

THE EARTH INSTITUTE COLUMBIA UNIVERSITY

TRANSFORMING PERSONAL MOBILITY

Lawrence D. Burns, Director, Program on Sustainable Mobility

William C. Jordan, President, Jordan Analytics LLC

Bonnie A. Scarborough, Program Manager, Program on Sustainable Mobility
The Earth Institute, Columbia University

August 10, 2012

ABSTRACT

This paper presents three distinctly different case studies to assess the personal mobility system that results from **combining**:

- The “Mobility Internet”¹
- Self-driving/driverless vehicles
- Shared vehicle systems
- Specific-purpose vehicle designs
- Advanced propulsion systems

Results indicate that this transformational mobility system:

- Is capable of supplying **better mobility experiences at radically lower cost** under a wide range of circumstances
- Offers **substantial sustainability benefits** through improved roadway safety, reduced roadway congestion, increased energy efficiency, reduced emissions, improved land use and enhanced equality of access.

Models that generalize the relationship between the characteristics of geographical regions and the cost and performance of coordinated/shared/driverless mobility systems are also provided.

¹ The term “Mobility Internet” was first coined in William J. Mitchell, Christopher Borroni-Bird, and Lawrence D. Burns, Reinventing the Automobile: Personal Urban Mobility for the 21st Century. The MIT Press, 2010. The “Mobility Internet” does for the movement of people and goods what the Internet has done for the movement of information by coordinating large amounts of real-time spatial and temporal connectivity and infrastructure data.

INTRODUCTION

Global roadway transportation has evolved to an enormous scale since the first motorized vehicle was invented over a century ago. Today, nearly a billion cars and trucks move people and goods along the world's roadways and consumers spend trillions of dollars each year on personally owned vehicles (including the costs of fuel, depreciation, insurance, taxes, parking, and time) to experience the resulting mobility benefits.

The growth of automobile transportation has occurred with virtually no disruptive change to the fundamental system conceived of by Karl Benz and popularized by Henry Ford. While this mobility system provides considerable personal freedom for those who can afford it and enables substantial economic activity, it is associated with serious side effects in terms of safety, energy, the environment, land use, traffic congestion, time use and equality of access.

In this context, a wide range of technology and business enablers are emerging that, when combined in innovative ways, promise to transform the way people and goods move around and interact economically and socially. Roadway transportation is now as *ripe for transformation* as the telecommunications, photography, computer, media, television and pharmaceutical industries were over the past two decades.

THE OPPORTUNITY

It is now possible to supply ***better mobility experiences at radically lower cost*** to consumers and society. This opportunity results from innovatively combining five emerging technology and business enablers:

- *The "Mobility Internet"* does for the movement of people and goods what the Internet has done for the movement of information by coordinating large amounts of real-time spatial and temporal connectivity and infrastructure data.
- *Self-driving/driverless vehicles* operate without human control enabling passengers to use their time as they please (e.g., texting, talking on the phone, eating, or watching a movie) without endangering themselves or others.
- *Shared vehicles* are used by several people throughout the day rather than being used exclusively by single individuals and being parked 90 percent of the time.
- *Specific-purpose vehicle designs* are tailored to the type of mobility they supply and the number of occupants they serve, making them more energy-, space-, and cost-efficient compared to general-purpose vehicles.
- *Advanced propulsion systems* move cars and trucks using alternative energy sources and power systems in addition to oil and combustion engines; they typically entail electric drive, electric motors and electronic and digital controls.

Individually, each of these building blocks promises incremental improvements over today's roadway transportation services. When combined in innovative ways to enhance the mobility experiences of consumers, the improvements are radical and the changes are transformational.

Transformational Opportunity

The “Mobility Internet”
+
Self-Driving/Driverless Vehicles
+
Shared Vehicles
+
Specific-Purpose Designs
+
Advanced Propulsion

***Better Mobility Experience at
Radically Lower Cost***

To illustrate the opportunity, consider the following three scenarios:

Scenario #1:

Joe, a typical car owner living in a U.S. city such as Ann Arbor, Michigan, spends about \$16 per day to own his mid-size sedan, and an additional \$5 per day to operate it, for a total of \$21 per day. These costs include depreciation, insurance, gasoline, maintenance, license fees, finance charges and taxes. He also spends money on parking and must use 60 to 90 minutes of his time each day focused on driving his car to and from work, and to access shopping, school, recreational and community activities. He has a busy schedule and his time is valuable.

Scenario #2:

Bob and his wife live in Babcock Ranch, a new eco-city in Southwest Florida. Before moving to Babcock Ranch, they were a two-car household, spending about \$16 per day to own each vehicle, with additional operating costs depending on how many miles they drove each day. Now that they have moved to Babcock Ranch, many of their daily activities (work, grocery store, health club, golf course) will be located within the town. They are wondering if it makes sense to continue to own two cars.

Scenario #3:

Anne lives in Manhattan with her husband and two children. Anne's family sold their car when they moved to Manhattan due to the cost and hassle of parking. Instead they use the wide range of public and shared transportation options available. They use the subway or buses daily on their way to school and work, and use yellow taxis when they need to get somewhere faster or are unable to use public transportation. All in all, each adult spends an average of \$200 per month for transportation via bus, subway, and yellow taxi. Although Manhattan has an excellent public transportation system, there are limitations. During rush hour, buses and trains are uncomfortably crowded. During non-rush hours, wait times for buses can be long. Traveling north south is easy, but going cross-town is either time-consuming (2 to 3 buses and subways) or expensive (yellow taxi). Anne and her family currently organize their schedule around the constraints of the public transportation system.

NEW MOBILITY SYSTEM

Now consider a new mobility system that can improve the experiences of Joe, Bob, and Anne. This new mobility system combines recent developments in driverless vehicle technologies with the emerging "Mobility Internet" that can coordinate the movements of these vehicles through space and time.

The new mobility system works as follows:

- Optimally sized shared fleets of driverless, coordinated, specific-purpose vehicles are available.
- Customers request a ride using an app on their smartphone.
- An autonomous vehicle arrives at their door within minutes and transports them directly to their destination.
- During the trip, the customer can use his/her time as desired (reading, eating, talking on the phone, watching a movie, sending email).
- Upon arrival, there is no need to park the vehicle, because it continues on to a nearby location to pick up another rider.

In Ann Arbor, Joe is able to spontaneously request a ride using an "app" on his smart phone. An autonomous vehicle arrives at his door within minutes and transports him directly to his destination. During the trip, Joe uses his time as he pleases. Upon arrival, he doesn't need to park the vehicle, which continues on to a nearby location to pick up another rider.

At Babcock Ranch, Bob and his wife have spent the day golfing. They used the new mobility system to get to and from the golf course, without having to use their personally owned cars. At the weekend, when they want to visit friends who live outside Babcock Ranch, they can still use their personally owned car.

In Manhattan, Anne needs to get across town to a doctor's appointment or a business meeting. Instead of trying to figure out whether she has time to take the bus or is willing to spend money on a yellow taxi, she requests a ride using an "app" on her smart phone. An autonomous vehicle arrives at her door and transports her directly to her destination.

HOW MUCH DOES THIS SERVICE COST TO SUPPLY?

Remarkably, initial estimates indicate that the cost to supply this service to Ann Arbor customers like Joe could be as little as \$2 per day. This cost reduction is a result of:

- *Better capital utilization:* Far fewer shared, driverless vehicles are needed to provide the same level of service as personally owned vehicles.
- *Better capacity utilization:* During peak travel times, the shared vehicles are occupied more than 75 percent of the time, compared to Joe's car which is in use less than 5 percent of the time; and
- *More efficient energy use:* When Joe travels alone or with one other person, the 1-to-2 passenger, purpose-designed vehicle that they ride in weighs 75 percent less than a conventional car, thereby using significantly less energy.

In addition to radically lower cost, Joe's mobility experience is significantly better in terms of safety, convenience, time use and peace of mind. In fact, this new mobility experience is so good and meets his needs so well that Joe no longer owns a car.

Similarly, at Babcock Ranch, initial estimates indicate that the new mobility system can be provided for \$3 per day per customer or about \$1 per trip. Bob and his wife enjoy the convenience of this mobility service so much that they have sold their second car.

Finally, in Manhattan, the new mobility system could operate as an alternative mode of transportation, competing with both yellow taxicabs and public transportation. Yellow taxicab fares are about \$5 per mile. Initial estimates indicate that a fleet of shared, driverless vehicles would cost about \$0.40 per mile to operate. This presents an appealing option in the current portfolio of mobility services. Compared to the bus or subway, shared, driverless vehicles would give Anne and her family superior comfort, convenience, and route flexibility. Compared to yellow taxicabs, they would be more convenient and less expensive as a result of purpose-built vehicle design, energy efficiency, and reduced labor costs.

SYSTEM MODELING

Analytical and simulation models were used to arrive at these conclusions. These models estimate the cost and performance of shared fleets of driverless vehicles serving trips within a region, and compare these estimates to current cost and performance.

Consider a region such as those in our scenarios where most of the trips for work, shopping, recreation and other purposes are intra-regional. We assume that a shared, driverless vehicle fleet operates within this region in the manner described above. The fleet operator optimizes the assignment of vehicles to customers so that low customer wait times are achieved and operational costs are minimized. (Note that while our discussion focuses on a single vehicle fleet, this doesn't preclude the possibility that there may be several competing vehicle fleets in an area). Since customers can spontaneously request a vehicle as described above, the operator must control the assignment of vehicles to customer trip requests in real-time. When Joe requests a shared vehicle using his smart phone app, the fleet operator will have information on Joe's current location, destination, and vehicle type needed. We assume Joe needs the vehicle immediately, since this would be the hardest type of travel to accommodate

(If Joe wants to schedule trips ahead of time, the shared fleet could certainly handle this, but this would only reduce the costs we're calculating here).

As shown in Figure 1, upon receiving this trip request, the fleet operator immediately determines which vehicle in their fleet could reach Joe first. The operator knows (via the Mobility Internet) the current location and state of all vehicles, can estimate when those vehicles currently serving trips will complete them, and can estimate when every vehicle could potentially reach the new customer's location. The closest vehicle could be one that is currently idle, or as shown in the figure, it could be a vehicle currently serving another trip. In either case the assigned vehicle travels to the new customer as shown as soon as possible. Through the Mobility Internet, the operator keeps Joe fully informed about the assigned vehicle's current location and expected arrival time. Note that assigning the closest vehicle to a customer not only minimizes wait time, but also approximately minimizes empty vehicle miles, an important part of operating costs. Also, if trips are scheduled ahead of time or if the fleet operator can forecast travel patterns, then customer wait time could be reduced by pre-positioning the vehicles where the operator thinks (or knows) they will be needed.

