B.S. Civil and Environmental Engineering
Massachusetts Institute of Technology, 2017

SUBMITTED TO THE
DEPARTMENT OF CIVIL AND ENVIRONMENTAL ENGINEERING
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF ENGINEERING
IN CIVIL AND ENVIRONMENTAL ENGINEERING

AT THE

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

JUNE 2019

Signature redacted

Signature of Author..
Department of Civil and Environmental Engineering
May 3, 2019

Signature redacted

Certified by ....
Carolina Osorio
Associate Professor of Civil and Environmental Engineering
Thesis Supervisor

Signature redacted

Accepted by ....
Heidi Nepf
Donald and Martha Harleman Professor of Civil and Environmental Engineering
Chair, Graduate Program Committee

Exploring the Impact of Incorporating On-Demand Drivers to a Child Rideshare Platform

by

Katherine Dieppa

Submitted to the Department of Civil and Environmental Engineering on May 3, 2019
in Partial Fulfillment of the Requirements for the Degree of Master of Engineering
in Civil and Environmental Engineering

Abstract

Traditional transportation options for children consist of school bus rides and car rides from parents or friends. With an increase in children's involvement in after-school activities and a rise in the number of households with two working parents, people are looking for a better solution to child transportation. Many rideshare companies have already sprung up with a focus to serve this younger population (examples include HopSkipDrive, Zum, and Kango). While these companies provide flexibility by allowing parents to book both pre-scheduled and on-demand rides, they do not provide the necessary reliability. Much like the rideshare giants of today, Uber and Lyft, these child-geared services do not hire their drivers as employees. Instead, the drivers are treated as independent contractors. Contracting drivers means that the driver supply is never guaranteed at any given time since it is up to each driver to accept or deny the ride requests that they receive. This study proposes a driver supply model in which all drivers are hired as employees, as opposed to contractors. Hiring drivers as employees allows for a company to train drivers extensively while on the job, has been proven to result in higher retention levels, and most importantly, enables the company to have clear insight into their supply at any given time. This study aims to identify how a driver pool should be divided (if at all) into two groups: one dedicated to serving pre-scheduled rides and another reserved for serving on-demand rides. Using the Gurobi™ optimization solver, we observe expected profits under a variety of scenarios in which we vary the proportion of drivers reserved to handle on-demand requests. Our results show that the optimal proportion of on-demand drivers is dependent on both the total number of requests received and the proportion of on-demand requests received. Results are discussed in detail.

Thesis Supervisor: Carolina Osorio

Title: Associate Professor of Civil and Environmental Engineering

Acknowledgements

I would like to extend my sincerest appreciation for my thesis advisor, Carolina Osorio, for her patience and support. She also served as my academic advisor throughout my time as a Master of Engineering candidate and as such, served as my sounding board and voice of reason during numerous moments of distress. Carolina: I could not have done this without you.

I would also like to express sincere gratitude for Tianli Zhou, a PhD candidate in transportation (also advised by Carolina), who sat alongside me for countless hours as I worked to debug my optimization model. Tianli: Your unrelenting willingness to help a fellow student is admirable and represents everything I love about MIT. I could not have done this without you.

Lastly, I would like to thank my incredible family and friends for their unwavering support and encouragement throughout this endeavor. To my sister, Evy: You are my role model. Thank you for constantly reminding me that I am capable of so much more than I think. I could not have done this without you. To my mom, mami: Thank you for teaching me that I can do anything I put my mind to. You never let me feel that there was a limit to what I was capable of, and for that, I thank you. I could not have done this without you. To my love, Alex: You make me feel invincible each and every day of my life. I wouldn't be as strong as I am without you by my side, and I certainly could not have done this without you. And lastly, to my late dad, papi: This one's for you. I could not have done this without you. I hope to make you proud up there. Love, your Catalina Peluca.

Table of Contents

List of Tables

List of Figures

1. Introduction

1.1. Background and motivation

1.1.1. Introduction to child rideshare services

The rise of ride-hailing apps like Uber and Lyft have made it very easy for urban dwellers without a car to get around without the use of public transportation. Users simply download the app, upload their payment information and request a ride from any location within the service's coverage area. There's just one major caveat: these rideshare services are limited to those who are 18 years or older [1]. Transportation options for minors have remained more or less the same throughout the past few decades. Kids have to rely on school buses, after school program shuttles, parents and friends' parents to get around. Luckily though, a handful of startups have set out to offer rideshare convenience to the youngest demographic. Companies like HopSkipDrive, Zum, and Kango are the leading the pack in the child rideshare space.

Unlike major rideshare companies such as Uber and Lyft, recent companies focusing on transportation for children allow for customers to schedule recurring rides, in addition to one-off rides. Uber has recently provided customers with the ability to schedule rides in advance, though these individual rides still need to be accepted by individual drivers and ultimately hinge on that driver's ability to follow through. Since these rideshare companies contract their drivers, the responsibility is placed on the driver. Many emerging rideshare companies tailored to rides for children have followed this model. Drivers are contracted and have the option to accept ride requests on a ride-by-ride basis. This differs greatly from traditional transportation services for children. In the past, parents have had to depend on the routes provided by school buses or hope that their student's after-school program offered transportation. Now, parents have the ability to plan their student's rides based on their individual needs. The only problem is that today's busy kids, perhaps more-so than the average adult, have complex schedules. They are irregular and highly subject to change. Luckily though, a large portion of a child's travel remains predictable. This sort of demand is much easier to supply. If you hire drivers who know exactly when they are expected to drive from day to day, then they can easily incorporate this work into their daily schedules. The difficulty comes in working around the irregular rides (such as those to soccer practice, dance, or the friend's house), the last-minute requests when a parent can't leave work on time, or the change in routes when a parent cancels a ride because their student is sick. The focus of this study is to explore a driver supply model that can best accommodate both pre-scheduled and on-demand requests.

1.1.2. Distinction between employees and contractors

Distinguishing between "standard" and "contingent" (or contracted) employees is a common way to differentiate between workforce members [2]. Stirpe et al. state that "the differences between the two types of employees reside in two fundamental criteria: the duration of their employment relationship and the quantity of HR investments the firm directs to them" [3][4].

A number of human resource (HR) incentives are offered to standard employees, including, but not limited to, economic incentives, training, empowerment and leadership opportunities [3]. As a result, standard employees tend to remain at their employer for longer periods of time than their contingent counterparts. For example, when hired as a standard employee, individuals may not possess all of the necessary skills for a job. Instead, employers often make informed decisions on the potential quality of an individual and invest time and resources to train and empower that individual. As a company employee, individuals are often exposed to the potential of "moving up" in the company. I.e., there are opportunities for growth and the employee has incentive to work hard in order to earn a higher position in the company, which is often linked with higher compensation, increased benefits, and most importantly, higher quality work [3].

In contrast, contingent workers (which I may refer to as "contractors," "contracted workers," or "contracted employees" from this point on) typically work for the employer for a limited amount of time, or on distinct, independent projects. For example, contracted workers are prevalent in the construction industry. Construction companies are often general contractors. This means that while they oversee and manage the entire construction process, they actually employ a variety of "subcontractors" to perform much of the specialized work, such as electrical, plumbing, and flooring projects. Given that construction is a highly complex, disaggregated process, the subcontractor system works well for this industry. It would be difficult for a company to take on the challenge of training its employees in each of these specialties. Instead, it capitalizes on the training and effort imparted by the specialized companies. Despite these benefits, hiring contractors can be a source of stress for a company as well. For example, lack of oversight and ability to manage the workers directly can result in work not getting done the right way. This may be frustrating for a company and can result in construction delays or increased costs when mistakes are made.

1.1.3. Employees vs. contractors for child rideshare applications

Given the sensitive nature of working with children, and the fact that extensive vetting and training is required for the job, we can see why employees would make the best choice for child transportation. As described in section 1.1.2, being an employee implies a much greater sense of association with a company. By nature, it brings the employee closer to the company and aims to connect it with the company's values and mission. Employee status means that the company can invest in training and even provide opportunities for growth. The employee's role can evolve over time, leading them to become an expert in their role so they can then serve as examples for other, less-tenured employees. A sense of community within the company can be instrumental in keeping employees around longer (thus increasing the average tenure). In the realm of child rideshare, employees with a long tenure and demonstrated commitment to the job can make the difference between a parent choosing one service over another.

Employing drivers for ride-hailing has proven successful for child applications since it allows for ample training and procedural requirements to be enforced. The only downside is that employees are not typically called upon at the company's will. In order to ensure your supply of drivers, you need to be paying them for times that are scheduled in advance. This makes it possible to supply drivers for rides that are scheduled well in advance. However, when supplying the rides that are received "last-minute" (day before or even day-of), a company cannot rely on their employee to be available if they are not "on shift." It is not economical for companies to pay an abundance of employees for hours that may or may not be put to use.

Hiring drivers for on-demand child transportation is challenging for two main reasons: (1) demand is last-minute and can be highly variable so it is nearly impossible to schedule these workers in advance, and (2) candidates must undergo a rigorous vetting and training process in order to work with minors. Uber and Lyft are able to contract their drivers because essentially no training or supplemental information is required. You need a phone that can download their driver application, you need a car that meets their requirements, and you need to pass a background check. By nature of employing independent contractors, companies such as Uber and Lyft cannot require their drivers to undergo training. They are expected to have the experience required to do the job.

Contracting drivers is an attractive option for any company operating in the shared ride-hailing economy. The low barrier to entry makes it easy to bring people on board, and the fact that contractors are only paid when they render services to the customer makes it economical for the employer. However, in the much more nuanced world of ride-hailing for children, the cons far outweigh the pros. When hiring drivers for children, it is imperative that the employer enforce strict procedures. The nature of child transportation is such that drivers need to be in constant communication with the employer in order to ensure that tasks are completed and information is relayed properly. In the delicate world of child transportation, drivers are constantly receiving specific direction from their employer. It is necessary since, as minors, their "clients" cannot be the ones to provide that information. This is where many of the existing child rideshare companies have failed. Entities such as Shuddle and Zemcar have relied on this contractor model, ultimately resulting in their demise or near collapse [5][6].

