

THE TRAVEL AND ENVIRONMENTAL IMPLICATIONS OF SHARED AUTONOMOUS VEHICLES, USING AGENT-BASED MODEL SCENARIOS

Daniel Fagnant
The University of Texas at Austin
6.9E Cockrell Jr. Hall
Austin, TX 78712
danfagnant@hotmail.com

Kara M. Kockelman
(Corresponding author)
Professor and William J. Murray Jr. Fellow
Department of Civil, Architectural and Environmental Engineering
The University of Texas at Austin
kkockelm@mail.utexas.edu
Phone: 512-471-0210

Presented at the 93rd Annual Meeting of the
Transportation Research Board in Washington DC, January 2014
and published in *Transportation Research Part C*, Vol 40 (2014): 1-13

ABSTRACT

Carsharing programs that operate as short-term vehicle rentals (often for one-way trips before ending the rental) like Car2Go and ZipCar have quickly expanded, with the number of U.S. users doubling every one to two years over the past decade. Such programs seek to shift personal transportation choices from an owned asset to a service used on demand. The advent of autonomous or fully self-driving vehicles will address many current carsharing barriers, including users' travel to access available vehicles.

This work describes the design of an agent-based model for Shared Autonomous Vehicle (SAV) operations, the results of many case-study applications using this model, and the estimated environmental benefits of such settings, versus conventional vehicle ownership and use. The model operates by generating trips throughout a grid-based urban area, with each trip assigned an origin, destination and departure time, to mimic realistic travel profiles. A preliminary model run estimates the SAV fleet size required to reasonably service all trips, also using a variety of vehicle relocation strategies that seek to minimize future traveler wait times. Next, the model is run over one-hundred days, with driverless vehicles ferrying travelers from one destination to the next. During each 5-minute interval, some unused SAVs relocate, attempting to shorten wait times for next-period travelers.

Case studies vary trip generation rates, trip distribution patterns, network congestion levels, service area size, vehicle relocation strategies, and fleet size. Preliminary results indicate that each SAV can replace around eleven conventional vehicles, but adds up to 10% more travel

distance than comparable non-SAV trips, resulting in overall beneficial emissions impacts, once fleet-efficiency changes and embodied versus in-use emissions are assessed.

INTRODUCTION

Autonomous vehicles may very well be on our streets and highways by the end of the decade. As of September 2013, Google had logged over 500,000 miles driven on public roadways using cars equipped with self-driving technology (Fisher 2013). California, Nevada and Florida have issued enabling autonomous vehicle legislation, with Washington, D.C. and nine other states following close behind (CIS 2013). This new technology has the power to dramatically change the way in which transportation systems operate. While autonomous vehicle impacts for traffic safety and congestion have been predicted in some detail, potential behavioral shifts and resulting environmental impacts have received little attention.

Once such avenue for behavioral shifting is carsharing. Carsharing programs (such as ZipCar and Car2Go) exist around the globe, and the number of U.S. users has grown from 12 thousand in 2002 to over 890 thousand in January of 2013 (Shaheen and Cohen 2013). Car sharing programs operate as short-term rentals, where members are able to rent a vehicle, typically located at an on-street parking location, drive to a nearby destination, and then release the rental when finished with their trip. Some programs like Car2Go allow rentals by the minute, while others like ZipCar require longer rental intervals, and additional annual membership and/or applications fees are often required.

U.S. National Household Travel Survey (NHTS) data (FHWA 2009) suggest that less than 17% of newer (10 years old or less) household vehicles are in use at any given time over the course of an “average” day, even when applying a 5-minute buffer on both trip ends; this share falls to just 10% usage when older personal vehicles are included and no buffers applied. In short, *Americans have many more cars than they need* to serve current trip patterns in most locations. Yet carsharing members still face nearby vehicle availability barriers: if a person is worried that he may be stranded, wait a long time, or walk a great distance in order to access a vehicle, he may opt to drive his own car instead.

Shared autonomous vehicles (SAVs), also known as autonomous taxis or aTaxis (Kornhauser et al. 2013), provide a solution, with members able to call up distant SAVs using mobile phone applications, rather than searching for and walking long distances to an available vehicle. Moreover, they provide carsharing organizations with a way of seamlessly repositioning vehicles in order to better match demand. These SAVs are assumed to be fully self-driving without any need of human operation, other than information regarding a traveler’s destination. In this way, SAVs could transform transportation for many: from an owned asset to a subscription or pay-on-demand service, at least in areas where population densities make such systems economically viable.

Such advances may provide significant environmental benefits, particularly in the form of reduced parking and vehicle ownership needs. Moreover, there is the potential for additional vehicle-miles traveled (VMT) reduction: Shaheen and Cohen (2013) estimate that North American carsharing members reduced their driving distances by 27%, with approximately 25%

of members selling a vehicle and another 25% forgoing a vehicle purchase. It is not clear how much of these effects would apply to SAVs, which may be more convenient and potentially more used than human-driven shared vehicles. Each SAV also can/will move itself (unoccupied) to the next traveler or relocate while unoccupied to a more favorable location, for lower-cost parking and faster future passenger service.

Some researchers have sought to model this phenomenon. Ford (2012) developed a shared autonomous taxi model, and relied on travelers to walk to fixed taxi stands, rather than allowing the SAVs to travel driverless to their next passenger, or relocate to more optimal locations. Kornhauser et al. (2013) investigated this idea further, exploring dynamic ride sharing implications for all person-trips across New Jersey. In this model, one or more passengers boarded at fixed stations, where taxis wait a given time before departing, and all passengers having similar destinations share a ride (with an unlimited SAV fleet size assumed). Burns et al. (2013) also investigated this setup, with case study examples of travelers in Ann Arbor (Michigan), Manhattan (New York), and Babcock Ranch, a new small town in Florida. A number of Burns et al.'s modeling assumptions were similar to this paper's investigation, though their focus was on cost comparisons with personal cars. For example, they estimated that per-mile costs fall between \$0.18 to \$0.34 per mile when switching to SAVs, depending on annual mileage driven, in the Ann Arbor case study. Other benefits may be estimated using Fagnant and Kockelman's (2013) suggestions for AV savings (on insurance, parking, time value, and fuel), versus technology costs (\$10,000), alongside an extra assumed \$5,000 per-shared-vehicle-year operating and management cost. In this scenario, when assuming a \$5-per-hour travel-time benefit (from less burdensome travel, lowering one's effective VOTT) and \$5-per-weekday parking costs, total realized added benefits over a two-year SAV lifespan more than double the extra added technology and operating costs.