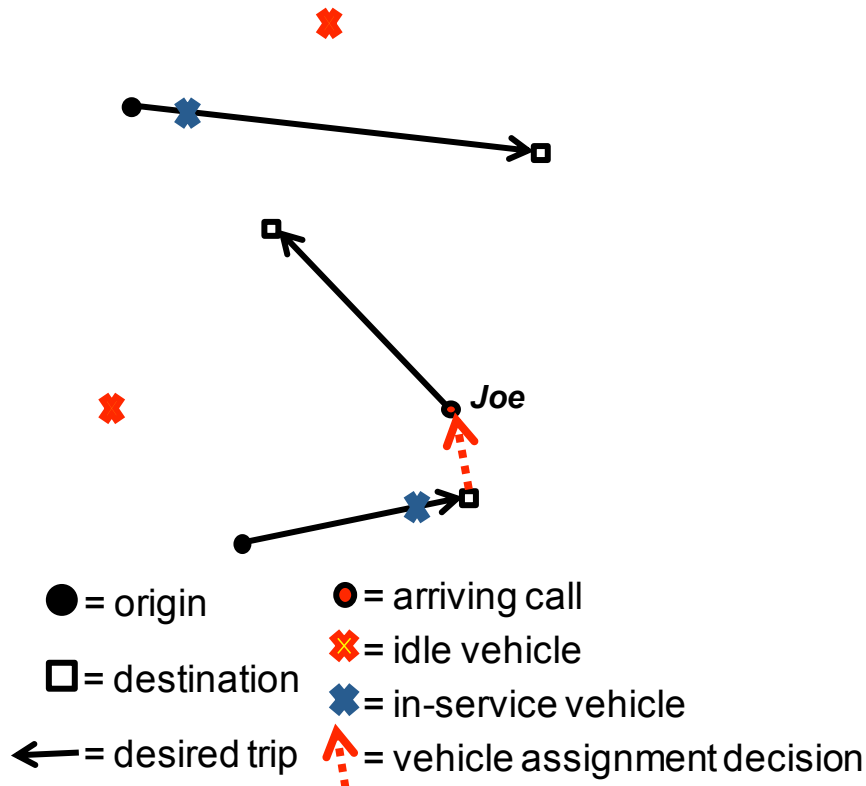


Figure 1. Shared, Driverless Fleet Assignment Decisions

We want to estimate the cost of a shared, driverless vehicle system. The primary costs are the cost of owning vehicles and the cost of operating them. Ownership costs include vehicle purchase, registration, and insurance costs; operating costs including fuel, maintenance, and repair costs. These operating costs depend not only on the length of trips served by the fleet, but also on the empty distance vehicles travel in getting to customer locations.

We also must estimate the customer performance of the system. In particular, we need to calculate how the shared vehicle fleet size (i.e., # of vehicles) and assignment decisions impact customer wait time from when they request a vehicle until one arrives at their location. To provide customers with a better mobility experience than they currently have, customer wait

times for a vehicle need to be short. The definition of short depends on the situation. Joe's and Bob's alternatives are usually their cars, which are immediately accessible when they are home, but may have substantial walk time to access when parked at a shopping center or other location. Anne in Manhattan usually has substantial walk and/or wait time when accessing public transportation or cabs. We'll estimate shared vehicle fleet costs assuming average wait times from calling for a vehicle until it's arrival should be 2 minutes or shorter. This is a fairly stringent standard that will lead to conservative cost estimates.

To estimate system performance and cost, we develop relatively simple models that capture the key variables that impact them. These are: the area of the region, the mean trip length, the mean trip rate and how this varies throughout the day, mean vehicle speed, the average fixed time needed per trip, the fleet size, and vehicle cost parameters. Using queueing and network modeling approaches we develop an analytical model to relate these key variables to shared, driverless vehicle fleet performance and cost. The Appendix develops this model, shows the results of validating it with simulation models, and presents results of applying the model to the three case studies and to more general situations.

Results for the three case studies are described in the following sections.

ANN ARBOR CASE STUDY: COST COMPARISON TO PERSONAL VEHICLE OWNERSHIP

To compare the potential benefits of a shared, driverless vehicle fleet with personally-owned vehicles, a case study was done for Ann Arbor, Michigan. Travel patterns and the cost of personal vehicle ownership were analyzed to determine whether such a system could provide residents with a less expensive and more convenient way of getting around.

Ann Arbor has a population of 285,000 and covers an area of 130 square miles. Ann Arbor was selected for the case study because it is representative of other small to medium-sized cities in the United States, based on data from the 2009 National Household Travel Survey. The other cities used in the comparison were: Austin, Texas; Orlando, Florida; Rochester, New York; Sacramento, California; and Salt Lake City, Utah. Although the population of Ann Arbor is smaller than that of the other cities (which range from 650,000 to 1.5 million), the travel data are very similar, including: average trip time, average trip length, average trips per vehicle, average vehicles per person, average vehicle occupancy, and average vehicle usage. When the data were compared, vehicle usage patterns were found to be remarkably consistent between all six cities. The absolute size of a shared, driverless vehicle fleet would vary depending on the population and area of each city. However, because the travel data for the cities is so similar, the costs per consumer per day would be nearly the same as for Ann Arbor. The estimated cost savings per consumer per day would therefore hold as well.

Ann Arbor, Michigan



- 285,000 people
- 130 square miles
- 200,000 personally owned vehicles
- 740,000 trips per day
 - 528,000 internal trips (<70 miles)
- Average trip
 - 8.3 miles
 - 16.8 min
 - 30 mph
 - 1.6 people
- Vehicles used an average of 67 minutes/day (5%)

Source: 2009 National Household Travel Survey data for Detroit-Ann Arbor-Flint area

Methodology

To evaluate whether there is consumer value to be gained by operating a shared, driverless vehicle fleet in Ann Arbor, the following methodology was used:

- Obtained travel data for Ann Arbor, including number of trips per day, average trip time, average trip distance, average trip speed, and average number of passengers.
- Used queuing, network, and simulation models to estimate how a system of shared, driverless vehicles would perform in meeting the demand for daily trips within Ann Arbor.
- Determined the number of shared, driverless vehicles needed to ensure adequate coverage and acceptable wait times during peak periods.
- Once the fleet size was known, available cost estimates for owning and operating mid-sized vehicles were used to estimate the cost of providing mobility services.
- Finally, the estimated cost of providing mobility services was compared to the cost of personal car ownership.

Trip Data for Ann Arbor

In 2009, Ann Arbor residents owned a total of 200,000 passenger vehicles. According to the National Household Travel Survey² these vehicles were driven for 740,000 trips per day, for an average of 3.7 trips per vehicle per day. Each trip averaged 8.3 miles and took 16.8 minutes at an average speed of 30 mph. The average number of occupants per trip was 1.6 and vehicles were in use a total of 67 minutes per day on average (about 5 percent of the time).


A shared driverless fleet would most likely compete for trips taken within the Ann Arbor urban area. The analysis therefore focused on the 120,000 vehicles that were driven less than 70 miles per day. These vehicles were responsible for 528,000 trips per day, for an average of 4.4 trips per vehicle per day. These internal trips averaged 5.8 miles in distance with 1.4 occupants.

Determining Shared, Driverless Fleet Size

To determine the size of the shared, driverless fleet that would be needed to serve the internal trips taken by the Ann Arbor community, queuing, network, and simulation models were used to estimate system performance. Results indicate that the same number of internal trips could be provided by a drastically reduced fleet size. The size of the shared, driverless fleet would vary depending on the acceptable wait time for consumers. For example, with a fleet size of 18,000 vehicles, consumers would expect to wait less than one minute for a vehicle to arrive, and the vehicle fleet would be utilized 75 percent of the time on average.

Ann Arbor Case Study:

Customer Wait Time is Short, Empty Miles are Low




- One shared fleet serves all 120,000 customers who currently use their car < 70 miles/day
- Provides almost instantaneous access to a vehicle while reducing the number of vehicles by 80+%
- Shared fleets decrease ownership costs while increasing operating costs to reposition empty vehicles
 - In widespread use, however, the empty mile cost would be small
- Pick a fleet size of 18,000 (~70% utilization) for cost calculation

Results also indicate that economies of scale are reached quickly. The analysis looked at what the vehicle fleet size and wait time would be if the geographical area remained the same, while

² U.S. Department of Transportation. 2009 National Household Travel Survey. Data for Detroit-Ann Arbor-Flint area.

the number of participating consumers decreased. Economies of scale are reached at about 1,000 consumers. The benefits of shared, driverless fleets can therefore also be achieved for consumer groups with niche interests.

Ann Arbor Case Study: Scale Economies Are Reached Quickly



- Fleet sizes set so that customer wait time < 2 min.
- Scale economies in operating a shared fleet are reached quickly (at about 1,000 customers)
- Most of empty mile reductions achieved with small fleets
- Benefits of shared fleets can be achieved for smaller pools of customers needing specialized vehicles tailored to their needs

Potential Cost Savings from Shared, Driverless Fleet

Cost Savings from Reduced Vehicle Ownership Costs

AAA³ has estimated that in 2012, the ownership costs for a medium-sized sedan were \$16 per day.⁴ If this sedan were retrofitted with driverless technology, ownership costs could increase by about 20 percent, for a total of \$19 per day. A shared, driverless fleet, however, would serve the needs of 120,000 consumers with 18,000 vehicles instead of 120,000. Even when the cost of the driverless technology is factored in, vehicle ownership costs are reduced to about 15 percent of the cost of the personally owned fleet, or under \$3 per day per consumer.

Cost Savings from Reduced Operating Expenses

Additional cost savings can be achieved in both energy consumption and vehicle ownership if, instead of a fleet of general-purpose vehicles, the fleet consists of multiple vehicle designs that

³ AAA, 2012 "Your Driving Costs: How Much Are You Really Paying to Drive?"

⁴ Ownership costs including full-coverage insurance, license, registration, taxes, depreciation (15,000 miles annually) and finance charges were estimated to be \$5,771 per year or \$15.81 per day.

are purpose-built for the types of trips taken. For the vehicles driven within the urban area of Ann Arbor, approximately 90 percent of trips taken involve 1 to 2 occupants, with the remaining 10 percent involving 3 to 6 occupants. If a shared, driverless fleet included vehicles purpose-built to carry 1 to 2 occupants, energy consumption and vehicle cost could be significantly reduced.

To estimate these cost savings, assume that general-purpose vehicles (such as the mid-sized sedan described above) weigh approximately 3,000 to 4,000 pounds and that vehicles purpose-built for 1 to 2 occupants weigh 800 to 1,000 pounds. Vehicles that are one-fourth the mass can be estimated to cost approximately one-fourth to purchase. The reduced mass combined with improved energy efficiency from electric drive results in one-tenth the energy consumed.⁵ AAA has estimated that in 2012, the operating costs of a mid-size sedan are \$0.20 per mile.⁶ The Ann Arbor data indicate that vehicles traveling within the urban area average 26 miles per day. The operating cost per day for these general-purpose vehicles is therefore approximately \$5 per day. By using purpose-built vehicles, we can reduce the operating cost to under \$1 per day and the ownership cost to under \$1 per day.

Cost Savings from Parking Fees

The cost of parking a personally owned vehicle in Ann Arbor averages \$5 per day (costs vary depending on location). Parking costs would be saved with the use of a shared, driverless fleet. Such a fleet would be in use approximately 75 percent of the time, as opposed to personally owned vehicles that are in use 5 percent of the time on average and are parked the remaining 95 percent of the time.

Cost Savings from Value of Time

In addition to the reduced costs of ownership, operation and parking, driverless vehicles will free up the travel time that is currently focused on driving and parking the vehicle. Consumers will be able to spend that time more productively on work or leisure activities. If consumers spend one hour less each weekday on driving and parking, and if the value of consumer time is estimated at \$25 to \$50 per hour, then the annual value of time saved ranges from \$6,250 to \$12,500 per year per consumer.⁷


⁵ Multiply one-fourth the vehicle mass with an electric drive that uses two-fifths of the energy of current engines to achieve energy consumption of one-tenth of the general-purpose vehicle.

⁶ Operating costs including gas, maintenance, and tires were estimated to be \$0.2006 per mile.

⁷ The median annual income in the USA is \$50,000, which translates into approximately \$25 per hour. The typical new car buyer has an annual income of \$80,000 to \$100,000 per year, which translates into \$40 to \$50 per hour. If a consumer saves 1 hour per day for 250 days out of the year and their time is worth \$25 per hour, then the annual value of time saved will be \$6,250. If their time is worth \$50 per hour, then the annual value of time saved will be \$12,500.