Companies who hire drivers as employees (as opposed to contractors) are crucial for child transportation because the nature of the job requires detailed, involved training and a true commitment to the job. Many reports show the extremely high driver turnover rate experienced by rideshare companies [7]. Such high turnover would be highly detrimental to companies with an employee-based supply model. The amount of time and training invested in each individual is extensive and requires that each driver stay a relatively long period of time with the company in order to obtain a return on that investment.

In a rideshare system that supplies both pre-scheduled and on-demand rides, it is beneficial to hire employees, not just contractors, that the system can rely on to supply rides. On-demand ride-hailing services, such as Uber, allow customers to call upon drivers at will. Drivers then have the ability to deny service to a customer for any reason. When both pre-scheduled rides and on-demand rides are provided, an administrative entity is necessary in order to ensure that the supply is available to provide the promised rides. This requires having employed drivers who you can count on to provide the pre-scheduled rides in addition to contracted drivers who can provide on-demand rides. Again, this study aims to determine how we might add contracted, "on-demand" drivers to a driver pool in order to maximize profits from both pre-scheduled and on-demand rides.

1.1.4. Benefits of employees over contractors

In the case of standard employees, there has been a push in recent decades for the implementation of high-performance work strategies (HPWS) [8]. Huselid writes: "An increasing body of work contains the argument that the use of HPWS, including comprehensive employee recruitment and selection procedures, incentive compensation and performance management systems, and extensive employee involvement and training, can improve the knowledge, skills, and abilities of a firm's current and potential employees, increase their motivation, reduce shirking, and enhance retention of quality employees while encouraging non performers to leave the firm." In an increasingly competitive job market, these HR strategies have been crucial in order to attract and retain talented individuals.

Huselid cites many HPWS implementation studies, and in all cases, finds that HPWS has a net positive impact on the company [8]. Human resource management (HRM) practices were found to benefit companies in three main ways: reduced turnover, increased productivity, and improved corporate financial performance [8]. In particular, the factors affecting turnover include perceptions of job security, organizational tenure, number of dependents, and whether a job meets an individual's expectations [8]. Productivity was found to depend, in part, on training and goal-setting procedures put in place by the company [8]. And lastly, many studies support that HRM practices have a substantial and positive impact on financial performance. Specifically, Terpstra and Rozell found a significant and positive link between extensive recruiting practices and firm profits, among other things [8].

To investigate the effects of mixing standard and contracted employees in the workplace, Huselid poses the following question: Are HPWS more effective when used in contexts where a contingent (i.e., contracted) workforce is also deployed, or in contexts of a uniform culture where the entire workforce consists of standard employees? Huselid's study examined HPWS data provided by 229 British firms through the WERS2004 Cross Section Management Questionnaire, which is completed through face-to-face interviews with the manager responsible for HR at the workplace. His study found that contracted employment contributes positively to overall employment for companies where below-average levels of high-performance work practices are in place. However, contracted employment negatively impacts overall performance in organizations with high levels of high-performance work practices. In even simpler terms, a company that actively implements HRM practices is negatively impacted by hiring contracted employees. Unfortunately, hiring contracted workers within a tight-knit organization can adversely impact that trust and loyalty of the company's employees [9].

1.1.5. Combining the advantages of an employee with the convenience of a contractor

In the case of rideshare services for kids, the benefits of standard employment and HPWS practices are endless. The sensitive, highly contextual, and mission-oriented nature of a driver's job requires that the employee be committed and closely connected to the company. Ongoing training and performance reviews are a must. As proven during my time working for a small rideshare startup named Sheprd, the logistical complexities of child transportation are much greater than that of "traditional" rideshare for adults. Even employees with past experience as bus drivers require extensive training in order to learn the ins and outs of the service. Having worked closely with the drivers throughout their tenure at Sheprd, I witnessed firsthand how the employees became better and better at their job over time. Drivers are closely monitored and mentored when they first join the company. Over time, they earn autonomy and are then able to offer their own insights to new hires. This benefit is not afforded by the use of contractors who are expected to masterfully complete the job at hand upon start of their contract.

Given the benefits of reduced turnover, increased productivity, and improved corporate financial performance, we conclude that a child rideshare company should employ standard employees and not contactors. The ideal supply model should define a subset of employees that are assigned to pre-scheduled requests, while the remaining employees are reserved to handle the on-demand rides that would traditionally be served by contractors. The focus this study is then on the following question: in order to maximize profits, what proportion of employees should be pre-scheduled vs. on-demand for a child ridesharing company that offers both pre-scheduled and on-demand rides?

1.2. Objectives

1.2.1. Optimal split between pre-scheduled and on-demand drivers

This study primarily aims to identify the impact on profit made by adding on-demand drivers to the supply pool for a child rideshare service. Due to the pre-planned and routine nature of many children's schedules, most transportation options for children are static. For example, school bus routes are typically designed at the beginning of the school year and do not change thereafter. However, family lives can be hectic: plans change, kids get sick, and commitments fall through. The ability to book ad-hoc rides for children is of increasing importance as the proportion of households with two working parents continues to rise [10]. The child rideshare company, Sheprd, upon which this project was based, allowed for parents to book rides in advance as well as on the day of the ride. However, Sheprd did not specifically set aside drivers to serve these day-of (or "on demand") requests. Drivers' schedules were created one night in advance, taking into account all of the ride requests that had been received up until that point (typically 8pm). Drivers received their route assignment that same night, as soon as the schedule was set. Any ride requests that came in after 8pm the night before would be accepted if the ride requests could be added to existing routes already taking place that day (i.e., not impacting the start and end times of a driver's shift). If the ride requests would cause a driver to start earlier or end later than his pre-scheduled ride route, we would ask them if they were okay with the proposed schedule change. If accepted by the driver, we would then incorporate the ride requests that did not initially fit into an existing route. As Sheprd employees began to notice an increase in the proportion of on-demand ride requests, we considered whether the addition of dedicated on-demand drivers would have a positive impact on profits.

This study aims to identify whether dividing the driver pool between those who serve pre-scheduled rides and those who served on-demand rides could be financially beneficial for the company. Specifically, given a fixed number of drivers, we aim to determine the proper proportion of the driver pool that should be set aside to handle on-demand requests.

1.2.2. Impact of location layout

As a secondary objective, this study aims to identify the impact of spatial distribution of ride requests on overall profit. Assuming a fixed coverage area, a fixed number of locations, and a fixed number of ride requests, we present the following three location layouts: (1) uniform, (2) neighborhoods, and (3) city center. Within each layout, we have designated one-fourth of the locations as "schools", another one-fourth as "after-school activities", and the remaining one-half as "home" locations. In the uniform layout, we have mapped all locations onto a square grid, ensuring that schools and after-school locations are evenly distributed throughout the grid (see figure 1) . In the neighborhood layout, we have divided all locations into four equal groups (or neighborhoods), ensuring that location types (home, school, and after-school) are evenly distributed across neighborhoods. The exact position of each location is calculated such that the latitude and longitude of the locations are selected at random from a continuous uniform distribution between the latitudinal and longitudinal boundaries of the neighborhood it belongs to (see figure 2). And lastly, we have the city center layout which places all school and after-school locations in the center and places all homes around the perimeter of the coverage area. The exact positions of the non-home and home locations are chosen such that the latitude and longitude of the locations are selected at random from a continuous uniform distribution between the latitudinal and longitudinal boundaries of the center and the perimeter, respectively (see figure 3).

In all cases, the temporal distribution of demand is chosen at random from a continuous uniform distribution between the start and end time of operations. However, once the ride's start time is set, there is some logic that defines the location type of the origin and destination. For example, if a ride takes place within the first two hours of operation (i.e., first thing in the morning), then we assume that the ride is taking place between a home location and a school location. The home and school locations for this ride are then chosen at random from the set of home locations and school locations, respectively.

1.2.3. Impact of location count

The third and final objective of this study is to identify the impact of location count on profit. Location count is the number of physical locations from which customers are able to choose when designating the origin and destination of their request. Location count is independent of the number of ride requests. Assuming a fixed coverage area, a set location layout, and a fixed number of ride requests, we present three different location counts: 16, 64, and 256. As mentioned in the previous section, we have designated one-fourth of the locations as "schools", another one-fourth as "after-school activities", and the remaining one-half as "home" locations. In order to make sure that the location counts satisfy this location type designation, as well as our location layout definitions from section 1.2.2, we had to choose locations that were both divisible by 4 and were square numbers. Thus, we chose 16, 64, and 256 (the squares of 4, 8, and 16, respectively).

2. Project Design

2.1. Model formulation

2.1.1. High level model

To determine the impact of designating on-demand drivers within a driver pool, we formulate two separate models. The first of the two, the pre-scheduled ride assignment model, provides us with the maximum profit for a given number of pre-scheduled drivers, pre-scheduled ride requests, location count, and location layout. The second of the two, the on-demand ride assignment model, provides us with the maximum profit for a given number of on-demand drivers, on-demand ride requests, location count, and location layout. In order to determine the overall profit earned for a given number of total drivers, we take the sum of the objective values from these two models under a variety of scenarios. In each scenario, we vary the number of each type of driver while ensuring that the total number of drivers, as well as other other inputs, remain the same. We define the following fixed variables: k_p as the number of pre-scheduled drivers, k_d as the number of on-demand drivers, n_p as the number of pre-scheduled ride requests, n_d as the number of on-demand ride requests, l as the number of locations, and α as the location layout. We then define our optimized profit, π, as the output of our high-level model, f, as follows:

$$\pi = f(k_p, k_d, n_p, n_d, l, \alpha) = f_p(k_p, n_p, l, \alpha) + f_d(k_d, n_d, l, \alpha)$$

where π_p and π_d are the outputs of the pre-scheduled ride assignment model and the on-demand ride assignment model, respectively.