This paper focuses on SAVs' travel and environmental implications, and uses different assumptions to model a much smaller share of such trips (around 3.5%, rather than 100%), while also modeling directional distribution effects by time of day. The investigation tests four different SAV relocation strategies, seeking to relocate unused SAVs to more favorable locations in order to reduce future traveler wait times. This paper examines ownership and VMT questions, investigating how many household-owned vehicles may be replaced by a fleet of SAVs, how much new travel may be induced (due to unoccupied-vehicle travel, rather than more or longer trips, which may emerge from lower perceived travel-time costs or new travel by those currently without driver's licenses), what factors influence such outcomes, and the resulting travel and environmental implications across a variety of reasonable model scenarios. SAVs are assumed to be shared serially in this investigation by different travel parties, and future work could incorporate dynamic ride-sharing as well.

An SAV model was specified, programmed (from the ground up, in C++), and applied, to gauge the environmental impacts of this novel opportunity. Here, the model operates by generating person-trips in each zone, on 5-minute intervals, over 100 days, with each trip having a destination and departure time. These trips are served by SAVs, which ferry travelers between origins and destinations as quickly as they can. Traffic congestion is modeled by time of day, with vehicle speeds slowed at peak hours. When individual SAVs are not in use, a virtual fleet manager uses a variety of strategies to relocate SAVs to more favorable positions, attempting to

reduce wait times for future travelers. Travel implications and emissions inventories are assessed by estimating changes in life-cycle inventories, based on changes in the overall vehicle fleet (with any pickup trucks and SUVs substituted for mid-sized SAV sedans), Total VMT is tracked (with added VMT emerging from unoccupied SAVs travel, for relocation or new traveler pickup), along with cold-starts savings (thanks to busy SAVs) and parking benefits (thanks to a shared fleet).

Twenty-five scenario variations were also run, in order to appreciate the impacts of changing many of the base-case scenario assumptions. These scenario variations examine the impacts of altering trip generation rates, the degree of trip centralization (how many trips are generated in the city center versus the outlying areas), the size of the overall service area, the frequency by which travelers return home via SAV, the extent of peak congestion, the effects of various SAV relocation strategies used and the impacts of limiting the overall size of the SAV fleet. While this model is by no means perfect (for example, an actual city's transportation network and origin-destination tables could be used to refine these evaluations in future work, and just a single vehicle type is used), this investigation provides a preliminary look into the potential implications of a new SAV system operating within the urban environment.

MODEL SPECIFICATION

The system operates by first generating trips throughout *a gridded city*. The program runs through the vehicle-assignment model 20 times, to determine approximately how many SAVs will be needed and where they should be placed at the start of the day. This process links trips to SAVs, generating a new SAV for every traveler that has been waiting at least 10 minutes (two time steps). The (integer-rounded) average number of SAVs generated in these 20 model runs is used as the fixed fleet size for all subsequent multi-day analyses. After this warm-start, the program is re-run, but SAVs can no longer be generated.

The city is composed of quarter-mile by quarter-mile grid cells or zones, to generate and attract trips. The base-case scenario's service area is a ten-mile by ten-mile square area ($40 \times 40 = 1600$ zones), about twice the size of Austin, Texas's Car2Go geofence (the area where the vehicle needs to be by the end of the rental). Other Car2Go geofences range from around 40 square miles (in Portland, Oregon) to 60 square miles (Washington, D.C.) and 75 square miles (Seattle, Washington). Zones in the outermost service-area corners are assumed to generate trips at one rate, zones in the city's "core" (the central area with dimensions half the city's length and width) at a higher rate, and zones within 2.5 miles of the city center (the outer urban core) a third, intermediate rate. Trip generation rates in zones lying between these three points are linearly interpolated between the closest two points, depending on their centroids' Euclidean distances to the city center. That is to say, for zones lying within the analysis area but more than 2.5 miles from the city center, rates will be linearly interpolated between the outermost service area and outer core rates; while for zones at 2.5 miles or closer to the city center, rates will be interpolated between the inner and outer core rates. Trips are generated using Poisson distributions per *5-minute time step* and assigned throughout a 24-hour period, based on the temporal distribution of NHTS trip-start rates (FHWA 2009).

Each generated trip is then assigned a destination, using the following steps:

- 1) Sample trip distance ($D = 1$ to 15 miles).
- 2) Sample E-W direction: $\Pr(E) = \alpha \times (\#zones\ E) / (\#zones\ E + \#zones\ W) + (1 - \alpha) \times 0.5$
- 3) Sample N-S direction: $\Pr(N) = \alpha \times (\#zones\ N) / (\#zones\ N + \#zones\ S) + (1 - \alpha) \times 0.5$
- 4) Choose destination: $\Pr(E-W = 0 \ \& \ N-S = D) = 1 / (4D + 1)$;
 $\Pr(E-W = 0.25\ mi. \ \& \ N-S = D - 0.25\ mi.) = 1 / (4D + 1)$,
 $\Pr(E-W = 0.50\ mi. \ \& \ N-S = D - 0.50\ mi.) = 1 / (4D + 1)$, ...
 $\Pr(E-W = D \ \& \ N-S = 0) = 1 / (4D + 1)$.
- 5) If destination is not in the service area, return to step 2.
- 6) If 20 consecutive destinations selected outside of service area, return to step 1.
- 7) Valid destination has been located. Proceed to next trip.

Each of these steps (including α 's derivation) is explained in greater detail below:

Step 1: Trip distances (D) are randomly drawn (in quarter-mile increments) from 1 to 15 miles away, based on NHTS (trip-distance) data, as shown in Figure 1.

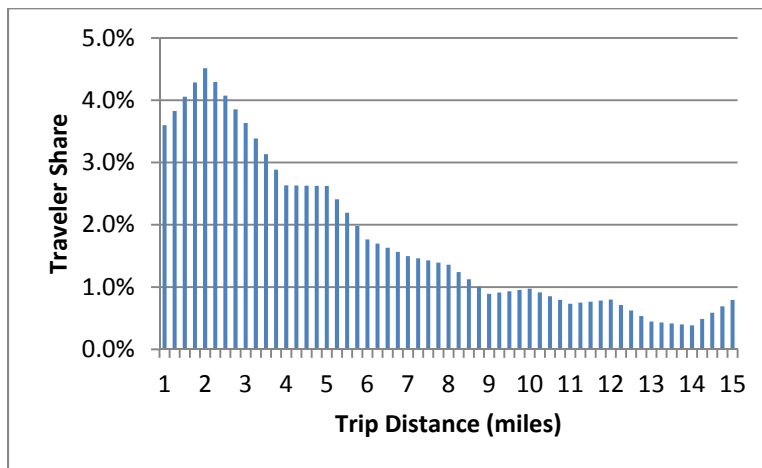


Figure 1: Trip Distance Distribution (Source Data: NHTS [FHWA 2009])

The 15-mile limit was established because most trips cannot travel for much more than 15 miles without leaving the service area, and forcing higher limits would lead to overly-high trip attractions at far corners (since the corners would be the only reachable destinations at certain, long-trip distances). In real-world business applications, out-of-service area rentals could potentially be made, with added user costs for vehicle relocation or extended-time rentals, though this phenomenon is not modeled here.