Ann Arbor Case Study

Costs per customer per day would be dramatically lower with shared, driverless fleet, from \$21 to \$2



- A shared driverless vehicle fleet can provide the same mobility as personally owned vehicles at far less cost
- Based on cost of today's vehicles
- Based on trips in vehicles driven within region
- Reduced parking expenses and value of time gained by not having to drive are added benefits

Ann Arbor Conclusions

A shared, driverless vehicle fleet would provide *better mobility experiences at radically lower cost* when compared to personally owned vehicles. Consumers would benefit from increased spontaneity, door-to-door convenience, more productive and relaxing travel time, more predictable and shorter travel times, and improved safety. The cost of operating and owning a mid-sized car is estimated to be \$21 per day for 2012. Analysis indicates that these costs can be reduced to under \$2 per day, representing a significant business opportunity. If parking expenses and the value of time spent driving are included, the total cost of today's mobility system comes to about \$50 per day, making the savings from a shared, driverless fleet even more dramatic.

BABCOCK RANCH CASE STUDY: COST COMPARISON TO PERSONAL VEHICLE OWNERSHIP

The Ann Arbor case study indicates that a shared, driverless fleet could provide better mobility experiences at radically lower cost when compared to personally owned vehicles and that these results can be extended to other medium-sized U.S. cities. What about smaller urban or suburban areas? Could the cost savings and increased convenience of the new mobility service be provided for these communities as well? To find out, a case study was done for the planned eco-city of Babcock Ranch in southwest Florida.

The city of Babcock Ranch is being built on a former cattle ranch. In 2006, developers Kitson & Partners purchased the 91,000-acre ranch, selling 73,000 acres back to the state of Florida to create a nature preserve, and retaining 17,000 acres on which to build an eco-city. When completed (around the year 2030), Babcock Ranch will have a population of 50,000 residents, as well as retail, commercial, and office space.

Babcock Ranch, Florida



- 50,000 people at buildout
 - 20,000 by 2022
- 20 square miles total area, of which 10 square miles developed
- 115,000 internal trips per day
 - 2.3 trips per person
- Average trip
 - 3.5 miles
 - 25 mph
- 9,800 internal trips during peak hour

Source: Analysis done for Babcock Ranch by David Plummer & Associates, 2011

Babcock Ranch is envisioned as a “living laboratory,” combining a philosophy of sustainability with cutting-edge technology to create a highly attractive living environment. Over half of the 17,000 acres will be permanently protected as greenways and open space, and residents will have access to the adjacent 73,000-acre Babcock Ranch preserve. All commercial buildings and homes will be certified as energy-efficient and constructed according to Florida Green Building Council standards. Babcock Ranch plans to be a platform for sustainable clean energy technology, as well as a proving ground for emerging “smart city” products and services. To that end, Florida Power & Light has committed to building the nation’s largest solar power facility (75 MW) at Babcock Ranch. In addition, an integrated “smart grid” will provide greater efficiencies and allow residents and businesses to monitor and control their energy consumption.

Methodology

To evaluate whether there is consumer value to be gained by operating a shared, driverless vehicle fleet in Babcock Ranch, the following methodology was used:

- Obtained travel data for Babcock Ranch, including number of trips per day, average trip time, average trip distance, average trip speed, and average number of passengers.

- Used queuing, network, and simulation models to estimate how a system of shared, driverless vehicles would perform in meeting the demand for daily trips within Babcock Ranch.
- Determined the number of shared, driverless vehicles needed to ensure adequate coverage and acceptable wait times during peak periods.
- Once the fleet size was known, available cost estimates of owning and operating mid-sized vehicles were used to estimate the cost of providing mobility services.
- Finally, the estimated cost of providing mobility services was compared to the cost of personal car ownership.

Trip Data for Babcock Ranch

When fully built out in 2030, the population of Babcock Ranch will be 50,000. The city will cover a total area of 20 square miles (17,000 acres), with 10 square miles retained in the form of parks and green spaces. The area within which vehicles will travel is therefore the remaining 10 square miles of developed property.

The analysis assumed that Babcock Ranch residents would use the shared, driverless fleet for trips taken within Babcock Ranch, and that they might continue to own one car for trips beginning and ending outside Babcock Ranch. Projected travel data was obtained from an analysis done by David Plummer & Associates for Babcock Ranch in 2011.⁸ This study estimated a total of 115,000 internal trips per day at build out, or an average of 2.3 trips per day per person. The average trip length was estimated to be 3.5 miles with an average vehicle speed of 25 mph.

Determining Shared, Driverless Fleet Size

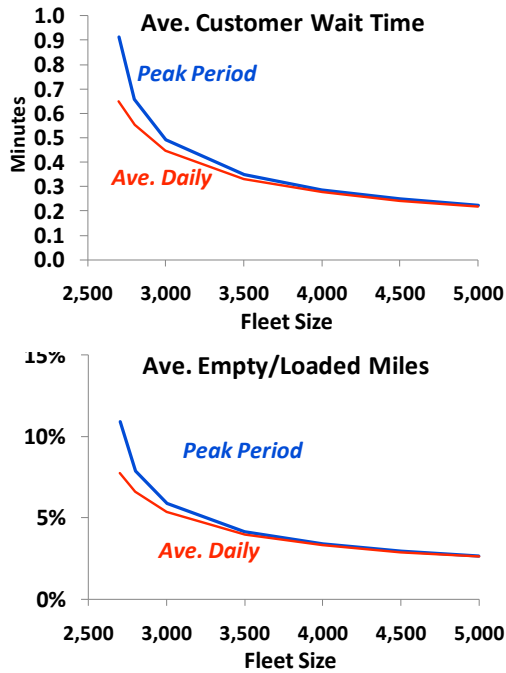
By using the Babcock Ranch trip data combined with modeling of a shared, driverless fleet, the number of vehicles needed to cover the expected demand can be estimated. Model assumptions include random origins and destinations within Babcock Ranch, and a circuitry factor of 1.4 applied to the straight-line distance when computing trip lengths. Peak hour trips were considered when determining the fleet size, as well as customer wait time and total vehicle utilization.

Results indicate that a fleet of 3,000 to 4,000 shared, driverless vehicles would be needed to serve a fully developed Babcock Ranch during the peak daily period. With a fleet this size, the average wait time for travelers would be well under one minute. A customer would always be close to an empty vehicle. For example, if there were 3,500 vehicles in the fleet, an average of 1,000 vehicles (100 per square mile) would be available at any time. Empty travel distances would also be very low. Estimates indicate that a shared fleet of 3,500 vehicles would be utilized 70 percent during peak periods, and would average only slightly less throughout the day.

⁸ David Plummer & Associates, Inc. "Babcock Ranch Community Development of Regional Impact Question 21. Transportation Master Traffic Study Update" (July 21, 2011).

Babcock Ranch Cast Study:


Customer Wait Time is Short, Empty Miles are Low



- A shared vehicle fleet of 3,000-4,000 vehicles would be needed to serve a fully developed Babcock Ranch in the peak period
- Average times travelers wait for vehicles would be well under a minute

As with Ann Arbor, scale economies are reached quickly in terms of the number of vehicles needed and the cost of owning and operating them. Costs are relatively constant once the number of customers rises above 5,000, and increase very little as the customer base changes from 5,000 to 1,000. In an area the size of Babcock Ranch, a relatively small population could therefore be efficiently served by a shared fleet.

Babcock Ranch Case Study: Scale Economies Are Reached Quickly



- Costs are relatively constant once the number of customers rises above 5,000.
- Cost rises a little as the customer base drops from 5,000 to 1,000
- In an area the size of Babcock Ranch, a relatively small population could be efficiently served by a shared fleet.

Cost Estimate for Shared, Driverless Mobility Service

For the Ann Arbor analysis, the number of vehicles used for internal trips was known, making it possible to compare the cost of internal trips using personally owned vehicles with the cost of providing the same service using a shared, driverless fleet. For Babcock Ranch, the number of personally owned vehicles that would be used to travel internally is not known, making it impossible to estimate cost savings. However, it is possible to estimate the cost of providing the mobility service.


As noted above, AAA has estimated 2012 ownership costs for a medium-sized sedan to be \$16 per day. For Babcock Ranch, we make the conservative assumption that the smaller, specific-purpose vehicles with driverless technology also have ownership costs of \$16 per day.⁹ If the 50,000 residents of Babcock Ranch were served by a shared, driverless fleet of 3,500 vehicles, the ownership costs for this fleet would be about \$1 per day per resident.

Similarly, we can estimate the operating costs using the 2012 AAA estimate of \$0.20 per mile. Residents are expected to travel internally an average of 8 miles per day, resulting in operating costs of about \$1.50 per day per resident. The mobility service could therefore be provided to consumers for a total cost of under \$3 per day per person, or just over \$1 per trip on average.

⁹ In reality, ownership costs will most likely be lower because the 20 percent increase in costs due to the driverless technology will be more than offset by the decrease in costs of the smaller vehicle.

Babcock Ranch Case Study:

Cost of mobility service could be as low as \$3 per day per customer and as low as \$1 per trip



- Assume vehicle cost = \$16/day
 - 2012 AAA estimate for mid-size sedan depreciated over 5 years
- Operating cost = \$0.20/mile
 - 2012 AAA estimate for fuel and maintenance for a mid-size sedan
- For a fleet of 3,000 – 4,000 vehicles, total cost is under \$3/person served/day
- At 2.3 trips/person, this is just over \$1/trip
- Average trip length is 3.5 miles, so total costs are about \$0.30/person-mile

Babcock Ranch Conclusions

As with Ann Arbor, the Babcock Ranch analysis concludes that a shared, driverless fleet could provide improved mobility experiences at a very low cost. A shared, driverless vehicle fleet would provide increased spontaneity, door-to-door convenience, more productive and relaxing travel time, more predictable and shorter travel times, and improved safety. The shared, driverless fleet could complement the use of personally owned vehicles, with Babcock Ranch residents using the new mobility service when convenient. The cost of providing these services would be about \$3 per day per person, or \$1 per trip on average.

MANHATTAN CASE STUDY: COST COMPARISON TO YELLOW TAXICABS

The island of Manhattan is home to a population of 1.6 million people living in an area of 23 square miles. Manhattan is one of the five boroughs in New York City, along with Brooklyn, Queens, the Bronx, and Staten Island. Because parking a car is difficult and expensive, only about 25 percent of Manhattan residents own a car. Manhattan has one of the world's most well developed systems of public and private transportation. The Metropolitan Transit Authority (MTA) runs bus and subway services throughout Manhattan and the other boroughs. In addition, there are several types of for-hire car services available, including yellow taxicabs, dispatch cars, "black car" services, and limousine services. Finally, residents have access to hourly rental cars (ZipCar, Hertz on Demand) and traditional rental cars.

Manhattan, New York



- 1.6 million people
- 23 square miles
- 410,000 trips per day
 - 0.3 trips per adult
 - 88% of all taxicab trips internal to Manhattan
- Average Trip
 - 2 miles
 - 11 minutes
 - 11 mph
 - 1.4 passengers
- Taxicab Fleet
 - 13,000+ vehicles

Sources: 2009 National Household Travel Survey data; 2010 US Census; NYC taxi data; Schaller Consulting, *The NYC Taxicab Fact Book* (2006)

A shared, driverless vehicle fleet could compete strongly with the mobility services provided by yellow taxicabs and, to varying degrees, by buses, the subway, other for-hire car services, hourly rental cars and personally owned cars. For simplicity, this analysis focuses on comparing the cost of using a shared, driverless vehicle fleet to the cost of using yellow taxicabs within Manhattan. Because yellow taxicabs are now required to have GPS on board, excellent information on taxi trips is available.