Each of these models solves a version of the *Dial-a-Ride-Problem* which consist of designing vehicle routes and schedules for n users who specify pickup and drop-off requests between origins and destinations [11]. Whereas Cordeau and Laporte assume that all ride requests are fulfilled and present a model that minimize costs, we remove the ride fulfillment constraint and set out to maximize profits.

2.1.2. *Pre-scheduled ride assignment model*

We consider a rideshare service through which caregivers book rides for their children. The optimization problem consists of assigning rides to drivers such as to maximize profit. We refer to this problem as the pre-scheduled ride assignment problem. The model formulated here is nearly identical to that of Cordeau [12]. The only difference is that our model has removed two of the constraints presented by Cordeau: the maximum route time and the maximum passenger ride time [12]. In this formulation, we only deal with ride requests that are received at least one day in advance of the trip. To formulate the problem, we introduce the following notation:

Model parameters:

n	number of ride requests
P	set of pickup vertices, P_d={1,2,...,n}
D	set of drop-off vertices, D_d=(n+1,n+2,...,2n}
L	set of physical locations (each vertex is mapped to a physical location)
V	vertex set, V={0, 2n+1, P, D} where 0 and 2n+1 are start and end depots
A	arc set, $A = \{(i,j) : i = 0,\ j \in P\ ,\ or\ i,j \in P \cup D,\ i \neq j\ and\ i \neq n+j,\ or$ $i \in D,\ j = 2n+1\}$
$(i, n+1)$	ride request, with i being the pickup vertex and n+1 the drop-off vertex
Q_k	vehicle capacity
K	set of all vehicles (i.e., drivers)
q_i	load for request i (i.e., number of passengers being picked up at vertex i)

c_{ij}	cost to traverse arc (i,j)
t_{ij}	travel time for arc (i,j)
$[e_i, h_i]$	permissible arrival time window for vertex i
γ	price per ride
t_{init}	operation start time (i.e., beginning of work day)
t_{end}	operation end time (i.e., end of work day)

Model variables:

x_{ij}^k	binary variable indicating whether arc (i,j) is traversed by vehicle k **(decision variable)**
u_i^k	time at which vehicle k arrives at vertex i
w_i^k	load of vehicle k upon leaving vertex i
r_i^k	ride time of passenger i (of request i) in vehicle k

The pre-scheduled ride assignment problem is formulated as follows:

Maximize $$((\sum_{k\in K}\sum_{i\in P}\sum_{j\in V} \times \gamma) - (\sum_{k\in K}\sum_{i\in V}\sum_{j\in V} \times c_{ij})) \qquad (1)$$

subject to

$$\sum_{i\in V} x_{0i}^k = \sum_{i\in V} x_{i,2n+1}^k = 1 \qquad (k \in K) \qquad (2)$$

$$\sum_{j\in V} x_{ij}^k - \sum_{j\in V} x_{n+i,j}^k = 0 \qquad (i \in P, k \in K) \qquad (3)$$

$$\sum_{j\in V} x_{ji}^k - \sum_{j\in V} x_{ij}^k = 0 \qquad (i \in P \cup D, k \in K) \qquad (4)$$

$$u_j^k \geq (u_i^k + t_{ij})x_{ij}^k \qquad (i, j \in V, k \in K) \qquad (5)$$

$$w_j^k \geq (w_i^k + q_j)x_{ij}^k \qquad (i, j \in V, k \in K) \qquad (6)$$

$$r_i^k \geq u_{n+1}^k - u_i^k \qquad (i \in P, k \in K) \qquad (7)$$

$$e_i \leq u_i^k \leq l_i \qquad (i \in V, k \in K) \qquad (8)$$

$$max\{0, q_i\} \leq w_i^k \leq min\{Q_K, Q_K + q_i\} \qquad (i \in V, k \in K) \qquad (9)$$

$$x_{ij}^k = \{0, 1\} \qquad (i, j \in V, k \in K) \qquad (10)$$

The objective function represents the expected profit for a given ride assignment matrix, x. It is defined as the difference between the expected revenue and the costs. Constraint (2) guarantees that the each driver starts and ends their route at the depot. Constraint (3) ensures that both the pickup and drop-off location are visited for each request that is served. Constraint (4) ensures that each driver departs from every location at which it arrives (excluding the depots). Constraint (5) defines the arrival time as being greater than or equal to the arrival time at the beginning of the arc plus travel time across the arc. Constraint (6) defines the load upon leaving vertex j as being greater than or equal to the load at vertex i plus the load associated with end vertex j. Constraint (7) defines the ride time of passenger i in vehicle k as being greater than or equal to the time of arrival at the destination (vertex $i+n$) minus time of arrival at the origin (vertex i). Constraint (8) ensures that the arrival time at vertex i lies within the permissible arrival time window for vertex i. Constraint (9) ensures that the load upon leaving vertex i is greater than or equal to the max of 0 and the load at i. Constraint (9) also ensures that the load upon vertex i is less than or equal to the minimum of the vehicle capacity and the vehicle capacity plus the load at vertex i. Lastly, constraint (10) defines x as a binary variable.

2.1.3. *On-demand ride assignment model*

We now consider a rideshare service through which caregivers book "on-demand" rides for their children. For our purposes, an "on-demand" ride request is one that requests a trip to take place on the same day the request is made. This optimization problem also consists of assigning rides to drivers such as to maximize profit. We will refer to this problem as the on-demand ride assignment problem. The model formulated here is nearly identical to that of Cordeau [12]. The only difference is that our model has removed two of the constraints presented by Cordeau and has added two additional parameters [12]. We have removed the maximum route time and the maximum passenger ride time, and we have added hourly wage and on-demand surcharge. These two additional parameters are used to calculate revenue in the objective function of the on-demand ride assignment model. In this formulation, we only deal with ride requests that are received on the same day of the desired trip. To formulate the problem, we introduce the following notation:

Model parameters:

n	number of ride requests
P	set of pickup vertices, P_d={1,2,...,n}
D	set of drop-off vertices, D_d=(n+1,n+2,...,2n}
L	set of physical locations (each vertex is mapped to a physical location)
V	vertex set, V={0, 2n+1, P, D} where 0 and 2n+1 are start and end depots
A	arc set, $A = \{(i,j) : i = 0,\ j \in P,\ or\ i,j \in P \cup D,\ i \neq j\ and\ i \neq n+j,\ or$ $i \in D,\ j = 2n+1\}$
$(i, n+1)$	ride request, with i being the pickup vertex and n+1 the drop-off vertex
Q^k	vehicle capacity

K set of all vehicles (i.e., drivers)

q_i load for request i (i.e., number of passengers being picked up at vertex i)

c_{ij} cost to traverse arc (i,j)

t_{ij} travel time for arc (i,j)

$[e_i, h_i]$ permissible arrival time window for vertex i

w hourly driver wage

s on-demand surcharge

γ price per ride

t_{init} operation start time (i.e., beginning of work day)

t_{end} operation end time (i.e., end of work day)

Model variables:

x_{ij}^k binary variable indicating whether arc (i,j) is traversed by vehicle k

(decision variable)

u_i^k time at which vehicle k arrives at vertex i

w_i^k load of vehicle k upon leaving vertex i

r_i^k ride time of passenger i (of request i) in vehicle k

The on-demand ride assignment problem is formulated as follows:

$$\text{Maximize } \left(\left(\sum_{k \in K} \sum_{i \in P} \sum_{j \in V} \times \gamma \times s \right) - \left(\sum_{k \in K} \sum_{i \in V} \sum_{j \in V} \times c_{ij} \right) - \left(K \times (t_{end} - t_{init}) * w \right) \right) \tag{1}$$

subject to

$$\sum_{i \in V} x_{0i}^k = \sum_{i \in V} x_{i,2n+1}^k = 1 \qquad (k \in K) \tag{2}$$

$$\sum_{j \in V} x_{ij}^k - \sum_{j \in V} x_{n+i,j}^k = 0 \qquad (i \in P, k \in K) \tag{3}$$

$$\sum_{j \in V} x_{ji}^k - \sum_{j \in V} x_{ij}^k = 0 \qquad (i \in P \cup D, k \in K) \qquad (4)$$

$$u_j^k \geq (u_i^k + t_{ij})x_{ij}^k \qquad (i, j \in V, k \in K) \qquad (5)$$

$$w_j^k \geq (w_i^k + q_j)x_{ij}^k \qquad (i, j \in V, k \in K) \qquad (6)$$

$$r_i^k \geq u_{n+1}^k - u_i^k \qquad (i \in P, k \in K) \qquad (7)$$

$$e_i \leq u_i^k \leq l_i \qquad (i \in V, k \in K) \qquad (8)$$

$$max\{0, q_i\} \leq w_i^k \leq min\{Q_K, Q_K + q_i\} \qquad (i \in V, k \in K) \qquad (9)$$

$$x_{ij}^k = \{0, 1\} \qquad (i, j \in V, k \in K) \qquad (10)$$

The objective function represents the expected profit for a given ride assignment matrix, x. It is defined as the difference between the expected revenue and the costs. The on-demand ride assignment problem builds upon the pre-scheduled ride assignment problem by introducing a surcharge parameter, s, in the first term of the objective function. This represents the price markup on rides that is incurred as a result of making the request on the same day of the ride. In the on-demand problem formulation, the objective function has a third term that calculates the total driver wages paid. In contrast to the drivers who serve pre-scheduled rides and are only paid for time spent serving a request, our on-demand drivers are paid for all hours of operation.

2.1.4. Model parameters

The bulk of the model parameters used in this study are based off of the operations of the Newton-based rideshare company, Sheprd [13].