Steps 2 & 3: East-west and north-south cardinal directions are then selected. Probabilities are linearly based on two components: the share of zones lying in each direction from the origin (effectively resulting in higher attraction levels towards the city center, particularly at high values of α) and a constant factor of 0.5 (effectively resulting in similar probabilities when choosing between north and south, and east and west directions, particularly at low values of α). The “attraction factor” (α) specifies relative degree of strength that each of these factors has when determining destination direction. Before noon, trip attractions pull more strongly towards the city center ($\alpha = 1$); after noontime, the core zones’ attraction factor (α) lessens (to about

0.77), such that the total number of trips entering the urban core (defined as the innermost 5-mile by 5-mile area) roughly equals the number leaving over the course of the day. This also helps keep trip-making balanced, such that nearly all travelers return “home” before the next day’s simulation. This trip attraction factor is determined for each 24-hour period by generating a series of sample trips and using the method of successive averages (Bell and Iida 1997) to pick a value for α . While this generates different values of α from day-to-day, due to the randomization processes embedded with the trip generation procedure, overall inflows and outflows remain close to even over each 24-hour period.

Step 4: The destination zone is then selected using a uniform distribution from among all possible combinations of the chosen travel distance and cardinal (north-south and east-west) directions. For example, if Steps 1 through 3 produced a 2.25 mile trip distance in a north and west direction, there will be equal 10% chances of traveling 2.25 miles due north, 2 miles north and 0.25 miles west, 1.75 miles north and 0.5 miles west, and so forth. As such, all possible destinations in the north and west directions at 2.25 miles away will have an equal chance of being selected.

Step 5: A new location is drawn if the chosen destination lies outside the service area. In such situations, the trip distance is retained, and a new trip direction is sampled.

Step 6: A new trip distance is drawn if 20 consecutive destinations lie outside the service area. In such settings, a new trip distance and a new trip direction are sampled.

Step 7: When a trip destination is located within the service area, the destination is assigned to the traveler and the program proceeds to the next trip.

Each assigned trip then looks, in turn, for the nearest available SAV, up to a maximum (5-minute) travel time. This process operates by iterating through all travelers using random ordering (in each 5-minute interval) to find available SAVs in the same zone, which are claimed if found. Next, all travelers who did not find an available SAV in their same zone look one zone out and claim that vehicle, if located. This process continues until all travelers have either found an SAV, or cannot locate one within a 5-minute travel distance. Those who cannot be served within the 5-minute period are moved to a “wait list” that is serviced at the beginning of the next period, before any new traveler assignments are made. In the event that two or more locations have free SAVs the same distance away, an SAV is chosen from the zone with the most available SAVs, to help prevent stockpiling too many SAVs in a single zone.

SAVs travel a fixed distance per time period, with effective speeds equal to 3 times the number of zones that SAVs are assumed to travel (in that scenario), in miles per hour, in order to reflect intersection delays and other factors (e.g., an SAV that travels 10 zones in one 5-minute time step has an effective speed of 30 mph¹). These fixed speeds are pre-set by time of day to determine the maximum number of zones an SAV can travel in 5-minutes, reflecting higher-speed off-peak periods and lower-speed peak periods. SAVs may travel in the horizontal or vertical directions to their destinations (which lie in all directions from their origins), but not

¹ This notion of 10 zones per 5-minute interval comes from the fact that $30 \text{ mph} \times 4 (1/4 \text{ mi. zones per mile}) / 12 (5\text{-minute intervals per hour}) = 10$.

diagonally, reflecting a strictly *grid-based street pattern*. Once an SAV is assigned to a traveler, it drives to the traveler (if not already in the same zone), then begins traveling to the final destination. This continues for more than one period, as necessary, and the SAV is released once it arrives at the destination. After dropping a traveler off, SAVs do not continue moving for the rest of the period, reflecting non-moving pick-up and drop-off time (and averaging a little over two minutes in the base-case scenario). SAVs can park in any zone and for any duration of time while waiting for the next traveler. The program also determines whether each traveler wishes to return home via SAV; if so, a dwell time (an activity duration) is chosen from a one of eight distributions, depending on arrival time (grouped by 3-hour periods, as shown in Figure 2) and based on NHTS dwell-time durations (FHWA 2009). This process allows travelers journeying at earlier hours to stay at their destination longer, on average, than those traveling later in the day.

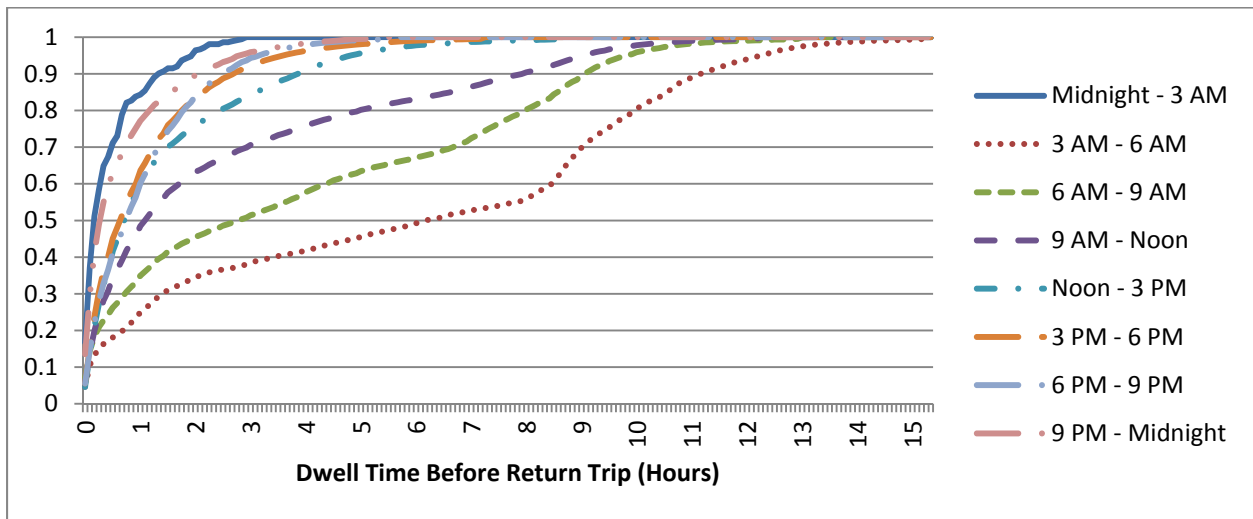


Figure 2: Cumulative Distributions for Dwell Times Before Return Trips
(Source Data: NHTS 2009)

Once all travelers that can be reached in a given 5-minute interval been served, the remaining, unoccupied vehicles that have not traveled during the interval consider relocating to a more advantageous position to shorten wait times in the next interval. This is conducted using a virtual fleet manager that determines where all unoccupied vehicles will be during the beginning of the next interval, and issues relocation commands to all SAVs. Though additional relocation efficiencies could be incorporated to further minimize traveler wait times, some limits are placed in order to serve the secondary goal of reducing unoccupied VMT. Four strategies (labeled R1 through R4) are pursued to these ends, with increasingly smaller evaluation areas for vehicle relocation. Each of these strategies is used in subsequent ordering during this relocation phase: R1 is followed by R2, and then R3, with R4 acting last. Any SAV that has moved during the 5-minute interval will not move again during the same interval; thus, if a vehicle relocates using strategy R1, it will not relocate again using strategy R3 until the next time interval.