Manhattan's Yellow Taxicabs

There are multiple types of car services for hire in Manhattan. The best-known are the taxicabs, characterized by their yellow exteriors, their metered fares, and the fact that they can only be hailed by passengers from the street (i.e., you can't call a yellow taxicab to come and pick you up). They are the only type of car service legally permitted to pick up passengers without prearrangement¹⁰. In order to operate a yellow taxicab, an owner must purchase a license from the city. Only a limited number of these licenses are available (13,237 in 2011).¹¹ They are often called "medallions" because of the metal medallion attached to the hood of a taxicab, indicating that the car is licensed.

In addition to the yellow taxicabs, dispatch, "black car", and limousine services exist. For dispatch car service, a customer can call a central dispatch office and request that a car pick them up and take them to a specified location. Dispatch offices are in contact with available cars

¹⁰ The Taxi & Limousine Commission (TLC) voted in April 2012 to allow livery cabs to be licensed to make street pick-ups in the outer boroughs and northern Manhattan where yellow taxicab service is low, a rule scheduled to take effect in the summer of 2012.

¹¹ Wikipedia, "Taxicabs of New York City". Accessed June 2, 2012.

via radio. These car services range widely in terms of quality of service, price, and types of vehicles used. Customers using dispatch car services include passengers traveling to and from the airports, and residents of neighborhoods where yellow taxicabs are hard to find. “Black car” services are generally fleets of higher quality vehicles that contract with companies to transport employees occasionally or daily to and from work and the airport. Limousine services are often used for special events. All in all, there are over 40,000 for-hire vehicles operating in the five boroughs of New York City, in addition to the 13,000+ yellow taxicabs.¹²

Methodology

To evaluate whether there is consumer value to be gained by operating a shared, driverless vehicle fleet in Manhattan, we used the following methodology:

- Obtained taxicab trip data for Manhattan, including number of trips per day, average trip time, average trip distance, average trip speed, and average number of passengers.
- Used queuing, network, and simulation models to estimate how a system of shared, driverless vehicles would perform in meeting Manhattan’s daily taxicab trip demand.
- Determined the number of shared, driverless vehicles needed to ensure adequate coverage and acceptable wait times during peak periods.
- Once the fleet size was known, available cost estimates of owning and operating mid-sized vehicles are used to estimate the per mile cost of providing mobility services.
- Finally, the estimated cost of providing mobility services were compared to existing fares charged by taxicabs.

Trip Data for Manhattan

Manhattan residents use taxicabs to get to a number of different destinations each day, including work, school, home, shopping, entertainment, and restaurants. All in all, taxicabs account for 470,000 trips per day in the five boroughs,¹³ with 88 percent of these trips internal to Manhattan (approximately 410,000 trips per day). Analysis of data obtained from FareShareNYC and the Taxi and Limousine Commission (TLC) was used to determine that during the weekday peak, over 300 taxicab trips are initiated per minute, with the average trip taking about 11 minutes to cover a distance of 2 miles. The average speed driven is between 10 and 11 mph depending on time of day, and the average number of passengers per trip is 1.4.¹⁴

Determining Shared, Driverless Fleet Size

By using the taxicab trip data combined with modeling of a shared, driverless fleet, the number of shared, driverless vehicles needed to cover the demand currently met by yellow taxicabs can be estimated. Assumptions include random origins and destinations within Manhattan, and a circuitry factor of 1.4 applied to the straight-line distance when computing trip lengths. Peak hour trips were considered when determining the fleet size, as well as customer wait time and total vehicle utilization.

Results indicate that a fleet of 9,000 shared, driverless vehicles could meet the demand with much shorter wait times and significantly improved vehicle utilization. Currently, the average


¹² Schaller Consulting, *The NYC Taxicab Fact Book* (2006).

¹³ Ibid.

¹⁴ Data originally accessed at <http://faresharenyc.com/data-analysis/>, with additional information provided by Jeff Novich, Charles Komanoff, and Aaron Glazer of FareShareNYC. Additional data subsequently obtained from the Taxi & Limousine Commission (TLC).

wait time for a passenger hailing a yellow taxicab is 5 minutes. A fleet of 9,000 shared, driverless vehicles that were centrally coordinated to respond to passengers “hailing” them via their smartphones, could reduce the wait time to under 1 minute, while increasing the vehicle utilization.

Manhattan Case Study: Customer Wait Time is Short, Empty Miles are Low



- High customer and vehicle density in Manhattan means a coordinated system could serve customers with a very short wait time and low empty miles
- On average a fleet of 9,000 vehicles could yield (during the peak period)
 - wait times of about 0.6 minutes
 - 0.1 empty miles/trip in a balanced system (5% of 1.9 miles/trip)
 - Imbalances between trip origins and destinations increase empty miles as discussed on slide 14

Potential Cost Savings from Shared, Driverless Fleet

Passengers pay on average \$5 per trip-mile¹⁵ when riding in a yellow taxicab. This revenue covers the labor cost of the driver, vehicle ownership and operating costs, and owner income. If we assume owner profit to be 15 percent, the cost of providing the yellow taxicab service is approximately \$4 per trip-mile.¹⁶ In contrast, the cost of providing a similar service via a shared, driverless vehicle fleet is estimated to be \$0.40 per mile.¹⁷ These cost savings are the result of fewer taxis, less empty miles and reduced labor costs of the driver.

¹⁵ The average fare is \$8 plus a 15 percent tip (\$1.20) for a total of \$9.20 for a 1.9-mile trip. The average cost per mile of the trip is therefore about \$5.

¹⁶ \$9.20 minus 15 percent is \$7.80 per trip, divided by 1.9 miles per trip for a per trip-mile cost of \$4.11.

¹⁷ AAA estimates the 2012 cost of owning a mid-sized sedan to be \$16 per day. We make the conservative assumption that a smaller, specific-purpose vehicle with driverless technology will also cost \$16 per day to own. Daily ownership costs for a fleet of 9,000 vehicles would be \$144,000. If the total number of miles driven internally in Manhattan each day averages 779,000 (410,000 internal trips multiplied by 1.9 miles per trip), then ownership costs per mile equal \$0.19. AAA estimates 2012 operating costs to be \$0.20 per mile, for a total of \$0.39 per mile combined.

Cost Savings from Central Coordination


Significant cost savings are possible through improved coordination. Currently, yellow taxicabs aren't allowed to provide scheduled or pre-arranged service (dispatch and "black car" services do this), so all trips are "hailed" from the street. This lack of central coordination results in a high level of empty miles for taxicabs, in addition to a higher wait time for customers (5 minutes on average). A fleet of 9,000 shared, driverless vehicles could reduce the wait time to under a minute and increase vehicle utilization to 70 percent during peak periods.

Cost Savings from Reduced Ownership and Operating Expenses

A fleet of 9,000 shared, driverless vehicles could meet the demand currently supplied by 13,000+ yellow taxicabs. This reduced fleet size results in savings from vehicle ownership costs. In addition, a more efficiently operated fleet would result in fewer empty miles, thereby reducing the total operating costs per loaded mile.

Manhattan Case Study:

Shared, driverless fleet would have \$0.80 per trip cost compared to \$6.30 - \$7.80 per trip cost of taxicabs



- Coordinating the vehicle fleet with customer demand results in lower system cost by reducing both empty miles and fleet size while providing much better customer service
- A fleet of 9,000 vehicles would
 - be over 70% utilized in the peak period
 - have a total cost of under \$0.80/trip using AAA cost parameters (this compares to between \$6.30 and \$7.80 for the current taxi system)

Manhattan Conclusions

Manhattan's yellow taxicabs are rated poorly for quality of service, with passengers giving them low ratings, compared to other modes of transportation, for ability to get a cab when you want one, value for money, safety from accidents, driver understanding of directions, and driver

courtesy.¹⁸ A shared, driverless fleet could provide better mobility experiences at radically lower cost. Significantly reduced wait times can be achieved using central coordination from a dispatch system accessed via the Mobility Internet. The current per-mile fare for a yellow taxicab is \$5, compared with an estimated \$0.40 per-mile cost of operating a shared, driverless vehicle fleet.

Although the current analysis has focused on yellow taxicabs, these represent only 8 percent of trips taken by Manhattan residents each day. A convenient and affordable mobility service would most likely draw passengers from other modes of transportation as well, including buses, the subway, other for-hire car services, and hourly rental cars.

CONCLUSIONS

By combining five business model and technology enablers (the “Mobility Internet,” self-driving/driverless vehicles, shared vehicle systems, specific-purpose vehicle designs and advanced propulsion systems) ***better mobility experiences at radically lower cost*** can be realized in a variety of settings. This mobility system can successfully compete with other modes of public and private transportation in Ann Arbor, Babcock Ranch, and Manhattan, by providing mobility services that are more convenient, more energy and resource efficient, safer, and cheaper. More specifically:

- ***Shared, driverless vehicle fleets result in greater efficiencies.*** In Ann Arbor and Babcock Ranch, far fewer cars are needed to cover the same number of trips as personally owned vehicles. In Manhattan, central coordination of the shared vehicles would result in much greater efficiencies compared to the current system of “hailing” yellow taxis from the street.
- ***Greater efficiencies result in cost savings for the customer.*** These cost savings are a result of reduced capital investment (fewer vehicles) and reduced operating expenses (fewer empty miles). In addition, cost savings can be expected from reduced parking expenses and increased time productivity (no need to focus on driving).
- ***Economies of scale are reached quickly with a shared, driverless vehicle fleet.*** Cost savings per trip/mile/day occur even for small fleets of shared vehicles. In Ann Arbor, the benefits of shared fleets can be reached with as few as 1,000 customers. At Babcock Ranch, costs are relatively constant above 5,000 customers and rise only slightly as the customer base changes from 5,000 to 1,000.
- ***Consumers experience greater convenience with a shared, driverless vehicle fleet.*** In Ann Arbor, a fleet size of 18,000 would result in average customer wait times of less than 1 minute. In Babcock Ranch, a fleet of 3,000 to 4,000 vehicles would result in average wait times of well under a minute. Similarly, in Manhattan, a fleet of 9,000 vehicles would result in average wait times of less than a minute.
- ***Results are consistent for a wide range of residential areas.*** Ann Arbor is representative of the traffic patterns of other medium-sized US cities, including Austin, TX, Orlando, FL, Rochester, NY, Sacramento, CA, and Salt Lake City, UT. Babcock Ranch is representative not only of new developments, but also of smaller suburban and urban areas. Finally,

¹⁸ Schaller Consulting, *The NYC Taxicab Fact Book* (2006).

Manhattan represents a densely populated urban center with roughly 70,000 people per square mile.

- ***Sustainability benefits are substantial with fleets of shared, driverless vehicles.*** Such a mobility system would improve transportation safety, improve roadway capacity, reduce congestion, increase energy efficiency, and improve land use (less land needed for parking). Advanced propulsion systems and specific-purpose vehicle designs would enable use of renewable energy sources and further increase energy efficiency.

APPENDIX:

METHODOLOGY FOR MODELING PERFORMANCE AND COST OF A SHARED, DRIVERLESS VEHICLE FLEET

Introduction

Consider a region such as those in our case studies where most of the trips for work, shopping, recreation and other purposes are intra-regional. We assume that a shared, driverless vehicle fleet operates within this region. The fleet operator optimizes the assignment of vehicles to customers so that low customer wait times are achieved and operational costs are minimized.¹⁹ Since customers can spontaneously request a vehicle, the operator must control the assignment of vehicles to customer trip requests in real-time. When a customer requests a shared vehicle using a smart phone app, the fleet operator will have information on that customer's current location, destination, and the vehicle type needed. We assume the customer needs the vehicle immediately, since this would be the hardest type of travel to accommodate. If a customer wanted to schedule trips ahead of time, this could also be accommodated, but would only reduce the costs being calculated here.