The arrival time window for each request origin is 10 minutes. This means that if a ride is booked for 8:00 AM (let's call this request A), the arrival time window is 8:00-8:10 AM. The arrival time window for each request destination is 40 minutes. The beginning of the window is equal to the beginning of the origin's arrival time window *plus* the time between the origin and destination. In other words, if if take 20 minutes to get from request A's origin to request A's destination, then the arrival time window for the destination would be 8:20-9:00 AM. This larger time window on the destination end of the route allows for carpooling to take place between pickup and drop-off of request A.

The base price per ride is equal to $17. This price is fixed and is not dependent on ride length, ride distance, or ride time. The total price does depend, however, on when the ride is requested. If the ride is requested at least one day in advance, the base price is charged. If the ride is requested on the same day that the ride will take place, them a premium (20%) is added to the base price.

Each vehicle's capacity is set to 5. This is based off of the 7-passenger SUVs used by Sheprd and assumes that passengers are only allowed in the back seats of the car.

All ride requests are constrained to an 8 hour time window (i.e., daily hours of operations) and are chosen at random from a continuous uniform distribution between the start and end time of operations. All pickup times must fall within this time window, though the drop-off times may end up falling outside of the time window.

2.1.5. Model assumptions

Our model assumes that there are three types of locations: schools, after-school facilities, and homes. We assume that one-fourth of all locations are schools, one-fourth are after-school facilities, and one-half are homes. The pickup time for a ride request determines the location type of the origin and destination. Specifically, if the pickup time is within hours 0 and 2 of operations, then we assume that the ride is taking place from home to school. If the pickup time in between hours 2 and 4 of operations, then we assume that the ride is taking place from school to after-school facility.

If the ride is taking place between hours 4 and 6 of operations, then we assume that the ride is taking place between either school and after-school facility, school and home, or after-school activity and home. And lastly, if the ride is taking place between the last two hours of operations (hours 6 through 8), then we assume that the ride is taking place between after-school facility and home.

2.2. Model implementation

2.2.1. Gurobi

Both of the models formulated in this study were implemented in Gurobi, a commercial optimization solver [14]. With the use of the Gurobi Python interface, we were able to write the formulate the models within a Jupyter notebook (part of the Anaconda Python distribution). The code is organized as follows:

1) Set parameters
    a) Number of drivers
    b) Number of locations
    c) Number of requests
    d) Location layout
    e) Start of operating hours
    f) End of operating hours
    g) Price per ride
    h) Hourly operating cost
    i) Hourly driver wage
    j) On-demand surcharge
    k) Arrival time buffer

l) Ride time buffer

2) Define set of locations

a) Define subset of "home" locations

b) Define subset of "school" locations

c) Define subset of "after-school" locations

3) Assign coordinates to each location depending on the location layout:

a) Uniform (all locations equally spaced)

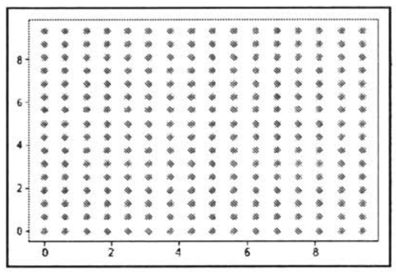

Figure 2-1: Example of uniform location layout.

b) Neighborhoods (locations are even split into 4 groups)

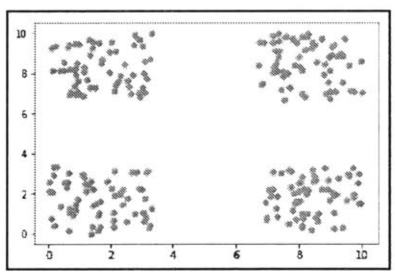

Figure 2-2: Example of neighborhood location layout.

c) City center (all school and after-school locations are at the center; all homes are located in the periphery)

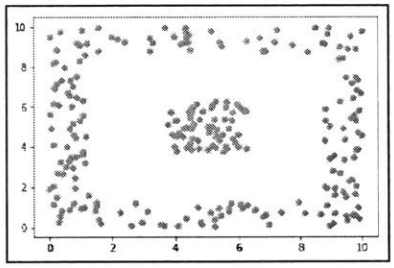

Figure 2-3: Example of city center location layout.

4) Define location distance matrix

    a) Translate matrix into dictionary

    b) Create dictionary that defines time between locations (proportional to distance matrix)

5) Define arrival time window for all pickup vertices

6) Define set of locations

    a) Define subset of "home" locations

    b) Define subset of "school" locations

    c) Define subset of "after-school" locations

7) Define dictionary that maps all vertices (both pickup and drop-off) to a location

8) Define dictionary for the load at each vertex

    a) All pickup vertices are set to a load of "1" (passenger is added)

    b) All drop-off vertices are set to a load of "-1" (passenger is removed)

9) Define arrival time windows for drop-off vertices

10) Create cost dictionary

    a) Uses time dictionary

11) Define model

    a) Both model formulations are identical in variable definition and constraint definition, but they use distinct objective functions.

The pre-scheduled ride assignment model and the on-demand ride assignment model are both implemented in Gurobi as they are presented above in section 2.1.2 and 2.1.3. However, since Gurobi cannot handle nonquadratic nonlinear constraints, constraints (6) and (7) were linearized according to recommendations in Cordeau (2006) [12].

3. Results and discussion

3.1. Location layout findings

We sought to identify the impact of location layout on profit. In order to do so, we looked at the expected profit from each location layout under a variety of scenarios. We focused on our discussion on the following LL (location layout) scenarios:

	Scenario LL1	*Scenario LL2*	*Scenario LL3*	*Scenario LL4*
Total number of requests (n)	Independent variable (x-axis)	Independent variable (x-axis)	Independent variable (x-axis)	Independent variable (x-axis)
Profit (π)	Dependent variable (y-axis)	Dependent variable (y-axis)	Dependent variable (y-axis)	Dependent variable (y-axis)
Location count (l)	16	16	16	16
Proportion of on-demand (OD) requests	50%	50%	0%	100%
Total number of drivers (k)	5	5	5	2
Proportion of on-demand (OD) drivers	20%	40%	0%	100%

Table 3.1: Location layout scenarios.

Figures 4 through 7 present profit as a function of the total number of ride requests. The number of requests is represented on the x-axis while the profit is represented along the y-axis. The solid line represents profit from the uniform location layout, the dotted line represents profit from the neighborhood location layout, and the dashed line represents profit from the city center layout.

In the first LL scenario (results illustrated in figure 4), we assume 50% OD requests and 20% OD drivers. We observe that the uniform location layout leads in profit when there are less than 24 total requests. With 16 total requests, the uniform layout produces a profit of $115.80 while the neighborhood and city center layouts are both unprofitable, each with losses of $19.86. At 24 total requests, all three location layouts produce profits between $94 and $100, but the uniform layout drops out of the lead once the total number of requests exceeds 24. After the 24-request mark, the city center layout leads profits, with neighborhood and uniform layout trailing behind, in that order. We observe a maximum profit of $472.13 under the city center layout when there are 48 total requests. At 48 total requests, the city center layout's profits are 37% higher than that of the uniform layout and 25% higher than that of the neighborhood layout.

In scenario LL2 (results illustrated in figure 5), we have defined 40% of the driver pool as OD. We observe that the city center location layout results in the highest profits, regardless of the number of ride requests. With a higher proportion of OD drivers than scenario LL1, we see that LL2 requires a larger number of requests in order to achieve profitability. Profits are positive for the uniform layout after the 24-request mark, but the other two layouts require at least 28 total requests to reach profitability. We observe a maximum profit of $299.84 under the city center layout when there are 48 total requests. At 48 total requests, the neighborhood layout's profits are 17% higher than that of the uniform layout and 11% higher than that of the neighborhood layout.

In scenario LL3 (results illustrated in figure 6), we have no OD requests or drivers. Since we only consider PS requests, the total number of requests ranges from 8 to 24 (as opposed to 16 to 48 in the previous two scenarios). Here, we observe that the city center location layout results in the highest profits until the 22-request mark, after which the neighborhood layout takes the lead. With no OD drivers, profitability is maintained throughout the entire range of ride requests. The rate at which profit increases with increased requests slows down for the city center layout at the 20-request mark while it actually speeds up for the neighborhood layout at this same point. This allows for the neighborhood layout to max out at a profit $298.83 with 24 requests, 40% higher than that of the uniform layout and 11% higher than that of the city center layout.

In the fourth and final LL scenario (results illustrated in figure 7), we have only OD requests and drivers. Since we only consider OD requests, the total number of requests ranges from 8 to 24 (as opposed to 16 to 48 in scenarios LL1 and LL2). Here, we observe that the city center location layout results in the highest profits (and smallest losses) throughout the range of requests. All three location layouts are unprofitable until reaching the 22-request mark. After this point, the city center layout takes the lead, maxing out at a profit of $85.19, 48% higher than that of the uniform layout and a whopping 401% higher than that of the neighborhood layout.

In almost all cases, the city center layout seems to lead with the highest profit. To explain this, we could argue that because the city center layout places all school and after-school locations close to one another, this layout is most conducive to carpooling. As a result, the city center layout increases the number of ride requests that can be served, increasing overall expected profit.