The *first vehicle-relocation strategy (R1)* attempts to match expected demand over large areas with available SAVs. The R1 strategy divides the city into 25 large blocks, each two miles square. A “block balance” value is calculated for each block, based on the expected demand for and supply of SAVs in the coming 5-minute period (using waiting travelers plus expected new

trips to be generated, minus the number of free SAVs that will be in the block at the start of the next period). The sum of all block balances is equal to zero, with each block balance calculated using the formula:

$$Block\ Balance = SAVs_{Total} \left(\frac{SAVs_{Block}}{SAVs_{Total}} - \frac{Demand_{Block}}{Demand_{Total}} \right) \quad (1)$$

If a given block has 10% or more vehicles than ideal², excess free SAVs are pushed into adjacent blocks; if a block has 10% or fewer vehicles than ideal, free SAVs (if available) are pulled from adjacent blocks.

Strategy R1 operates by first selecting the block with the highest absolute block balance value (that farthest from 0). If the block balance is positive, this block pushes SAVs that have not moved during the current 5-minute interval into adjacent blocks; if negative, it pulls similarly free SAVs from adjacent blocks. First, the number of available SAVs is determined (in the selected block if pushing, and in adjacent blocks if pulling), as well as the block balances of adjacent blocks. The number of SAVs to be moved in each direction is determined by first selecting an SAV to be pushed onto the nearby block with the lowest block balance, or pulled from the nearby block with the highest block balance and at least one available SAV. Next, block balances are updated, and this process is repeated until no available SAVs remain to be pushed or pulled, the 10% threshold is no longer exceeded, or SAV relocation will improve overall block balances by less than 1 (e.g., pushing an SAV with a +8.4 block balance onto another block with a +6.9 block balance).

After the number of SAVs to be moved in each direction is determined, a corner of the block is chosen using a random draw. Moving in a clockwise direction, the strategy seeks available SAVs one zone out from the (pulling) block. If an available SAV is located and there is demand in that direction, the SAV will be moved. Next, available SAVs two zones from each edge with remaining demand are sought and allocated. This process continues until each vehicle that will be moved in the block has been assigned a direction. Figure 3 illustrates one possible assignment process, where the numbers of SAVs to be moved in each direction are shown in adjacent blocks in parentheses, and the SAVs to be moved are shown using letter-number combinations in the central block.

² The ideal share is the proportion of expected demand divided by the proportion of available SAVs for a given block, thus the ideal share for any block is 1.0.

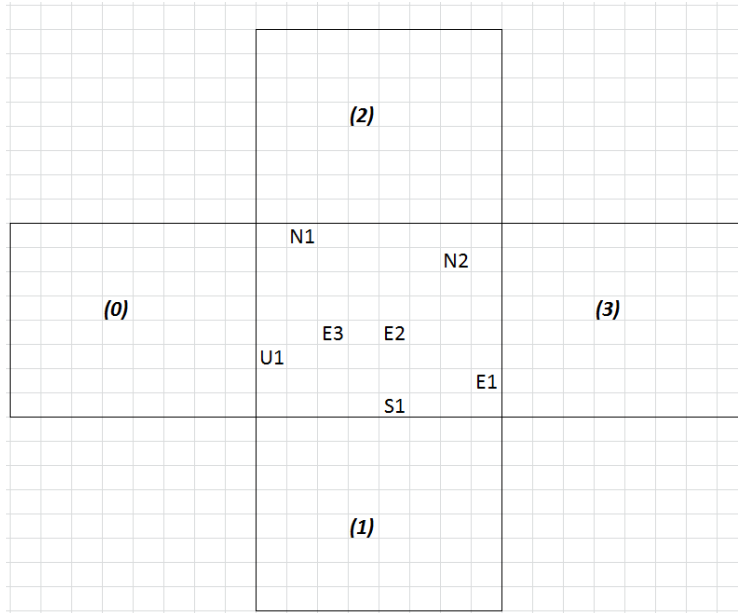


Figure 3: Selecting Vehicles for Relocation (Pushing) Using Strategy R1

In Figure 3’s example scenario, two SAVs are to be pushed northward, three eastward, and one southward; and the northwest corner is randomly drawn to begin this process. As such, N1 will be chosen first, to move north, followed by E1 to travel east, and S1 to travel south. Looking two zones away from the borders, N2 would be chosen to head north, then SAVs will be searched for at increasing distances from the eastern border, until E2 and E3 are identified. Available SAV U1 will remain unmoved.

Once SAVs are identified to relocate and their directions are assigned, they first move just over the boundary into the new block. If the initial zone of entry is unoccupied by any other (available) SAV, the pushed SAV stops. If one or more SAVs is already present, it will continue to move in the same direction if an unoccupied zone is reachable within the 5-minute travel distance interval. If an unoccupied zone is not reachable, the SAV looks for the nearest same-direction reachable zone with just one SAV (and then with just two SAVs, if no reachable zone with just one SAV can be located), and so forth, until a destination is identified.

At this point, all identified SAVs in the target block will have been pushed to (or pulled from) nearby blocks. The algorithm next identifies the block with the next-largest absolute block balance value, and this process repeats until all blocks have either been served (by pushing or pulling SAVs) or all remaining blocks have a discrepancy less than 10% (of excess demand or supply of SAVs) from their target value. The *second strategy (R2)* is essentially the same as R1, though 100 one-square-mile blocks are used instead of the 25 larger ones.

The *third strategy (R3)* seeks to “fill in the white (zero-SAV) spaces” using available SAVs occupying zones where two or more SAVs are predicted to be free at the start of the next interval. These SAVs look for unoccupied locations within a half mile (two zones) where there are no other SAVs in any adjacent zone. The *final strategy (R4)* acts as a “stockpile management” strategy, shifting available SAVs into adjacent zones when the free-SAV

imbalance (in the next interval) is three or more SAVs. For relocation strategies R3 and R4, a tie-breaker prioritizes moving vehicles to zones with higher generation rates. Figure 4 illustrates how all four strategies tend to work, for a given period's supply versus demand conditions, and using actual values from model runs.