As shown in Figure 1, upon receiving the trip request, the fleet operator immediately determines which vehicle in their fleet could reach the customer first. The operator knows (via the Mobility Internet) the current location and state of all vehicles, can estimate when vehicles currently serving trips will complete them, and can estimate when each vehicle could potentially reach the new customer's location. The closest vehicle could be one that is currently idle, or, as shown in the figure, it could be a vehicle currently serving another trip. In either case the assigned vehicle travels to the new customer as soon possible. Through the Mobility Internet, the operator keeps the customer fully informed about the assigned vehicle's current location and expected arrival time. Note that assigning the closest vehicle to a customer not only minimizes wait time, but also minimizes empty vehicle miles, thereby minimizing operating costs. If the fleet operator can forecast travel patterns or if trips are scheduled ahead of time, customer wait time could be further reduced by pre-positioning vehicles where the operator predicts that they will be needed.

¹⁹ Note that while our discussion focuses on a single vehicle fleet, this doesn't preclude the possibility that there may be several competing vehicle fleets in an area.




Figure 1. Shared, Driverless Fleet Assignment Decisions

We need to estimate the performance of the system (customer wait time). In particular, we need to calculate how the number of vehicles in a shared fleet and the assignment decisions impact customer wait time. To provide customers with a better mobility experience than they currently have, customer wait times must be short. The definition of short depends on the situation. In our scenarios, Joe in Ann Arbor and Bob in Babcock Ranch have the alternative of driving their own cars. These are immediately accessible when they are at home, but there may be substantial walk time to access the cars when they are parked, for example, at a shopping center. Anne in Manhattan usually has substantial walk and/or wait time when accessing public transportation or yellow taxicabs. We'll estimate shared vehicle fleet costs assuming that average wait times (between calling for a vehicle and its arrival) will be 2 minutes or shorter. This is a fairly stringent standard that will lead to conservative cost estimates.

We also need to estimate the cost of a shared, driverless vehicle system. The primary costs are vehicle ownership and vehicle operation. Ownership costs include vehicle purchase, registration, and insurance; operating costs include fuel, maintenance, and repairs. The operating costs will depend not only on the length of trips served by the fleet, but also on the empty distance vehicles travel to customer locations.

To estimate system performance and cost, we develop relatively simple models that capture the key variables: the area of the region, the mean trip length, the mean trip rate and how this varies throughout the day, mean vehicle speed, the average fixed time needed per trip, the fleet size, and vehicle cost parameters. Using queueing and network modeling approaches, we develop an analytical model to relate these key variables to the performance and cost of a shared, driverless vehicle fleet. In this Appendix, we will describe the development of the analytical model, show the results of validating it with simulation models, and present the results from applying this model to the three case studies and to more general situations.

Base Model Definition for Analytical Model

To calculate the performance (i.e., customer wait time) and cost of the shared vehicle system, we develop two types of models: (a) an approximate analytical model that can be run quickly to

be used for performing sensitivity analyses, and (b) a simulation model to both validate the analytical model and handle extensions to our shared fleet system that are difficult to model analytically.

We define the region being served by the vehicle fleet as a square with area A , as shown in Figure A1.




Figure A1. Square region with uniformly distributed trip origins and destinations.

Within this region, trips are requested at an average rate λ and occur within a period P (this could be a peak travel period or any portion of the day you want to analyze). Each trip i has an origin, O_i , a destination D_i , and a request time t_i . We make the following simplifying assumptions about trips:

- Trip origins and destinations are uniformly distributed over region A ;
- Trips are requested at a constant average rate, λ , and the time between requests is exponentially distributed²⁰;
- Vehicles travel at a constant speed, v ; and
- There is one type of vehicle used for all trips.

The shared, driverless vehicle fleet has M vehicles for serving this region. The fleet is operated in the following simple manner:

- Trip requests are received at random time intervals by a centralized fleet coordinator, which knows the current location and status of all vehicles.

²⁰ While average trip rates clearly vary throughout the day, assuming an exponential distribution for inter-request times is quite realistic.

- When a new request is received, the coordinator scans through all vehicles, both idle and in-service, and determines which vehicle can reach the requestor first. This vehicle is assigned to that request.
- When a vehicle completes a trip, if it has been assigned to another trip, it departs immediately for that trip origin. If it has not been assigned to another trip, it waits at its current location until another trip is assigned.
- We assume that initially vehicles are randomly scattered over the region. This is a simplification; in fact, there probably would be multiple parking/maintenance facilities scattered over the region.

To model the performance of the vehicle fleet, we can think of the fleet as a queueing system where the vehicles are servers and trip requests are customers. Figure A2 shows how a vehicle spends its time. The service time in this system does not only include the time that the vehicle is with the customer (loaded), it also includes the time that the vehicle is driving without any customer (empty). We define the total service time for each customer as starting when the vehicle starts traveling to the customer and ending when the vehicle reaches the customer's destination and drops them off.




Figure A2. Vehicle service cycle, including both empty and loaded travel.

System performance depends on the average empty and loaded travel distance and time, as well as the average idle time relative to in-service time, the system's average capacity utilization. The average straight-line distance between random origins and destinations in a square of area A is $0.52\sqrt{A}$.²¹ We modify this by applying a circuitry factor, α , which recognizes that travel is constrained to a road network. Studies have shown that the road network adds 10 to 50 percent to straight-line travel distance depending on the type of area. Average loaded distance is

$$\begin{aligned} \bar{d} &= \text{average trip length assuming random origins and destinations} \\ &= \alpha \cdot 0.52\sqrt{A} . \end{aligned} \quad (1)$$

Loaded travel time is the time that it takes for the vehicle to traverse this distance plus the average time that the vehicle is stopped at the trip origin for the passenger(s) to board and the time stopped at the destination for passengers to alight. Let b = boarding time and a = alighting time. We assume these are constant for all trips. So,

$$\begin{aligned} t_l &= \text{average loaded travel time} \\ &= \frac{\bar{d}}{v} + b + a . \end{aligned} \quad (2)$$

²¹ Christofides, N. and S. Eilon (1969). An Algorithm for the Vehicle Dispatching Problem. *Operations Research Quarterly*, 20, 309-318.

The average empty distance for each trip is not so easily expressed. A trip request can be responded to either by a vehicle that is currently idle or by a vehicle that is currently in service but ending its trip close to the new request. So, this distance actually depends on both the distribution of idle vehicles over the region and the rate at which vehicles are completing their trips, as well as the distribution of their destinations.

To develop equations for the average empty travel distance and for the expected time that customers wait for a vehicle after making a trip request, consider the status of vehicles when a request is made. As shown in Figure A3, idle vehicles are distributed over the region waiting to be assigned.




Figure A3. When a trip request is made, either the closest idle vehicle is assigned or an in-service vehicle if it is completing its trip closer (i.e., in the dashed circle) soon enough to reach the customer first.

When a trip request is made, the closest idle vehicle could be assigned, but since we want to minimize customer wait time, the system will also consider in-service vehicles that might be ending their trips close to the customer and that could respond quicker. By “close” we mean that this in-service vehicle must be completing its current trip closer to the customer than the nearest idle vehicle (i.e., within the circle shown in Figure A3). The average empty distance and customer wait time depend on the distance to the closest empty vehicle and the probability, \tilde{P} , that an in-service vehicle will end its trip closer to the new customer in time to reach the customer first.

This distance to the closest idle vehicle depends on the number of idle vehicles and where they are located within the region. Define ρ as the average capacity utilization of vehicles. Then $(1 - \rho)$ is the probability that any individual vehicle is idle at any point in time. With a fleet of M vehicles, the average number of idle vehicles when a trip request is made is

$$\bar{I} = (1 - \rho)M. \quad (3)$$

Applying results for the average distance from the closest of many points to a random point in a square,²² the expected empty travel distance for the closest idle vehicle is

$$\bar{e}_I = \alpha \cdot k \sqrt{A/I} , \quad (4)$$

where $k = 0.52$ if the vehicles are randomly distributed over the area. However, while trip destinations are random, the process of assigning the closest vehicle to each trip request results in a significantly non-random distribution of idle vehicles. Based on simulation studies, $k \approx 0.70$.

The empty travel time and the wait time customers would experience if an idle vehicle responds to the trip request is

$$\bar{W}_I = \bar{e}_I / v . \quad (5)$$

However, we must now assess the probability that an in-service vehicle could end its trip closer (within the circle of radius r in Figure A3) to the new request and arrive at the customer first. This radius is the straight-line distance to the closest idle vehicle,

$$r = \bar{e}_I / \alpha . \quad (6)$$

The rate at which trips end within this circle is simply

$$\tilde{\lambda} = \lambda \pi r^2 / A . \quad (7)$$

So we can approximate probability \tilde{P} as simply the probability that an in-service vehicle ends its trip within a distance r of the customer before we expect the closest idle vehicle to reach the customer. Recognizing that trip completions, like trip requests, will be exponentially distributed,

$$\tilde{P} \approx \int_0^{\bar{W}_I} \tilde{\lambda} e^{-\tilde{\lambda}t} dt = 1 - e^{-\tilde{\lambda}\bar{W}_I} \quad (8)$$

\tilde{P} is an upper bound of the probability that an in-service vehicle will service the customer, because a vehicle could land in the region but still not be able to reach the customer before the closest idle vehicle. Also there is a chance (although it is small for realistic trip rates and vehicle fleet sizes) that the in-service vehicle may already have its next trip assigned.

Figure A4 shows that the time it takes an in-service vehicle to reach a customer has two components: t_c , the time from when the trip request is made until the time the vehicle completes its current trip; and t_e , the time for the vehicle to travel empty to the customer. So, to be assigned to serve the customer, an in-service vehicle must not only complete its current trip within r of the customer, but $t_c + t_e$ must be less than \bar{W}_I .

²² Christofides, N. and S. Eilon (1969). An Algorithm for the Vehicle Dispatching Problem. *Operations Research Quarterly*, 20, 309-318.




Figure A4. Two components of time for an in-service vehicle to respond to a customer call.