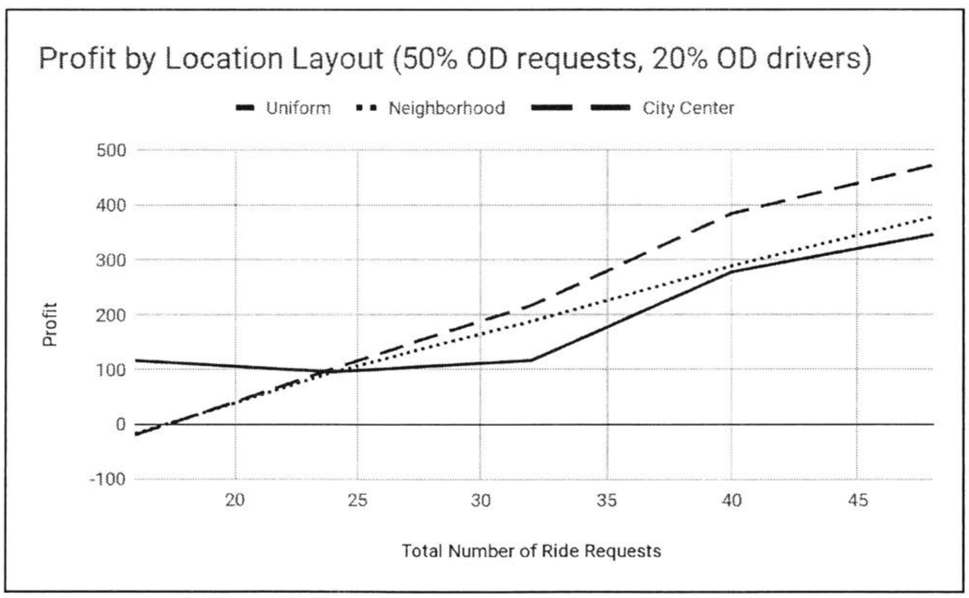

Figure 3-1: Profit by location layout. This scenario assigns an equal number of pre-scheduled and on-demand ride requests, as well as a driver pool with 20% on-demand drivers. The results show that city center layout produces a slightly higher overall profit after the 24 ride request mark.

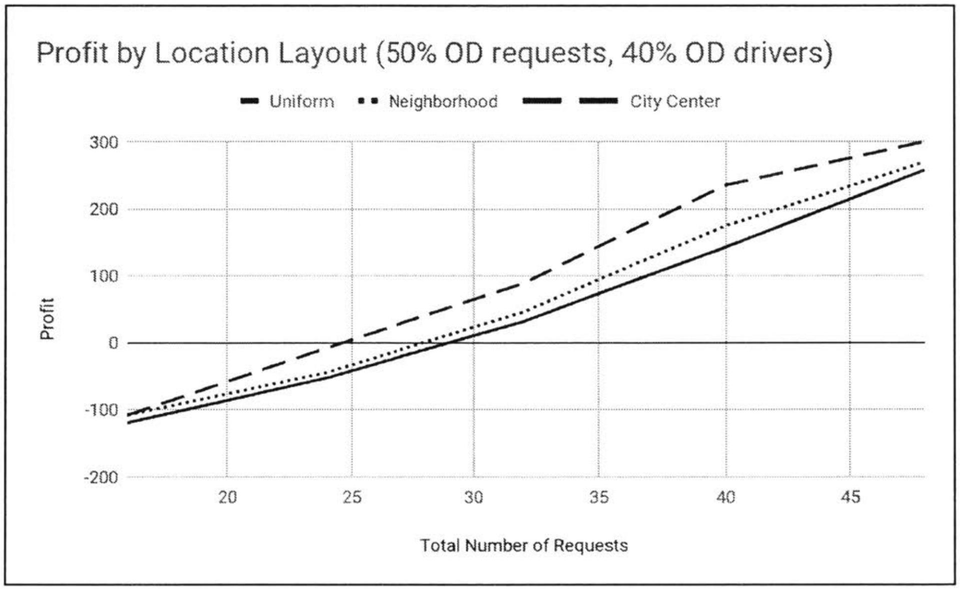

Figure 3-2: Profit by location layout. This scenario assigns an equal number of pre-scheduled and on-demand ride requests, as well as a driver pool with 40% on-demand drivers. The results show that city center layout consistently produces a slightly higher overall profit.

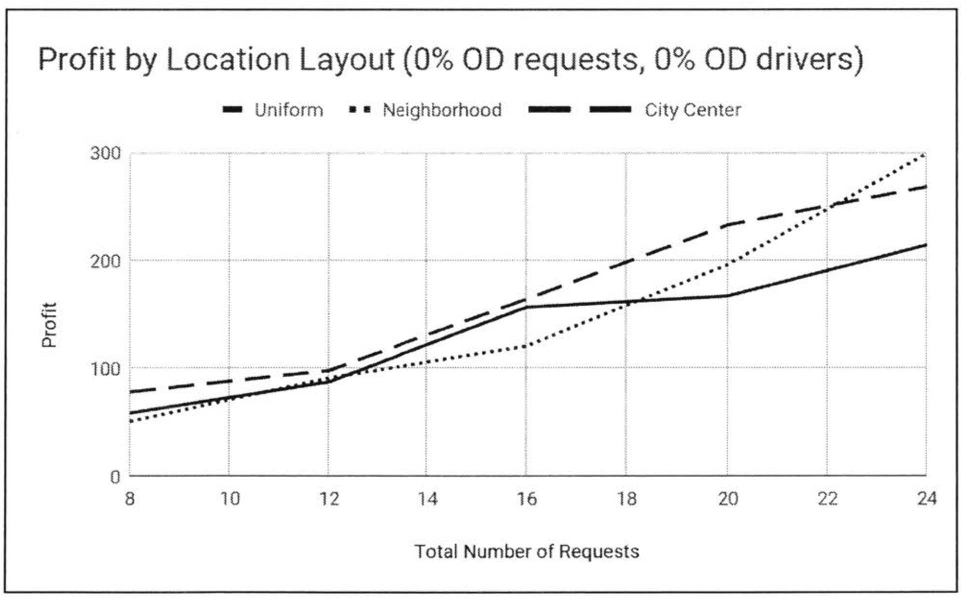

Figure 3-3: Profit by location layout. This scenario considers only pre-scheduled ride requests, as well as a driver pool with only pre-scheduled drivers. The results show that city center layout produces a slightly higher overall profit up until 22 requests. After this point, the neighborhood layout takes the lead.

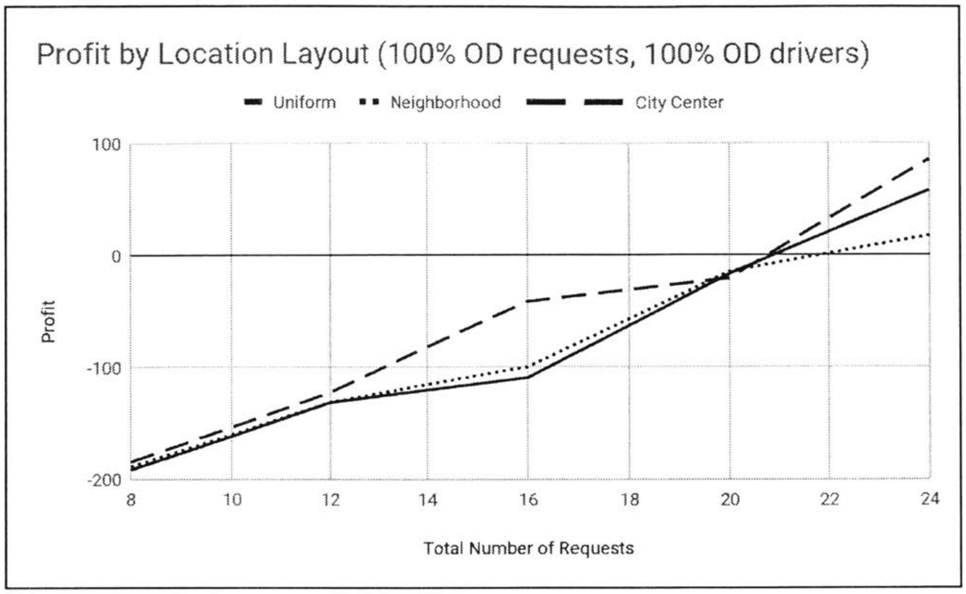

Figure 3-4: Profit by location layout. This scenario considers only on-demand ride requests, as well as a driver pool with only on-demand drivers. The results show that city center layout produces a slightly higher overall profit.

3.2. Location count findings

We sought to identify the impact of location count on profit. In order to do so, we looked at the expected profit from each location count under a variety of scenarios. We focused on our discussion on the following LC (location count) scenarios:

	Scenario LC1	*Scenario LC2*	*Scenario LC3*	*Scenario LC4*
Total number of requests (n)	Independent variable (x-axis)	Independent variable (x-axis)	Independent variable (x-axis)	Independent variable (x-axis)
Profit (π)	Dependent variable (y-axis)	Dependent variable (y-axis)	Dependent variable (y-axis)	Dependent variable (y-axis)
Location layout (α)	City center	City center	City center	City center
Proportion of on-demand (OD) requests	50%	50%	0%	100%
Total number of drivers (k)	5	5	5	2
Proportion of on-demand (OD) drivers	20%	40%	0%	100%

Table 3.2: Location count scenarios.

Figures 8 through 11 present profit as a function of the total number of ride requests. The number of requests is represented on the x-axis while the profit is represented along the y-axis. The solid line represents profit from the 16-location count, the dotted line represents profit from the 64-location count, and the dashed line represents profit from the 256-location count.

In the first LC scenario (illustrated in figure 8), we assume 50% OD requests and 20% OD drivers. We also assume a city-center layout for all LC scenarios. In LC1, we observe that the 64-location count leads profits for the first half of the request range. After the 32-request mark, the 16-location count results in a higher rate of profit increase and takes the lead. Though the rate of increase slows down at the 40-request mark, the 16-location count remains in the lead and maxes out at a profit of $472.13, 45% higher than that of 256-location count and 12% higher than that of 64-location count.

In scenario LL2 (illustrated in figure 9), we assume 50% OD requests and 40% OD drivers. Again, we assume a city-center layout. This scenario, which includes a higher proportion of OD drivers as compared to the last, does not result in profitability for any of the location counts until after the 24-request mark. After this point, the 16-location count takes the lead and pushes further and further into the lead as the total number of requests continues to increase. At its max, the 16-location count results in a profit of $299.84, 38% greater than that of 64-location count and 35% greater than that of 256-location count.