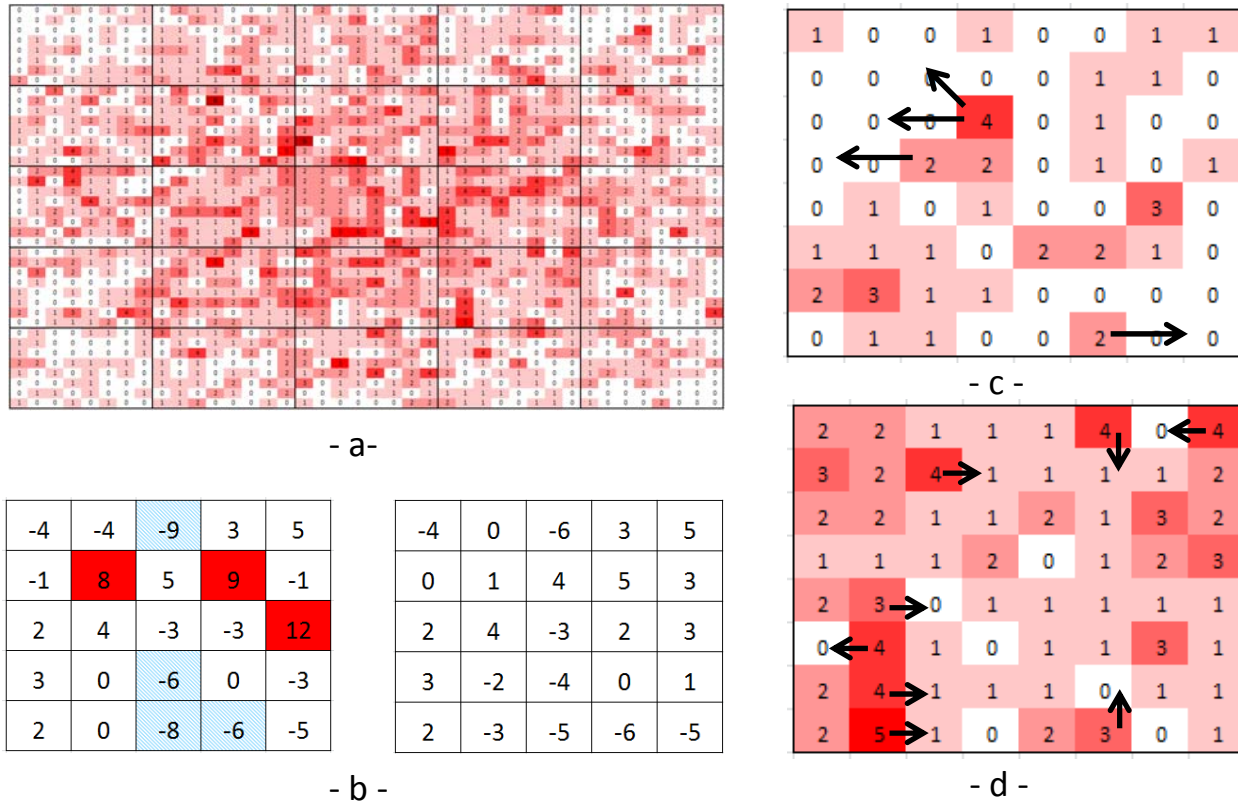


Figure 4: Relocation Strategy Examples: 4a. Example Starting Location of all SAVs; 4b. Block Balances Before Strategy R2 Relocation (Left) and After (Right); 4c. R3 Strategy Illustration on a Subsection of Zones; 4d. R4 Strategy Illustration on a Subsection of Zones

Figure 4a shows locations of all SAVs just before relocation begins, with darker zones representing greater numbers of SAVs in those zones. Figure 4b gives predicted block balances, and depicts how strategy R1 operates. Highlighted blocks with negative numbers represent blocks that seek to pull vehicles from adjacent blocks, while highlighted blocks with positive numbers attempt to push excess vehicles into adjacent blocks. Figure 4b (right side) shows block balances after rebalancing using strategy R1. Note that some balances remain somewhat uneven, due to SAV availability limitations during the time interval. Figures 4c and 4d demonstrate how vehicles would relocate using strategies R3 and R4, respectively, provided that there are available SAVs in the originating zones.

For further illustration of overall model operation, Figure 5 depicts the movement of one SAV's daily journey, beginning at S and ending at E. Solid arrows represent vehicle movements delivering passengers, dashed arrows represent vehicle relocation movements, the region's core zone is shaded (representing the area used to determine the relative "attraction factor" (α) for

trips starting after noon), and block divisions are shown as solid black lines. Note that each arrow represents net movement during a single, 5-minute interval, so short diversions to pick up travelers are not shown and a single trip may span multiple time intervals. Also note that all trips begin and end in a zone's centroid, so there may be a small degree of extra VMT from vehicles traveling internally within the origin and destination zones.

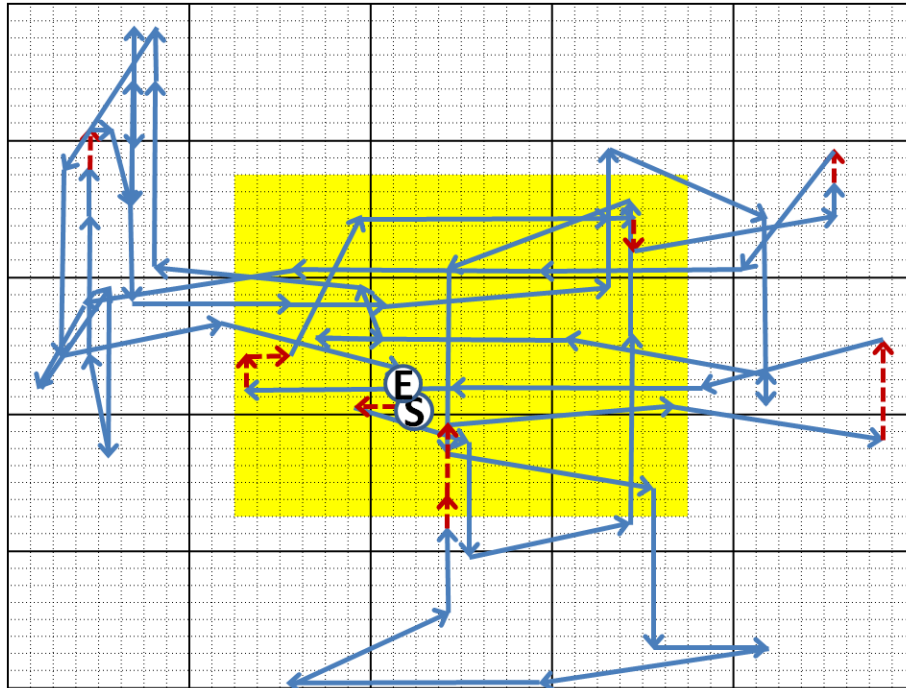


Figure 5: Travel Patterns and Operation for an Example SAV, 7 AM to Noon

Several observations can be made from this particular vehicle's journey, illuminating key model operations. First, all four relocation strategies are visible here. At the far right, strategy R1 is being used (with the SAV is relocating across 2-mile blocks, and thus farther than R2 allows). Just below and to the right of the starting location, an R2 move is visible (here, the SAV is traveling more than two zones but is not crossing a larger 2-mile block boundary). The remainder of relocations are just one-zone (R4) or two-zone (R3) moves, typically not crossing an R1 or R2 block boundary. Also, congestion impacts during peak periods can be observed in the upper-left blocks, where per-period maximum travel distances are shorter than observed elsewhere. Finally, also in the upper-left corner, backtracking may be observed, where the SAV drops off a passenger with a very short wait interval, before taking him back to his original starting location.