Given that an in-service vehicle ends a trip within r of the customer in the time interval $(0, \bar{W}_I)$ after the request is made, the expected time within the interval that it completes its trip is:

$$\bar{t}_c = \int_0^{\bar{W}_I} t \tilde{\lambda} e^{-\tilde{\lambda} t} dt / \tilde{P} = \left[1 - e^{-\tilde{\lambda} \bar{W}_I} (\tilde{\lambda} \bar{W}_I + 1) \right] / \tilde{\lambda} \tilde{P} \quad (9)$$

Assuming that the first in-service vehicle that ends its trip within r of the customer will be the one assigned (there is a small chance that another vehicle could a short time later end its trip even closer), the expected empty travel distance for the in-service vehicle would be

$$\bar{e}_N = \alpha \cdot \frac{2}{3} r. \quad (10)$$

The expected empty travel time will be

$$\bar{t}_e = \bar{e}_N / v. \quad (11)$$

The expected wait time by a customer served by a vehicle that is in-service when the trip request is made, given that an in-service vehicle is used, is

$$\bar{W}_N = \bar{t}_c + \bar{t}_e. \quad (12)$$

Now that we have the expected empty travel distance and customer wait times for either a currently idle or currently in-service vehicle serving the new trip request, the overall expected empty travel distance and wait time are just the sum of these weighted by the probability that each is used:

$$\bar{e} = (1 - \tilde{P}) \bar{e}_I + \tilde{P} \bar{e}_N. \quad (13)$$

$$\bar{W} = (1 - \tilde{P}) \bar{W}_I + \tilde{P} \bar{W}_N. \quad (14)$$

Starting with our input parameters A , λ , v , M , a , b , and α , equations (1) through (14) therefore give a way to estimate the performance of a shared, driverless vehicle fleet. However, using these equations assumes that we know the system's capacity utilization. Unfortunately, we don't. A basic queueing principle, however, is that in steady-state a queueing system's average capacity utilization is its average arrival rate divided by its average service rate, multiplied by the number of servers. The service rate is the inverse of the average service time \bar{S} so we can write

$$\rho = \lambda / \left(\frac{M}{\bar{S}} \right) = \lambda \bar{S} / M . \quad (15)$$

From the point of view of the vehicle fleet, a vehicle is serving a customer once it starts traveling empty to that customer, so the service time is the sum of the empty and loaded travel times,

$$\bar{S} = \frac{(\bar{e} + \bar{d})}{v} + b + a . \quad (16)$$

We can calculate the system's capacity utilization from (15) and (16) once we know \bar{e} . Unfortunately we need the capacity utilization to calculate \bar{e} . So, we can't get a closed form solution for our performance metrics in terms of our input parameters. However, a very simple iterative procedure can do this. The basic approach in this procedure is:

1. Compute the minimum possible value of ρ, ρ_{min} , by setting $\bar{e} = 0$ and using equations (15) and (16). Set $\rho = \rho_{min}$.
2. Using ρ , compute \bar{W} and \bar{e} using equations (1) through (14), then use the new \bar{e} to compute ρ_{next} from (15) and (16).
3. If $\rho = \rho_{next}$, then stop. This value of ρ is what we want to use in further calculations. If not, then increment ρ and return to step 2. However, if $\rho = 1$, then stop; the system doesn't have enough vehicles.

This procedure works very fast and the converged value of capacity utilization yields the average empty travel distance via (14) and the customer wait time via (13).

System cost can then be easily calculated. Let the input cost parameters be defined as follows:

- c_v = the daily amortized ownership cost for driverless vehicles, and
- c_m = operating cost per mile for driverless vehicles (including gas, maintenance, repairs).

Then the total system cost per day would be

$$TC = c_v M + c_m \lambda (\bar{e} + \bar{d}) . \quad (17)$$

In (17), the lambda should be in trips/day since we are dealing with daily cost.

For comparison with other systems, we'll often express cost as the cost per trip-mile:

$$\text{Cost per tripmile} = TC / \lambda \bar{d} . \quad (18)$$

Validation using Simulation Model

The base model specified above is approximate. It makes a number of assumptions to simplify the analysis, including:

- There are always idle vehicles somewhere in the system when a trip request is received;
- An in-service vehicle is assigned to a new customer if one completes its previous trip close to the customer before the closest idle vehicle would be expected to reach the customer; and

- If an in-service vehicle is assigned to a new customer, it is the first one that is expected to complete its trip close to the customer.

These are very reasonable assumptions and should introduce only small errors in estimating customer wait time, empty vehicle miles, and system costs.

Before applying the analytical model to our case studies, however, we want to validate it against a simulation of shared, driverless vehicle fleet operations that does not make these simplifying assumptions. The simulation will also assume that trip origins and destinations are randomly distributed over the square region, but will randomly generate specific origins and destinations. The simulation model randomly generates trip request times over a period, and, as requests are received, assigns vehicles to trip requests to minimize customer wait times as described earlier. That is, when a request is received the model evaluates the expected time when each idle vehicle could reach the customer and the expected time each in-service vehicle, after completing its assigned trips, could reach the customer. The vehicle that is expected to reach the customer first is assigned. The simulation model calculates customer wait times and empty vehicle miles, and we compare these results to those calculated by the analytical model.

Note that the simulation model takes much longer to run than the analytical model, especially for values of λ that are typical of our case studies. For large values of λ and M the simulation model can take well over an hour to run on a standard laptop computer whereas the analytical model runs in a fraction of a second.

The graphs in Figure A5 show system performance predicted by the analytical model compared to the simulation model²³ for the case where $A = 100$ square miles and $\lambda = 1000$ trips/hour. These values might be typical for a small urban area or a portion of a larger urban area in which a shared vehicle fleet might operate. In all cases we use $v = 20$ mph, $\alpha = 1.4$, and a and $b = 1$ minute. The graphs show the results for values of M (vehicle fleet size) between 520 and 1000. With fewer than 520 vehicles, the capacity utilization approaches 100 percent and the analytical procedure does not converge. Increasing vehicles beyond 1000 adds cost without improving performance. These results indicate:

- The analytical model predicts average capacity utilizations, wait times, and total costs typically within 5 percent of simulated values over a wide range of fleet sizes (30 to 90 percent).
- The analytical model appears to slightly over-predict empty miles per trip at larger fleet sizes. Empty mile estimates for the two models are within 2 percent over most of the range in Figure A5, but grow to about 8 percent different for 1000 vehicles.

²³ The simulation results are from running the simulation for ten replications of a 10-hour period starting with randomly distributed vehicles and then running the simulation for a one-hour warm-up period before collecting statistics.

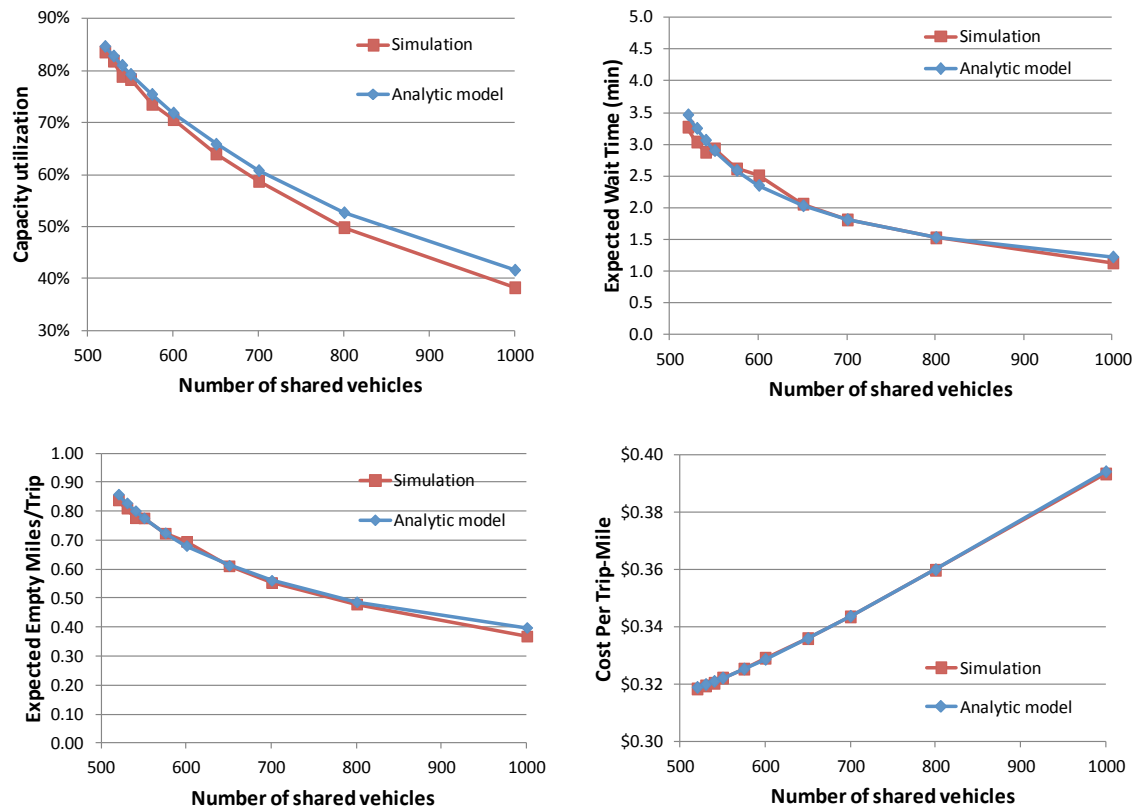


Figure A5. Analytical vs. simulation model results (capacity utilization, expected customer wait time, expected empty travel distance, and cost) for a region with an area of 100 square miles and a trip rate of 1000 trips/hour.

Table A1 compares analytical and simulation model results over a wide range of model parameters. The range tested covers realistic region areas and trip rates that we might expect to analyze. The table shows the values for the same four model outputs as shown in Figure A5 for these different parameter values with the percent difference between analytical and simulation model results in the last four columns of the table. The table shows:

- The analytical model generally gives excellent results: most values are estimated to within 5 percent of simulation model values (green boxes), with only 6 values out of 200 off by over 10 percent (red boxes).
- Vehicle cost estimates are extremely close to simulated values in all cases.
- Empty miles are usually extremely close, but the model over-predicts by over 10 percent in one case.
- Customer wait times are generally close but the analytical model diverges from the simulated values by about 10 percent in several cases.

In summary, comparison with simulated values shows that the analytical model can be used with confidence to analyze shared, driverless fleet performance.

Table A1. Comparison of Analytical and Simulation Model Results.