In the third LC scenario (illustrated in figure 10), we assume 0% OD requests and 0% OD drivers. Again, we assume a city-center layout. This scenario, which considers only PS requests and PS drivers, results in profitability throughout the entire request range. Since we are only considering PS requests, the total number of requests ranges between 8 and 24 (as opposed to a range of 16 to 48 when both OD and PS requests are considered). In this scenario, profits for the three location counts are very close to one another until around the 14-request mark. After this point, the 16-location count takes the lead and maxes out at a profit of $268.21, 23% greater than that of 256-location count and 9% greater than that of 64-location count. However, it is interesting to note that the 16-location count had the largest lead at the 20-request mark where its profits were $232.98, a whopping 45% above that of the 256-location count and 37% above that of the 64-location count. After the 20-request mark, the rate of profit increase slows down for 16-location count while it speeds up for 64 and 256-location count. Given the high rate of increase we observe for the 64-location count at the high end of our request range, it is possible that this location count will take the lead in profits as the number of requests continues to increase.

In the fourth and final LC scenario (illustrated in figure 11), we assume 100% OD requests and 100% OD drivers. Again, we assume a city-center layout. This scenario, which assumes only OD requests and drivers, does not exhibit profitability until after the 20-request mark. Even though we experience losses when there are fewer than 20 requests, the 16-location count results in the smallest losses throughout most of the range. Once we achieve profitability, the 16-location count comes out as the clear leader, maxing out at a profit of $85.19 (337% greater than that of 256 locations and 78% greater than that of 64 locations).

We set out to identify the impact of location count on overall profit. In the four location count scenarios described above, we observe that, in general, the smaller the location count, the higher the resulting profit. Though the results are mixed when request counts are low, each scenario seems to have its own "critical" number of requests after which the lowest location count comes out on top. For example, profits in scenario LC2 for the 16 and 64-location counts are nearly indistinguishable until the 32-request mark. After this point, the 16-location count is the clear frontrunner in terms of profit. We draw two main conclusions from our findings: (1) assuming a fixed number of requests, a smaller number of locations results in higher profits, and (2) there is a minimum number of requests that need to be received in order for the benefits of fewer locations to be realized. In the same way that a train or bus has limited stops but is only profitable if enough people use it, this service is most profitable with limited locations but only when demand exceeds some minimum level.

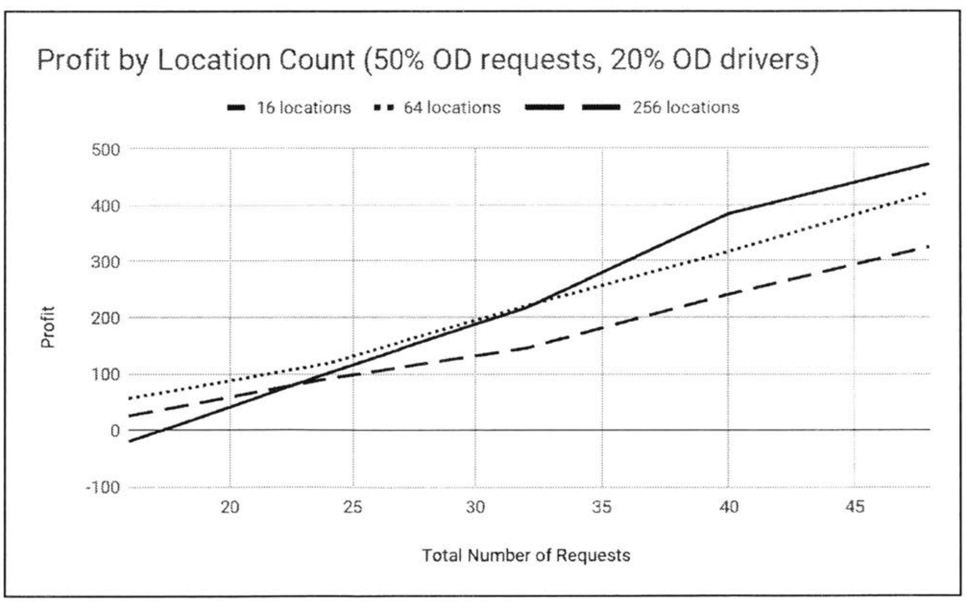

Figure 3-5: Profit by location count. This scenario assigns an equal number of pre-scheduled and on-demand ride requests, as well as a driver pool with 20% on-demand drivers. The results show the smallest location count (16) results in the highest profit once the number of ride requests exceeds ~32.

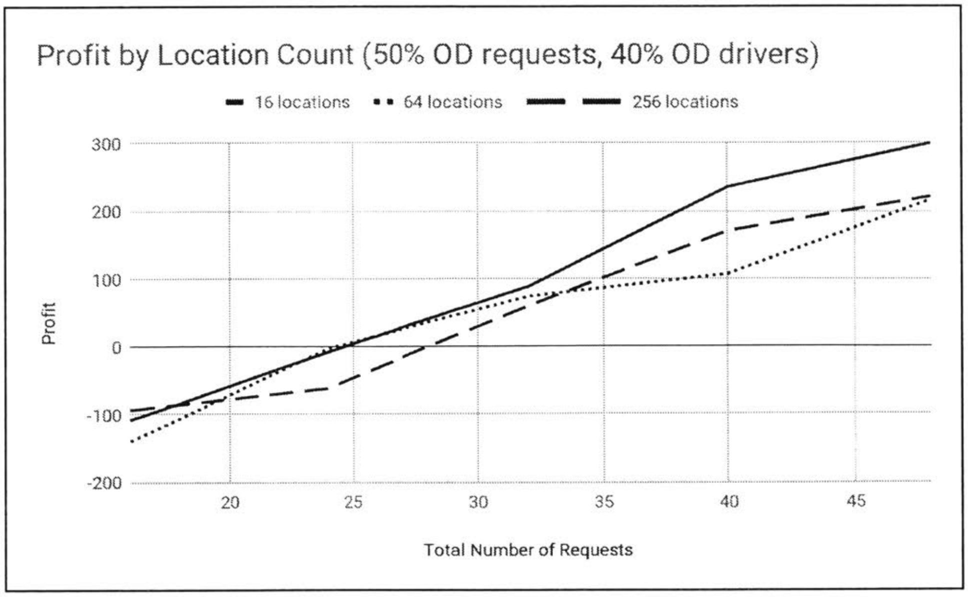

Figure 3-6: Profit by location count. This scenario assigns an equal number of pre-scheduled and on-demand ride requests, as well as a driver pool with 40% on-demand drivers. The results show the smallest location count (16) results in the highest profit.

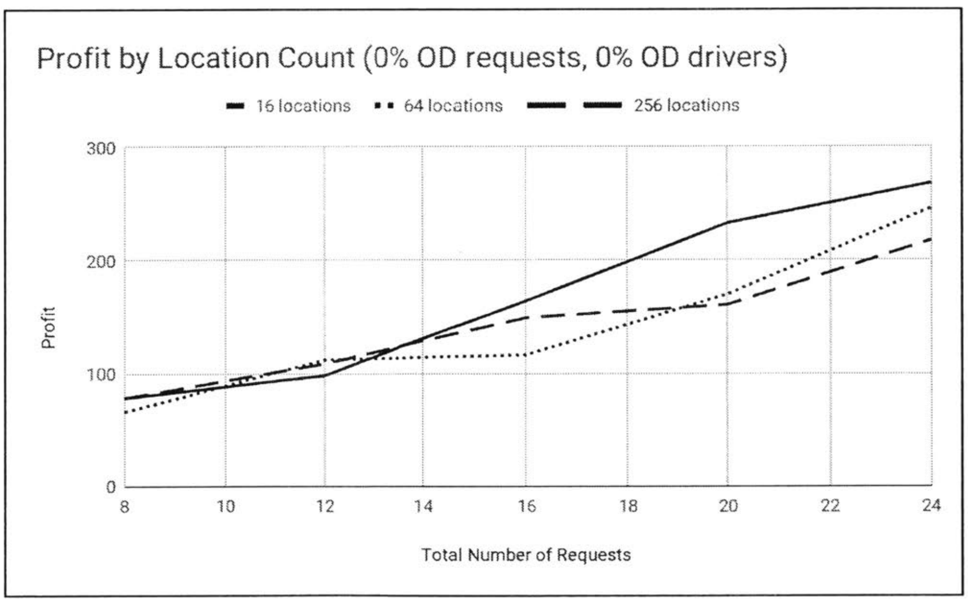

Figure 3-7: Profit by location count. This scenario assigns only pre-scheduled ride requests, as well as only pre-scheduled drivers. Differences in profit across location counts are minimal up until the 13 ride request mark. Beyond 13 ride requests, the smallest location count (16) produces a clear lead on profit.

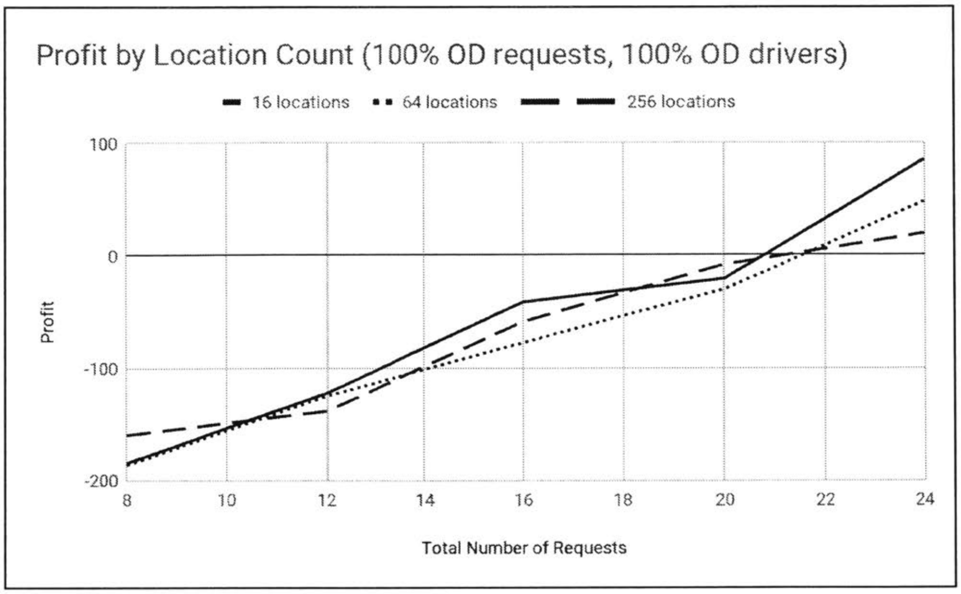

Figure 3-8: Profit by location count. This scenario considers only on-demand requests, as well as a driver pool with only on-demand drivers. Differences in profit across location counts are minimal but it seems that the smallest count (16) produces a clear lead on profit after we exceed ~20 ride requests.