In addition to relocation strategies, vehicle refueling, cleaning and maintenance were explicitly modeled. SAVs were assumed to be relatively fuel-efficient vehicles, operating at 40 miles per gallon with 12-gallon tanks. Once any SAV went below a 2-gallon reserve and had dropped its passengers off, a ten-minute wait period was incurred in the same zone, rendering the SAV temporarily unavailable for service, in order to allow for refueling and minor/routine cleaning. More serious SAV maintenance and cleaning was assumed to occur during low-demand times (e.g., overnight) and so not impact the simulated operations.

CASE STUDY RESULTS

A *base-case scenario* (S0) model run was conducted, using 100 simulated days with summary statistics reported upon completion. This scenario is intended to represent a mid-sized city, perhaps the size of Austin, Texas, with about 3.5% of formerly human-driven trips within the service area now being served by SAVs. Table 1 notes S0's base assumptions (all of which were later varied in sensitivity analyses, described below). According to 2009 NHTS data, commuter speeds averaged 28.7 mph for metropolitan statistical areas with populations between 1 and 3 million persons (Santos et al. 2011). Additionally, Schrank, et al.'s (2012) Urban Mobility Report estimates that Austin peak travel times are 32% greater than off-peak times. Using these data points, off-peak speeds were assumed to be somewhat higher than baseline commuter speeds (which were probably measured during congested times of day) and set at 33 mph for the base-case scenario. Furthermore, since SAVs are assumed to be traveling more often in the urban core (and therefore in generally more congested areas), differences in peak and off-peak speeds were assumed to be more pronounced, and a 21 mph base-case scenario congested speed was assumed. Relative trip generation rates were based on Austin's population patterns, with average census block population density 2.5 miles from the city center showing 90% of the core population density, and just 30% of core density observed 7 miles out. Table 2 summarizes the overall travel impacts.

Table 1: S0 Base-Case Model Parameters

Parameter	Value
Service area	10 mi. x 10 mi.
Outer service area trip generation rate	9 trips per zone per day
Outer core trip generation rate	27 trips per zone per day
Innermost zone trip generation rate	30 trips per zone per day
Off-peak speed	33 mph
Peak speed	21 mph
AM peak	7 AM - 8 AM
PM peak	4 PM - 6:30 PM
Trip share returning by SAV	78%

Table 2: S0 Model Results (Daily Averages & Standard Deviations of Key Behaviors)

Category	Measure	Mean	S.D.
Trips & SAVs	# Person-trips per day	60,551	336
	# SAVs	1,688	0
	# Person-trips per SAV per day	35.87	0.20(5.15)
Wait Time	5-minute wait periods	249	109
	Avg. wait time per person-trip	0.295	0.014(0.61)
	# Un-served person-trips (across all days)	0	0
	% Waiting 5 min +	0.4%	0.2%
Trip Miles	Total VMT per day	332,900	2,200
	Unoccupied VMT per day	32,060	350
	Avg. person-trip distance (mi.)	5.43	0.01(3.33)
	Unoccupied miles per person-trip	0.53	0.011
	% Induced travel (added VMT)	10.7%	0.1%
Usage	Min. # SAVs not in use	19	11
	Min. # SAVs unoccupied	54	22
	Max # share in use / moving	98.87%	0.63%
	Max # share occupied	96.80%	1.35%
Vehicle Starts	SAV warm starts per person-trip	0.73	0.01
	SAV cold-starts per person-trip	0.054	0.001
	# Warm-starts per day	44,190	370
	# Cold-starts per day	3,287	53

Note: Reported means represent daily averages and standard deviations represent averages of day-to-day standard deviations (or average within-day standard deviations when shown in parentheses) across 100 simulated days.

Complete model results show how each base-case SAV serves approximately 31 to 41 travelers per day, with average wait times under 20 seconds³. Less than 0.5% of travelers waited more than five minutes, only three persons per day (0.005% of travelers) waited ten minutes or more, and no traveler waited 15 minutes or more on any of the 100 days (so none was considered “un-served”, with a 30+ minute wait time). This very high level of service should be acceptable to almost anyone traveling by SAV, especially since many travelers are more willing to accept an occasional, longer wait during the PM peak traffic period, where congestion is present and demand is highest. Standard deviations show that, while the average and the “average of average” wait times is 0.30 minutes (or 18 seconds), the average standard deviation of wait times across a day’s trips is 0.61 minutes (due largely to relatively high peak-period wait times); and the average of these 100 standard deviations (as shown in Table 2) is just 0.01 minutes. About 11% of total VMT stems from vehicles relocating to new zones while unoccupied, and SAV relocation to cheaper parking areas during times of low demand was not modeled.

³ While a 20-second average wait time appears quite optimistic, it reflects the phenomenon that an available SAV will almost always be within one or two zones, for the great majority of travelers under these scenario conditions. There may be additional time required for internal-zone travel, pickup and dropoff; and this extra time is modeled by not allowing SAVs to continue traveling for the remainder of the 5-minute interval after dropping a passenger off.