	Model Parameters			Analytical Results				Simulation Results				% Difference			
				Ave.	Ave.	Total		Ave.	Ave.	Total		Ave.	Ave.	Total	
	Lambda	Area	Fleet	Cap.	Wait	Empty	Cost/	Cap.	Wait	Empty	Cost/	Cap.	Wait	Empty	Cost/
	(trips/hr)	(sq. mi.)	Size	Util.	Time (min.)	Miles/ Trip	Trip Mile	Util.	Time (min.)	Miles/ Trip	Trip Mile	Util.	Time	Trip	Trip Mile
Base Case	1,000	100	525	83%	3.30	0.84	\$0.32	84%	3.27	0.84	\$0.32	0%	1%	0%	0%
	1,000	100	530	82%	3.22	0.82	\$0.32	82%	3.04	0.81	\$0.32	1%	6%	2%	0%
	1,000	100	540	81%	3.03	0.80	\$0.32	79%	2.88	0.78	\$0.32	2%	5%	2%	0%
	1,000	100	550	79%	2.90	0.78	\$0.32	78%	2.93	0.78	\$0.32	1%	-1%	0%	0%
	1,000	100	575	75%	2.57	0.72	\$0.33	74%	2.62	0.72	\$0.33	2%	-2%	0%	0%
	1,000	100	600	72%	2.35	0.68	\$0.33	71%	2.51	0.70	\$0.33	2%	-6%	-2%	0%
	1,000	100	650	66%	2.03	0.61	\$0.34	64%	2.06	0.61	\$0.34	3%	-2%	0%	0%
	1,000	100	700	60%	1.80	0.56	\$0.34	59%	1.81	0.55	\$0.34	3%	0%	1%	0%
	1,000	100	800	53%	1.53	0.49	\$0.36	50%	1.53	0.48	\$0.36	6%	0%	2%	0%
	1,000	100	1,000	41%	1.22	0.40	\$0.39	38%	1.14	0.37	\$0.39	8%	8%	8%	0%
Low Demand Case	100	100	59	88%	9.91	2.46	\$0.38	85%	9.48	2.13	\$0.37	3%	5%	16%	3%
	100	100	60	84%	9.31	2.20	\$0.37	82%	8.23	2.01	\$0.37	3%	13%	10%	1%
	100	100	65	75%	7.67	1.88	\$0.37	75%	6.93	1.80	\$0.37	1%	11%	4%	0%
	100	100	70	69%	6.63	1.73	\$0.38	66%	5.97	1.67	\$0.38	4%	11%	4%	0%
	100	100	75	63%	5.90	1.62	\$0.38	63%	5.30	1.53	\$0.38	1%	11%	6%	1%
	100	100	80	59%	5.34	1.52	\$0.39	57%	4.88	1.44	\$0.39	3%	9%	5%	0%
	100	100	90	52%	4.61	1.37	\$0.40	50%	4.32	1.31	\$0.40	3%	7%	5%	1%
	100	100	100	46%	4.09	1.25	\$0.42	45%	3.77	1.18	\$0.41	3%	9%	7%	1%
	100	100	110	41%	3.73	1.16	\$0.43	39%	3.59	1.11	\$0.43	4%	4%	5%	0%
	100	100	150	29%	2.89	0.93	\$0.50	29%	2.80	0.90	\$0.50	3%	3%	3%	0%
High Demand Case	10,000	100	4,600	91%	1.49	0.38	\$0.29	91%	1.62	0.39	\$0.29	0%	-8%	-2%	0%
	10,000	100	4,700	88%	1.32	0.36	\$0.30	89%	1.52	0.37	\$0.30	0%	-13%	-5%	0%
	10,000	100	4,800	86%	1.20	0.34	\$0.30	86%	1.33	0.35	\$0.30	0%	-9%	-3%	0%
	10,000	100	5,000	83%	1.03	0.30	\$0.30	83%	1.15	0.32	\$0.30	0%	-10%	-4%	0%
	10,000	100	5,250	78%	0.90	0.27	\$0.30	78%	0.97	0.28	\$0.30	1%	-8%	-3%	0%
	10,000	100	5,500	74%	0.80	0.25	\$0.31	74%	0.88	0.26	\$0.31	0%	-9%	-4%	0%
	10,000	100	6,000	68%	0.68	0.22	\$0.32	67%	0.72	0.22	\$0.32	1%	-6%	-3%	0%
	10,000	100	7,000	58%	0.55	0.18	\$0.33	57%	0.56	0.18	\$0.33	2%	-2%	-1%	0%
	10,000	100	8,000	51%	0.47	0.15	\$0.35	49%	0.47	0.15	\$0.35	2%	0%	1%	0%
	10,000	100	10,000	40%	0.38	0.13	\$0.39	39%	0.37	0.12	\$0.39	4%	4%	4%	0%
Small Area Case	1,000	25	290	83%	2.22	0.55	\$0.34	83%	2.12	0.54	\$0.34	0%	4%	1%	0%
	1,000	25	300	80%	2.02	0.52	\$0.34	80%	2.06	0.53	\$0.34	0%	-2%	-2%	0%
	1,000	25	310	77%	1.83	0.49	\$0.34	77%	1.85	0.49	\$0.34	0%	-1%	-1%	0%
	1,000	25	325	73%	1.65	0.46	\$0.34	73%	1.66	0.46	\$0.34	0%	0%	0%	0%
	1,000	25	350	67%	1.42	0.42	\$0.35	68%	1.46	0.42	\$0.35	-1%	-3%	-1%	0%
	1,000	25	400	58%	1.16	0.36	\$0.37	58%	1.25	0.37	\$0.37	-1%	-7%	-4%	0%
	1,000	25	500	46%	0.90	0.29	\$0.40	45%	0.93	0.29	\$0.40	1%	-2%	-1%	0%
	1,000	25	600	38%	0.77	0.25	\$0.43	37%	0.77	0.25	\$0.43	3%	-1%	0%	0%
	1,000	25	800	28%	0.61	0.20	\$0.50	27%	0.61	0.20	\$0.50	3%	0%	1%	0%
	1,000	25	1,000	22%	0.53	0.17	\$0.58	21%	0.51	0.17	\$0.57	4%	3%	3%	0%
Large Area Case	1,000	400	950	87%	5.59	1.38	\$0.31	86%	5.68	1.38	\$0.31	1%	-2%	0%	0%
	1,000	400	975	85%	5.06	1.30	\$0.31	83%	4.94	1.29	\$0.31	2%	2%	1%	0%
	1,000	400	1,000	82%	4.63	1.24	\$0.31	80%	4.50	1.22	\$0.31	3%	3%	2%	0%
	1,000	400	1,050	78%	4.00	1.13	\$0.31	76%	3.98	1.12	\$0.31	2%	0%	1%	0%
	1,000	400	1,100	74%	3.59	1.05	\$0.32	72%	3.67	1.05	\$0.31	2%	-2%	0%	0%
	1,000	400	1,200	67%	3.03	0.93	\$0.32	65%	3.06	0.93	\$0.32	3%	-1%	0%	0%
	1,000	400	1,250	64%	2.85	0.88	\$0.33	62%	2.75	0.85	\$0.33	3%	4%	4%	0%
	1,000	400	1,300	62%	2.68	0.84	\$0.33	60%	2.69	0.83	\$0.33	3%	0%	1%	0%
	1,000	400	1,400	57%	2.43	0.77	\$0.34	55%	2.50	0.78	\$0.34	4%	-3%	-1%	0%
	1,000	400	1,500	53%	2.24	0.72	\$0.35	51%	2.14	0.69	\$0.35	4%	5%	5%	0%

Use of Models for Case Studies

Earlier sections of this document discussed the case studies to which we applied this basic model. These case studies examined the performance of a shared, driverless vehicle fleet as an alternative for personally owned vehicle trips within the urban area of Ann Arbor, Michigan and the proposed eco-city of Babcock Ranch, Florida, and as an alternative for yellow taxi trips within Manhattan, New York. The basic data used for these analyses was discussed earlier. This data was used to calculate parameters for the base model as shown in Table A2.

Table A2. Case Study Model Input Data

parameter	Case Studies					
	Ann Arbor		Babcock Ranch		Manhattan	
	peak period	average daytime	peak period	average daytime	peak period	average daytime
Area, A (square miles)	128	128	20	20	12	12
Trip request rate, λ (trips/hour)	55,000	44,000	9,800	9,500	26,000	21,000
Period length, P (hours)	1	12	1	12	1	18
Vehicle speed, v (miles/hour)	30	30	25	25	10	11
boarding time, b (minutes)	1	1	1	1	1	1
alighting time, a (minutes)	1	1	1	1	1	1
circuitry factor, cf	1.0	1.0	1.4	1.4	1.1	1.1
Vehicle ownership cost, c_v (\$/vehicle-day)	\$19	\$19	\$19	\$19	\$19	\$19
Vehicle operating cost, c_m (\$/vehicle-mile)	\$0.20	\$0.20	\$0.20	\$0.20	\$0.20	\$0.20

Data from the 2010 Census and the 2009 National Household Travel Survey were used to study the opportunity for a shared, driverless vehicle fleet to serve local Ann Arbor trips instead of privately owned vehicles. As discussed in the text, the analysis focused on the 90 percent of vehicles that are driven less than 70 miles per day, which was 120,000 vehicles on a typical day. The Ann Arbor urban area is 128 square miles and we assume that these vehicles were used for trips within this region. These vehicles were responsible for 528,000 trips per day, for an average of 4.4 trips per vehicle per day. These internal trips averaged 5.8 miles in distance. The pattern of vehicle usage throughout a typical weekday is shown in Figure A6. Most vehicle trips are made in the 12-hour period from 7am to 7pm, so approximating that all trips occur in this period, the average trip rate is 528,000 trips/12 hours = 44,000 trips/hour. The evening peak hour is about 25 percent above the average, so a peak trip rate is about 55,000 trips/hour. The circuitry factor is set at 1.0 so that the average trip length equals the observed 5.8 miles.




Figure A6. Vehicle Use by Time of Day for Privately Owned Vehicles on Weekdays in Ann Arbor Area - 2009 NHTS data

Figure A7 shows the performance and cost for a shared, driverless fleet serving these Ann Arbor trips. The left chart shows peak period and average daily wait time and empty miles. In the peak period a fleet of at least 16,000 needs to be in operation to meet demand. Customer wait times are quite sensitive to fleet size at very high utilizations, so using a slightly larger fleet of 18,000 vehicles (87 percent utilization in the peak period) yields an average customer wait of about one-third of a minute and an average empty travel distance of less than 0.20 miles. An 18,000-vehicle fleet would yield average wait times and empty distance over the entire day of 0.2 minutes and 0.1 miles/trip.

The average cost per trip-mile over the day for an 18,000 vehicle fleet would be \$0.31/trip-mile with two-thirds of this being operating cost and one-third being ownership cost.²⁴ Comparing the cost graph to the wait time and empty mile graph shows that cost is relatively insensitive to the size of the fleet. Total costs increase by less than 15 percent if the fleet is increased from 13,500 to 21,000 vehicles. While ownership costs clearly increase with fleet size, operating costs – the largest part of cost – decrease as empty miles come down. On the other hand, wait time performance is very sensitive to fleet size. As mentioned earlier, a fleet of less than 16,000 vehicles would not even be able to meet peak period demand, so performance would be unacceptable with a fleet smaller than this. Above 16,000 vehicles, peak period wait times come down rapidly even though costs only increase modestly to achieve this performance. So, sizing the fleet at 18,000 vehicles or slightly more makes sense to provide adequate service at reasonable cost.

²⁴ For comparison, the cost of owning and operating a car is \$0.20/mile operating cost plus \$16/day ownership cost for a mid-size car (2012 AAA estimates). For these Ann Arbor internal trips, the average vehicle was driving 26 miles per day. So this cost per trip mile for a personally owned vehicle is $\$0.20 + \$16/26 = \$0.82/\text{trip-mile}$.

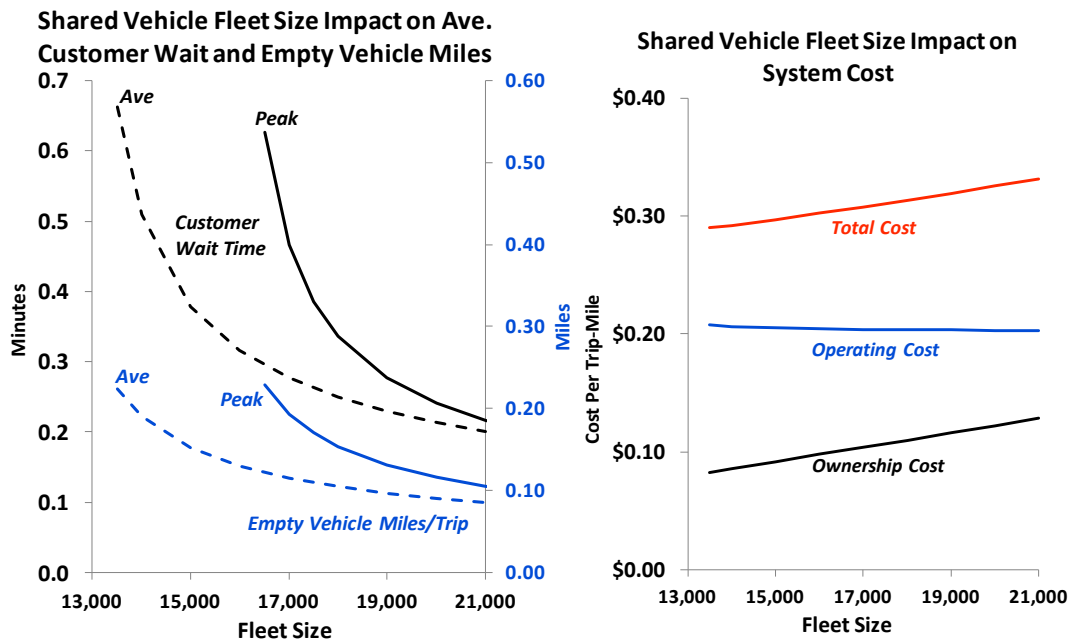


Figure A7. Shared fleet performance and cost for Ann Arbor Case Study.