3.3. Driver split findings

The main goal of this study was to identify the impact of adding on-demand drivers to a driver pool that was previously designed to serve pre-scheduled requests. We aim to understand whether the introduction of on-demand drivers made financial sense for the company. Our real-life example, Sheprd, permitted users to request on-demand (i.e., same-day) rides, but they only had the capacity to accept requests that could be adding to an existing route without negatively impacting arrival times of the stops already on the route.

In order to explore the impact of introducing on-demand drivers, we looked at the expected profit from each "driver split" under a variety of scenarios. In each scenario, we assume a total of 5 drivers and compare the impact of splitting the drivers up in one of the following 3 ways:

1. 0% OD drivers (i.e., 0 OD drivers and 5 PS drivers)
2. 20% OD drivers (i.e., 1 OD driver and 4 PS drivers)
3. 40% OD drivers (i.e., 2 OD drivers and 3 PS drivers).

We focused on our discussion on the following DS (driver split) scenarios:

	Scenario DP1	*Scenario DP2*	*Scenario DP3*
Total number of requests (n)	Independent variable (x-axis)	Dependent variable	32
Total number of PS requests (n_p)	Dependent variable	24	Dependent variable
Total number of OD requests (n_d)	Dependent variable	Independent variable (x-axis)	Dependent variable
Profit (π)	Dependent variable (y-axis)	Dependent variable (y-axis)	Dependent variable (y-axis)
Location layout (α)	City center	City center	City center
Location count (l)	16	16	16
Proportion of on-demand (OD) requests	50%	Dependent variable	Independent variable (x-axis)
Total number of drivers (k)	5	5	5

Table 3.3: Driver proportion scenarios.

Figures 12 presents profit as a function of the total number of ride requests. The number of requests is represented on the x-axis while the profit is represented along the y-axis. The solid line represents profit with 0% OD drivers, the dotted line represents profit with 20% OD drivers, and the dashed line represents profit with 40% OD drivers. This figure represents the results of scenario DP1.

Figures 13 presents profit as a function of the number of OD ride requests. The number of OD requests is represented on the x-axis while the profit is represented along the y-axis. The solid line represents profit with 0% OD drivers, the dotted line represents profit with 20% OD drivers, and the dashed line represents profit with 40% OD drivers. This figure represents the results of scenario DP2.

Figures 14 presents profit as a function of the percentage of OD ride requests. The percentage of OD ride requests is represented on the x-axis while the profit is represented along the y-axis. The solid line represents profit with 0% OD drivers, the dotted line represents profit with 20% OD drivers, and the dashed line represents profit with 40% OD drivers. This figure represents the results of scenario DP3.

In the first DS scenario (illustrated in figure 12), we assume 50% OD requests, a 16-location count and a city center layout. In this scenario, we observe the resulting profit as the total number of requests varies. When the total number of requests is small (less than 20), profitability is only achieved with the 0% OD driver split. However, expected profits for all three driver splits are seen to increase as the total number of requests increases, with the rate of increase for the 20% and 40% OD driver split being higher than that of the 0% OD driver split.. In fact, at the 24-request mark, the 20% OD driver split takes the lead and maintains the highest profit for all request counts in our range. Profits for the 20% OD driver split maxes out at $472.13, 76% greater than that of the 0% OD driver split and 57% greater than that of the 40% OD driver split.

In scenario DS2 (illustrated in figure 13), we assume 24 PS requests, a 16-location count and a city center layout. Here, we observe how profits are impacted by varying the number of OD requests, given a fixed number of PS requests. Profit is constant for the 0% OD driver split since the number of PS requests is fixed (and since the number of OD requests is irrelevant since there are no drivers available to serve them). Both the 20% and 40% OD driver splits are profitable at 8 OD requests. However, they both result in lower profits than the 0% OD driver series. Profits from the 20% and 40% OD driver splits are lower than that of the 0% driver split because the revenue from 8 OD ride requests is not enough to offset the fixed cost of paying these drivers to be available all day. As we increase the number of OD requests, we find that 20% OD driver split surpasses the 0% OD split at the 10-OD-request mark and remains in the lead throughout the remainder of the range. The

10th OD request marks the point at which it becomes more profitable to employ 1 on-demand driver than to not employ any. Once the 20% OD split take the lead at the 10-request mark, the 0% OD driver split remains in second place in terms of profit until the 23rd OD request mark. At the 23rd OD request mark, profits from the 40% OD driver split exceed that of the 0% OD driver split. The 23rd OD requests marks the point at which it becomes more profitable to employ two on-demand drivers than to not employ any. We can assume that there is a number of OD ride requests after which it makes more sense to employ 2 OD drivers as opposed to 1, but we are not able to determine that within the scope of our analysis.

In the third and final DS scenario (illustrated in figure 14), we assume a total of 32 ride requests, a 16-location count and a city center layout. In this scenario, we measure the profit that results from varying the proportion of OD requests. Starting with the smallest proportion of OD requests (25%), we find that the 0% OD driver split results in the highest overall profit of $268.21, 8x greater than that of the 40% OD driver split and 25% greater than that of the 20% OD driver split. As the proportion of OD requests increases, that initial profit of the 0% OD driver split ($268.21) is never surpassed. Profits for the 0% OD driver split steadily decline as the percentage of OD requests increases (this is expected since the number of PS requests are declining along with it). Profits for both the 20% and 40% OD driver splits are increasing between the 25% and 40% OD-request marks while profits for the 0% OD driver split head downwards. The 20% OD driver split actually surpasses the 0% OD driver split around the 35% OD-request mark and maintains higher profits throughout. Profit for both the 20% and 40% OD driver splits take a downward turn between the 40 and 50% OD-request marks before returning to a positive rate of change at the 50% OD-request mark. The 40% OD driver split surpasses the 0% OD driver split profit at around the 63% OD-request mark, securing it in second for the remainder OD request range. At the 75% OD-request mark, the 20% OD driver split results in the highest profit of $236.88, 203% greater than that of the 0% OD driver split and 46% greater than that of the 40% OD driver split. Though we observe the highest overall profit when there are no OD drivers, this does not necessarily mean that this is the smartest option. We have to consider the fact that we are rejecting a large number of ride requests (i.e., all of the OD ride requests), and predict the impact that those declinations could have on future demand. We did not assign penalties to unfulfilled ride requests in this model, but doing so in the future could correct for this and provide us with a more accurate understanding of the impacts of each scenario.

The scenarios described above lead us to conclude that for a ride assignment problem of the scale we have presented, a driver pool composed of 20% OD drivers is likely to increase overall profits. All three DS scenarios supported the conclusion that 20% OD drivers is the ideal split. First, we saw highest profits with 20% OD drivers in scenario DS1 once the total number of requests exceeded 24 (figure 12). Then, we saw highest profits with 20% OD drivers in scenario DS2 once the number of OD requests exceeded 10 (figure 13). And lastly, we saw profits take the lead with 20% OD drivers in scenario DS3 once the proportion of OD requests exceed ~35% (figure 14).

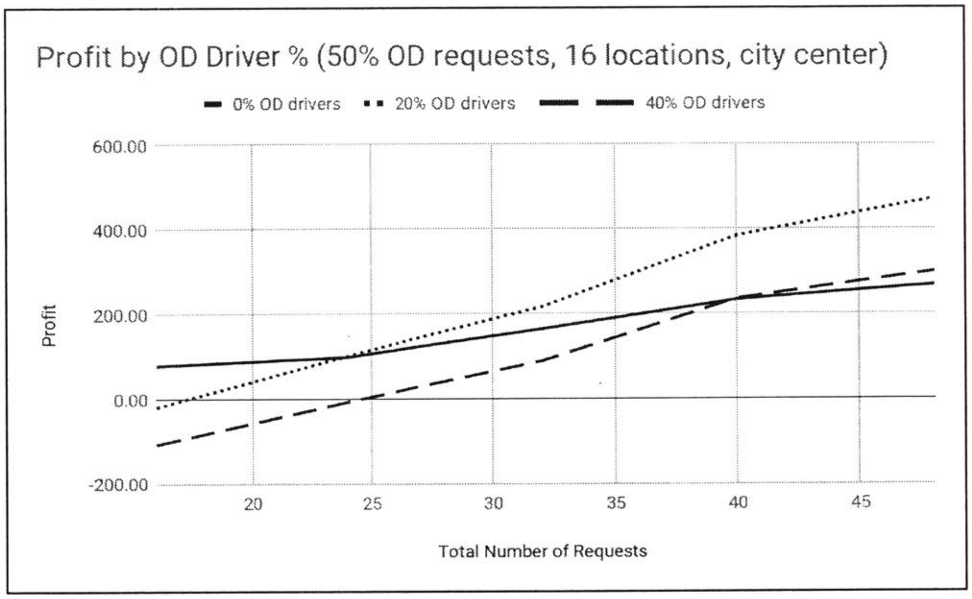

Figure 3-9: Profit by proportion of on-demand drivers in the driver pool. This scenario assigns an equal number of pre-scheduled and on-demand requests. Here, we have set the location count to 16 and the location layout to city center. After the 24 ride request mark, we see that the 20% on-demand driver scenario results in the highest profit.