During the heaviest-use interval (typically just after 5 PM), more than 97% of vehicles were occupied, and just 1% were idle (versus relocating), indicating very high SAV utilization levels (in contrast to the NHTS' maximum-use statistic of 10 percent, as discussed in this paper's introduction). There are just 0.054 cold-SAV starts per person-trip simulated here, versus about 0.64 cold starts per person-trip across the U.S. today.⁴

NHTS records show that each licensed driver in the U.S. averages 3.02 private motor-vehicle-trips per day and the U.S. has 0.99 owned or leased household vehicles per licensed driver (Santos et al. 2011). Thus, Table 2's total trip count results suggest that this SAV system has 20,049 member drivers who may normally possess 19,849 personal vehicles. Yet here just 1,688 SAVs are used, suggesting that each SAV has the ability to replace nearly 12 privately owned vehicles, on average. Such findings indicate that almost eleven parking spaces can be eliminated for every SAV. However, this analysis examines only shorter (under 15-mile) trips (to stay within the 10-mile x 10-mile modeled fence); if longer trips are included, SAV use rates may be lower. Moreover, vehicle use rates may be higher in specific contexts: for example, Seattle survey data (PSRC 2006) suggest that household vehicles are used more intensively than the U.S. fleet: 16% of vehicles were reported in use during the highest use 5-minute interval in Seattle compared to just 12% nationally, for vehicles 10 years old or newer. If vehicle utilization rates for are similarly higher to the Seattle profile, a single SAV may substitute for 8 or 9 vehicles (rather than nearly 12, as simulated here). Also, these results do not consider average vehicle occupancies (in SAVs or conventional vehicles). As such, reported person-trips shown here should be interpreted as person-trip parties, with one or more travelers, rather than as single-occupant vehicles only. SAV occupancies may be closer to national averages, around 1.67 person miles per vehicle mile (Santos et al. 2011).

Emissions and Energy Implications

Chester and Horvath's (2009) life-cycle inventory estimates were used to evaluate the SAV system's possible emissions and energy use implications, as shown in Table 3. This evaluation assumes that the existing U.S. light-duty-vehicle fleet (BTS 2012), consisting of passenger cars (53% of fleet), pickup trucks (14%), and SUVs and minivans (33%), will be replaced with conventionally fueled SAV sedans. No electric, hybrid-electric, or alternative fuels are assumed here (for the SAVs or the fleet being replaced). While this assumption may work well when replacing a relatively small share of the vehicle fleet, as investigated here, some pickup trucks, SUVs and other larger vehicles may be needed in any SAV fleet, especially as more travelers and trips shift to SAVs. New SAV sedans may also be smaller than the nation's current average sedan size, though this effect too was not modeled. Emissions and energy impacts are estimated based on vehicle operation (in-use, VMT-based emissions and energy), vehicle manufacture (embodied energy), vehicle parking infrastructure (embodied and via parking space maintenance), and trip-start emissions differences. All four impacts categories are influenced by fleet change, with vehicle operation also influenced by new/induced VMT (from travel of empty SAVs), parking influenced by reduced needs (due to fewer vehicles), and starting emissions influenced by changes in total number of vehicle starts, as well as the share of cold starts. While

⁴ Kang and Recker (2009) note that the U.S. EPA estimates that it takes approximately one hour of idle time for a catalytic converter to cool and be considered a "cold start", and they estimate that around 68% of U.S.-vehicle trips (with internal combustion engines) are cold starts.

the overall number of vehicles will be lower, this analysis assumes that the total manufacture rates for new vehicles will remain about the same. This is because SAVs will travel many more miles per year than conventional vehicles, and so will wear out much sooner (in terms of years) as their mileage rises faster. This also assumes that SAVs will be able to travel more miles before replacement due to unoccupied travel, but is roughly negated by shorter lifespans, since auto parts may wear out through differing combinations of miles driven and years of wear.

Table 3: Potential Environmental Impacts of Introducing SAVs (per SAV introduced)

Environmental Impact	Sedan (Passenger Car) Life-Cycle Inventories (Values not Shown for Pickup Trucks, & SUVs)				Average US Light-duty Vehicle vs. SAV Sedan Emissions Inventories			
	Operating (Running)	Manufac.	Parking	Vehicle Starts	Average LDV	SAV Total	Difference	% Change
Energy use (GJ)	890	100	15	0	1230	1087	-144	-12%
GHG (metric tons)	69	8.5	1.2	0	90.1	85.0	-5.1	-5.6%
SO ₂ (kg)	3.9	20	3.6	0	30.6	24.6	-5.9	-19%
CO (kg)	2100	110	5.2	1400	3833	2546	-1287	-34%
NO _x (kg)	160	20	6.4	32	243	200	-43	-18%
VOC (kg)	59	21	5.2	66	180	92	-88	-49%
PM ₁₀ (kg) ⁵	20	5.7	2.7	2.0	28.2	26.4	1.8	-6.5%

These results indicate beneficial energy use and emissions outcomes for all emissions species when shifting to a system of SAVs, assuming the same trip pattern/demand schedule is maintained (e.g., SAV users do not start making more or longer trips). The criteria pollutants evaluated here are sulfur dioxide (SO₂), carbon monoxide (CO), oxides of nitrogen (NO_x), volatile organic compounds (VOC), and particulate matter with effective diameter under 10 μm (PM₁₀). Greenhouse gas (GHG) reductions and total energy use reductions are also anticipated. Under our modeling and vehicle assumptions, using Chester and Horvath's estimates, VOC and CO emissions will experience the most significant reductions, largely due to the substantial quantities generated during vehicle starts (both cold and warm). PM₁₀ shows little reduction, along with GHG (even though the SAVs are assumed to be smaller and more fuel efficient than the average U.S. LDV, and so require less energy in their manufacture and operation), thanks to some added driving (for vehicles to access their travelers and, to a lesser extent, to relocate in anticipation of demand imbalances over space [thus shortening average response/passenger-wait time]). These results also assume that overall manufacturing needs will remain very similar, since per day SAV usage levels will be much higher, so there will be a greater (1.5 to 2 year) turnover rate. This being noted, newer vehicles may be more environmentally friendly than older ones, as U.S. Corporate Average Fuel Economy requirements tighten and technologies improve on conventional and alternative-fuel vehicles. Other factors may mitigate new running costs. For example, traditional carsharing arrangements in North America show members' average VMT

⁵ Using EPA's MOBILE6 model, Chester and Horvath (2009) estimated no cold start emissions of PM₁₀. EPA's new MOVES model corrects this deficiency, and suggests that each cold start emits about the same weight of PM₁₀ emissions that would be generated in a mile of travel, at 25 mph average speed. Similarly, the average cold start generates about the same level of PM_{2.5} that would come from 2.5 miles of travel, but we do not have embodied energy emissions estimates for PM_{2.5}, to define a total reduction here.

falling by 27% (Shaheen and Cohen 2013), and their use of carpooling and non-motorized modes (biking and walking) rising, though transit use often falls (Martin and Shaheen 2011). Moreover, Shoup (2007) estimates that an average of 30% of traffic in central business districts is generated by vehicles seeking to find a parking space close to their occupants' final destination, so reduced parking needs could lead to further emission reductions, along with congestion improvements (thus reducing idling and other driving-based emissions). Finally, SAVs should reduce crash occurrences (and delay vehicle replacement a bit, reducing manufacturing emissions), and may help ease congestion due to AVs being able to drive more smoothly and intelligently, particularly if other non-shared AVs are acting in concert (KPMG and CAR 2012). Such considerations suggest considerable potential for further energy and environmental savings, beyond what is estimated here.