The customer wait time and empty distance performance shown in Figure A7 may seem unrealistic. How can a shared fleet get a vehicle to a customer spontaneously making a trip request with so little wait time? To answer this question, we need to think about where vehicles are when a customer requests a trip. In the peak period, the 18,000 vehicles are 87 percent utilized. That means that 13 percent of the vehicles (over 2,300 vehicles) are idle. With 2,300 vehicles spread over 128 square miles, we have 18 vehicles per square mile. In addition, there are almost 1,000 vehicles per minute that are completing trips and are available to serve customers. Therefore, there are a lot of vehicles available to serve trips even in a peak period.

Figure A8 shows a snapshot of vehicle locations when a customer requests a trip during the Ann Arbor peak period from a simulation of the 18,000 vehicle shared fleet. The figure shows the customer location and the locations of idle vehicles when the request is made. Note that the locations of idle vehicles are not quite as smoothly distributed as one might expect. This is because the vehicle assignment process leads to some non-randomness in these locations. In fact, the customer in this case appears to be located in an area somewhat sparsely populated with idle vehicles. There are about 3,100 idle vehicles at the time represented in this figure (an above average amount), with the closest one being 0.34 miles from the customer. In addition, there is an average of almost 10 vehicles per minute completing trips in each square mile that could also potentially serve this customer. In this particular case, the closest idle vehicle can reach the customer first, so the customer waits 0.7 minutes for a vehicle. The intuition here is that with the large number of idle and soon-to-be-available in-service vehicles per square mile, it would be surprising if average wait times were not small.

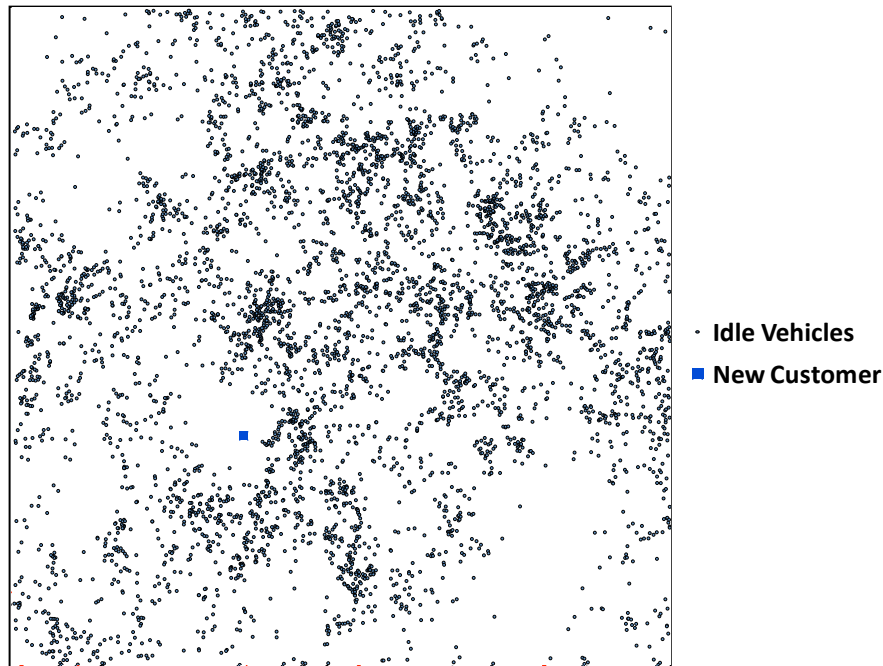


Figure A8. Location of idle when a customer places a trip request in the Ann Arbor peak period.

This customer experiences about twice the average peak period wait time. This raises the question about wait time variation. Although the average wait time is short, are there a significant number of customers that have very long wait times? Simulation analyses have shown that the standard deviation in wait time is 0.7 to 0.9 times the mean wait time and the 95th-percentile wait is less than 1 minute in the Ann Arbor case. So, the longest wait times are not that long.²⁵

The above analysis has shown that the cost of a shared vehicle fleet is low assuming that all car trips internal to Ann Arbor were made using the shared fleet. What if only a portion of these trips use the shared fleet? Also, we've assumed a single fleet. What if we have fleets of different types of vehicles for different trip needs? In either case, the trip demand rate would be well below the level assumed in this analysis. How would this change overall cost?

We did a sensitivity analysis on average daily trip demand rate, varying it from 50 to 50,000 trips per hour. At each demand rate, fleet sizes were set so that average customer wait times were 1 minute. Figure A9 charts the resulting cost per trip mile as a function of the demand rate. System cost turns out to be very insensitive to the demand rate. As the trip rate decreases from 50,000 trips per hour to 5,000 trips per hour, cost per mile increases by only 7 percent. However, when the trip rate falls below 1,000 trips per hour, costs start escalating fast.

²⁵ What is an acceptable customer wait time is an important question for operating a shared fleet. A wait time of several minutes may be acceptable if the customer is informed of the expected arrival time of the vehicle, with updates on the vehicles location as it approaches. In modeling the system, we've aimed to achieve very short average wait times to be conservative – i.e., this leads to a relatively higher shared fleet cost.

The intuition behind this is shown by what happens to capacity utilization in the chart on the right. When the trip rate is between 5,000 and 50,000 trips per hour, an average wait time for 1 minute can be achieved while keeping capacity utilization at 85 percent or above. However, maintaining this average wait time with lower trip rates requires substantial excess vehicles causing ownership cost to escalate significantly. Hence, scale economies for operating a shared fleet set in relatively quickly – at about 500 to 1,000 trips/hour for an Ann Arbor type area. This has significant implications for implementing shared fleets and for operating shared fleets of multiple vehicle types.

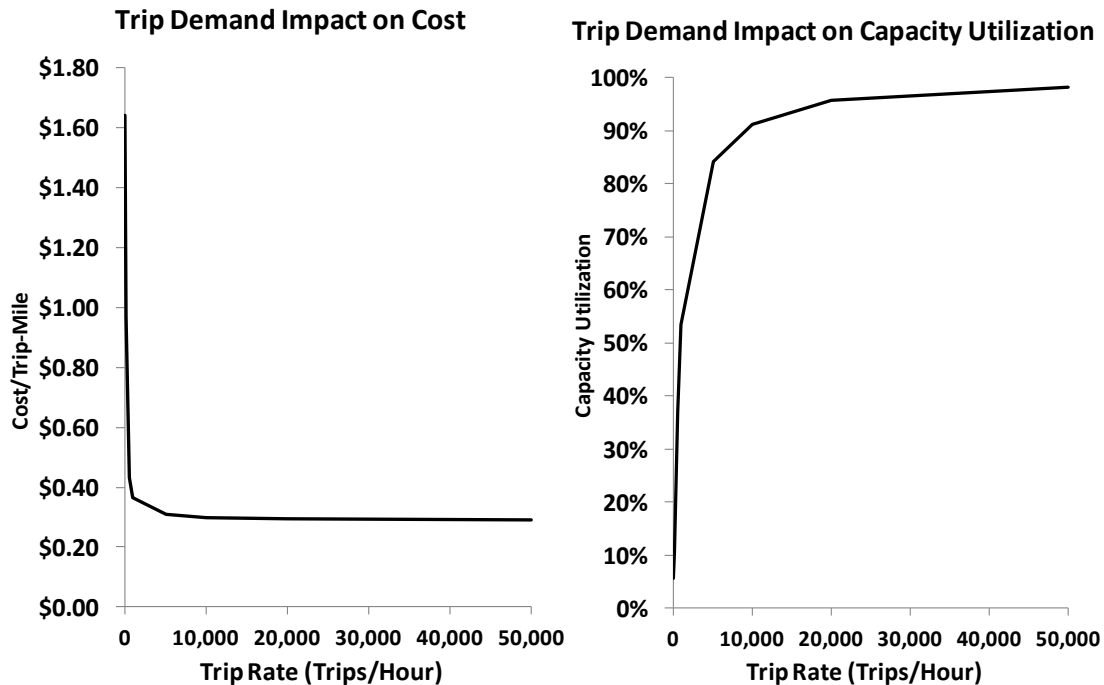


Figure A9 Sensitivity of total shared fleet cost and capacity utilization to trip rate. Fleet sizes set so that average customer wait times are one minute

The major modeling assumptions that we make are that:

1. Trips occur within a square region.
2. Origins and destinations are uniformly distributed over the region.
3. Trip requests have a constant average rate and inter-request times are exponentially distributed.

These are all simplifying assumptions that have varying impacts on the attractiveness of a shared vehicle fleet. Assuming a square region as opposed to a more realistic shape with the same overall area has very little impact on the results, based on simulation studies. Likewise, the assumption of exponentially inter-request times is fairly realistic in periods where the average trip request rate is constant.

The assumption of a constant average rate is a more critical assumption – as shown in Figure A6, the trip rate is far from constant over the course of a day. To some degree, this can be addressed by applying the model to short periods of say 1 to 4 hours over which the average rate is fairly constant and then aggregating the results. To a degree this is what we have done in the case studies by doing a separate analysis for the peak period versus the overall daily average.

By far the most critical assumption is the second one. Assuming that origins and destinations are uniformly distributed over the entire region implicitly assumes that the number of trip origins and destinations are roughly in balance in any portion of the area. So vehicles generally don't have to be moved far after completing their trips. This assumption also implies that at any point in time the next trip is equally likely to be anywhere within the region. This makes responding to trip requests difficult – vehicles can't be moved to where they are likely to be used next because you don't know where that is. So, customer wait times are likely to be at their highest in this case. By assuming uniformly distributed origins and destinations, this analysis may result in higher expected wait times, but lower empty miles than situations where origins and destinations are imbalanced.

Additional analyses of cases where trip origins and destinations are imbalanced are required. For example, analyses of cases where more trips start in one place and more trips end in another at a given time of day, such as the morning rush hour when trips start in residential areas and end in commercial or industrial areas. In the Babcock Ranch case study, the potential imbalances in trips were analyzed and these increased empty vehicle miles to 20 percent of loaded miles from the 5 percent that would be expected if origins and destinations were balanced. This increased overall cost for the shared fleet by 10 percent.

In addition to added cost, imbalances in trip origins and destinations require more sophisticated methods for moving vehicles after they have completed trips. That is, by collecting data on where imbalances in trip origins and destinations exist, vehicles can be pre-positioned after completing trips to areas where you expect them to be needed. This incurs more extra empty miles than you would need in a balanced situation, but this pre-positioning allows customer wait times to be reduced.

Summary

In this Appendix, we described the methodology used to calculate performance (customer wait time) and cost for a shared, driverless vehicle fleet. We developed two types of models: an approximate analytical model that could be run quickly to develop sensitivity analyses, and a simulation model to both validate the analytical model and estimate extensions to our system that were difficult to model analytically. When comparing results from the analytical and simulation models across a wide range of parameters, we found that the analytical model could be used with confidence to estimate performance and cost of the shared, driverless vehicle fleet. We then presented details of how the models were used in our case studies, with a focus on Ann Arbor. We described the sensitivity analyses undertaken to test our results, and found that the results were robust. Finally, we examined the major simplifying assumptions made in the models, and determined that they were all reasonable, although additional analysis of cases where trip origins and destinations are imbalanced is desirable.

ACKNOWLEDGEMENTS

We would like to acknowledge Jeffrey Sachs, Director of the Earth Institute, for being an enthusiastic and supportive champion of our work. Significant contributions were made to this project by Frances Liu and Nicole Olaes and we thank them for their input.