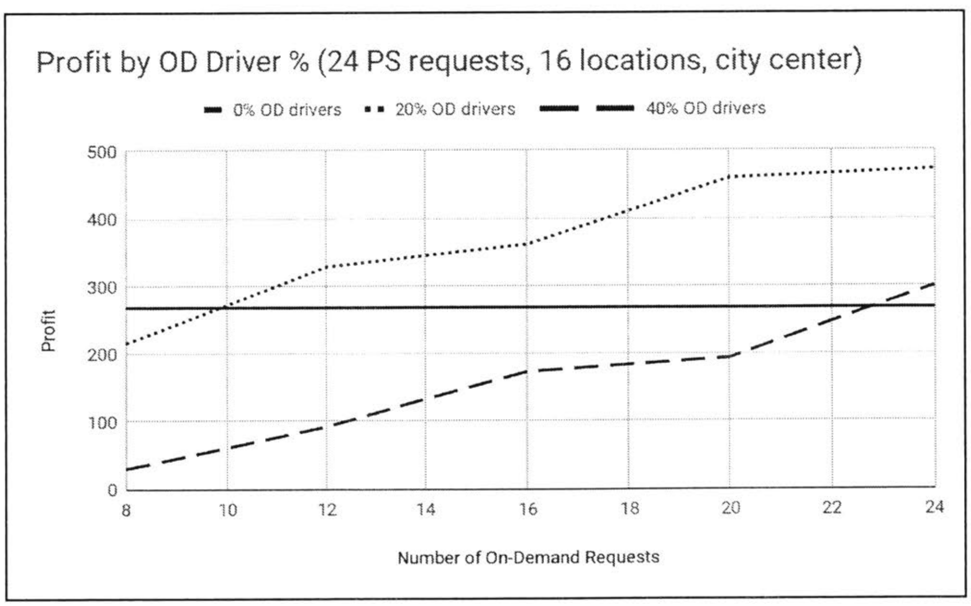

Figure 3-10: Profit by proportion of on-demand drivers in the driver pool. This scenario assumes 24 pre-schedule requests while varying the number of on-demand requests. Here, we have set the location count to 16 and the location layout to city center. We see that once we have at least 10 on-demand requests, the highest profit is achieved with a 20% on-demand driver pool.

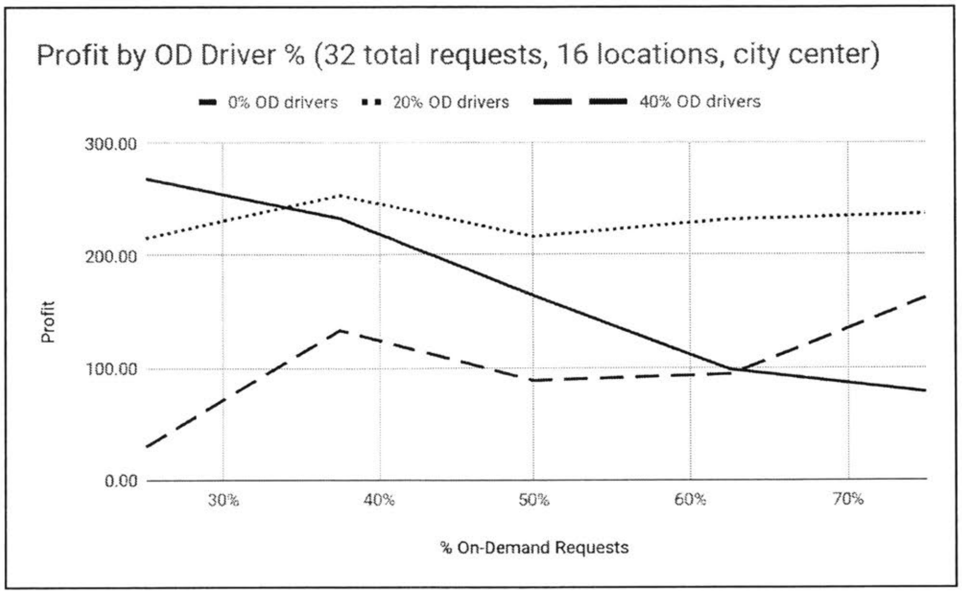

Figure 3-11: Profit by proportion of on-demand drivers in the driver pool. This scenario assumes a total of 36 ride requests while varying the proportion of on-demand requests. Here, we have set the location count to 16 and the location layout to city center. We see that once we on-demand requests make up 35% of the total request, the highest profit is achieved with a 20% on-demand driver pool.

3.4. Limitations

The main limitation encountered in this study was a lack of computational power. The ride assignment models implemented in Gurobi were too complex to run at the size of Sheprd operations (roughly 30 drivers and 150 ride requests). As a result, we were limited to small test scenarios with a maximum of 5 drivers and 32 ride requests. These were chosen as the upper bounds since the model was able solve the optimization problem with these inputs in roughly 1 minute or less. As a result of this limitation, we were not able to run the model using Sheprd ride data to see how the introduction of on-demand drivers would have affected profits. Instead, we ran smaller test scenarios to get a sense of how the introduction of on-demand drivers may impact overall profit.

Another limitation of the model is oversimplification of our ride requests. While we do take into account the time of day when choosing pickup and dropoff locations (i.e., rides that take place first thing in the morning are only from home to school), there are many other patterns we do not consider. For example, for passengers being picked up from school and taken to an after-school activity, it is likely that many students of the same age or from the same school are going to the same after-school activity. This would result in an increased likelihood of carpool (and greater profits for the company). However, our model assigns all pickups and dropoffs by selecting at random from the set of all possible locations for a given pickup time, without taking into account the likelihood that students being picked up in one place are more likely to be going to the same place (or nearby places). Similarly, since school zones often determine the school that a student attends, it is likely that students picked up for school in the same area would be attending the same school. Our model does not take these probabilities into account.

3.5. Further research recommendations

I would encourage further research to implement heuristics in this model in order to allow us to run the model at a larger scale. This would allow us to model operations at a more realistic size that would provide more meaningful insight into the impact of incorporating on-demand drivers into the driver pool.

Furthermore, I would take a closer look to find the peak time of on-demand requests. This way, we could strategically employ on-demand drivers to fulfill ride requests during the times when demand is highest for last-minute requests. From anecdotal experience at Sheprd, I would posit that the school pickup hour sees the highest number of on-demand requests. If we were to conclude that this is the case, we could schedule on-demand drivers to only handle the school pickup rides, thus lowering the fixed cost of employing on-demand drivers and lowering the ride request barrier for turning a profit.

Lastly, our results could be further explored by testing various size coverage areas. For all of our scenarios, we assumed a 10x10 degree (latitude, longitude) area. Future studies could use the dimensions of actual cities in order to more accurately estimate profits.

4. Conclusion

In the course of this study, we built a pair of models to estimate the profits earned by a child rideshare company under a variety of scenarios. We define pre-scheduled requests as those that are received at least one day in advance of the ride. We define on-demand requests as those received on the same day of the ride. We define pre-scheduled drivers as those who fulfill pre-scheduled ride requests. And we define on-demand drivers as those who fulfill on-demand ride requests.

We aimed to identify the impact of three different variables: location layout, location count, and proportion of on-demand drivers. We found that, in general, the city center layout resulted in the highest profit for the rideshare service. In this layout, all school and after-school locations where located in the center of the coverage area while all home locations were located in the periphery. In terms of location count, we found that, in general, a smaller location count results in the highest rideshare profit. And lastly, we found that as long as roughly 35% of total ride requests are on-demand, we achieve the highest profit when 20% of our driver pool is on-demand.

References

[1] "Requests from underage riders," www.uber.com, (2019).
[2] Joseph P Broschak and Alison Davis-Blake, "Mixing standard work and nonstandard deals: The consequences of heterogeneity in employment arrangements," *Academy of Management Journal* vol. 49, issue 2 (April 2006): 371-393.
[3] Luigi Stirpe, Jaime Bonache, and Antonio Revilla, "Differentiating the workforce: The performance effects of using contingent labor in a context of high-performance work systems," *Journal of Business Research* vol. 67 (2014): 1334–1341.
[4] Arne L. Kalleberg, "Organizing Flexibility: The Flexible Firm in a New Century," *An International Journal of Employments Relations* vol. 39, issue 4 (December 2001): 479-504.
[5] Greg Kumparak, "Shuddle, the Uber-like service for getting your kids around, is shutting down tomorrow," http://www.techcrunch.com, (2016).
[6] Private email sent from CEO Juliette Kayyem to Zemcar customers on September 20, 2018.
[7] Russell Hensley, Asutosh Padhi, and Jeff Salazar, "Cracks in the ridesharing market--and how to fill them," *McKinsey Quarterly* vol. 3 (2017): 48-55.
[8] Mark A. Huselid, "The Impact of Human Resource Management Practices on Turnover, Productivity, and Corporate Financial Performance," *Academy of Management Journal*, vol. 38, no. 3 (1995): 635–872.
[9] Jone L. Pearce, "Toward an Organizational Behavior of Contract Laborers: Their Psychological Involvement and Effects on Employee Co-Workers," *The Academy of Management Journal* vol. 36, no. 5 (October, 1993): 1082-1096.
[10] "Raising Kids and Running a Household: How Working Parents Share the Load," *Pew Research Center,* (November 4, 2015).
[11] Jean-François Cordeau and Gilbert Laporte, "The dial-a-ride problem: models and algorithms," *Annals of Operations Research* vol. 153 (May 2007): 29-46.
[12] Jean-François Cordeau, "A Branch-and-Cut Algorithm for the Dial-a-Ride Problem," *Operations Research* vol. 54, issue 3 (2006): 573-586.
[13] www.sheprd.com
[14] www.gurobi.com