SENSITIVITY ANALYSIS: MODEL AND PARAMETER VARIATIONS

In order to appreciate how different settings impact SAV benefits and travel outcomes, model specifications and assumptions were adjusted to define 26 distinct alternative scenarios (each run with 100 simulated days of data). Eight types of variation were tested. *Trip generation scenarios* (ST1, ST2 and ST3) changed the base case's (S0's) overall trip generation rates by doubling the inner-core, outer-core, and outer-corner trip generation rates (to 60, 54 and 18 trips per zone per day, respectively), by halving these rates (to 15, 13.5 and 4.5 trips), and quartering them (to 7.5, 6.75 and 2.25 trips), while holding all distribution patterns constant. *Demand-centralization scenarios* (SC1 and SC2) adjusted trip generation rates to reflect constant CBD trip generation rates with wider differences between inner and outer rates (all rates within 2.5 miles of the center at 30 trips per zone, and linearly falling for those outside the urban core to 3 trips per zone at the outer corner) and linear rates, with narrower differences between inner and outer rates (rates linearly falling between the inner core [24 trips per zone] and outer service area [20 trips per zone]). In these scenarios, the total number of trips was held approximately constant. It should be noted that the relative trip distributions occurring over the course of the day remained the same from scenario to scenario, tracking average NHTS behaviors.

Service-area scenarios expanded or contracted the total service area by 2.5 miles in north-south and east-west directions, resulting in 56% more area for the region in SA1 and 44% less area in SA2, while holding the Base Case's individual zone trip-generation rates constant. In the expanded-area scenario (SA1) trip generation rates were assumed to continue to fall in the newly served areas at the same per-mile (from city center) rate as scenario S0. Two additional *return-trip by SAV* scenarios (SH1 & SH2) increased and decreased the rate at which individual travelers use SAVs to return home by 11% (and -11%), though trip generation rates were adjusted to ensure that the total number of one-way trips remained constant across the city. A *less-congested-peak scenario* (SP1) increased peak-period speeds by 3 mph, while shrinking peak periods to just 0.5 hours during the AM and 1.5 hours during the PM. A more-congested peak scenario (SP2) decreased peak speeds by 3 mph, and increased peak periods to 2 hours in the AM peak and 3.5 hours in the PM. A *SAV demand scenario* (SD1) altered the SAV demand profile by time-of-day for SAVs to match Seattle travel patterns (PSRC 2006). This scenario exhibited much greater travel during the AM and PM peaks than NHTS national averages. *Relocation strategies* examined the effects of not using any vehicle-relocation strategies, using just a single strategy (SR1 through SR4), or using all but one strategy (SRM1 through SRM4) in

order to compare each strategy's relative effectiveness (from reduced wait times) and costs (from induced VMT). All other scenarios used all four relocation strategies. Finally, limits on the total number of SAVs were established in order to quantify how wait times and quality of service degraded as SAV supplies became limited (SL1600 through SL1200). Table 4 shows all strategies and results across several key measures.

Table 4: Alternative Scenarios' Results

Category	Scenario	Description	# SAVs in fleet	Person-trips per SAVs	5-minute wait intervals	Avg. wait time	% induced travel	Cold starts per person-trip
Base case	S0	Base case scenario	1,688	35.9	249	0.30	10.7%	0.054
Trip generation	ST1	Double trips generated	3,245	37.1	226	0.14	7.3%	0.052
	ST2	Half trips generated	859	33.9	*	0.56	12.2%	0.060
	ST3	Quarter trips generated	433	31.7	301	0.92	13.8%	0.068
Centralization	SC1	Greater trip centralization	1,652	36.2	341	0.31	*	*
	SC2	Lesser trip centralization	1,712	*	213	*	*	*
Service area	SA1	Greater service area	2,272	33.7	337	0.33	*	0.059
	SA2	Smaller service area	1,053	40.3	154	0.26	*	0.048
Return-home trips by SAV	SH1	Greater rate of trips returning home	1,674	35.3	206	0.29	*	*
	SH2	Lesser rate of trips returning home	1,676	*	240	*	*	*
Peak congestion	SP1	Lesser peak congestion	1,517	40.0	214	*	*	0.048
	SP2	Greater peak congestion	1,872	32.3	519	0.35	10.4%	0.060
SAV demand	SD1	Greater AM & PM peak SAV use	2,134	28.3	293	0.24	8.9%	0.096
Relocation strategies	SR0	No relocation	1,691	*	2,425	0.69	4.9%	0.067
	SR1	Only R1	1,674	36.5	433	0.46	7.3%	0.061
	SR2	Only R2	1,707	*	519	0.42	6.4%	0.061
	SR3	Only R3	1,677	*	1,750	0.51	7.0%	0.060
	SR4	Only R4	1,689	*	1,644	0.57	4.9%	0.066
	SRM1	All R minus R1	1,690	*	576	0.34	8.7%	*
	SRM2	All R minus R2	1,704	35.5	239	*	9.5%	0.057
	SRM3	All R minus R3	1,688	*	280	0.36	8.6%	0.056
Limitations on # of SAVs	SL1600	SAVs limited to 1600	1,600	37.9	694	0.38	*	0.050
	SL1500	SAVs limited to 1500	1,500	40.5	2,422	0.59	*	0.047
	SL1400	SAVs limited to 1400	1,400	43.5	9,505	1.24	*	0.043
	SL1300	SAVs limited to 1300	1,300	47.0	22,718	2.38	*	0.039
	SL1200	SAVs limited to 1200	1,200	50.9	41,469	3.98	*	0.035

Notes: * signifies that the resulting average outcome was either not statistically significant ($p > 0.05$) vs. the base-case scenario's results, and/or not practically significant (within +/-1% change in SAVs per trip or +/-2% for other measures).

The SAV demand scenario had more substantial impacts on the overall outcomes than any other scenario that did not limit the SAV fleet or available relocation strategies. In Scenario SD1, strong peak-hour demand for SAVs was evident, so the number of trips that a single SAV could serve was nearly 27% lower than in the base case. Moreover, the number of 5-minute wait intervals for travelers placing trip requests (or "5Is") grew by 18%, even with the expanded fleet.

Among the other scenarios that alter base-case parameters (i.e., all non-relocation and SAV-limiting scenarios), the extent of peak congestion/lowered speeds (shown in SP1 & SP2) and overall number of trips served (in ST1-3) impacted outcomes to a greater degree than other scenario changes. Changes in trip-pattern centralization (SC1 & SC2) and return-home rates using SAVs (SH1 & SH2) had relatively minimal impacts. As should be expected, greater peak congestion (SP1) negatively impacted travel outcomes. Even when more vehicles were present (vehicles per person-trip rose 11%, due to the warm-up period run before the first day to keep wait times under 10 minutes), the total number of 5Is rose by 108% under SP1. Conversely, the 5Is fell by 14% in the low-congestion scenario (SP2), even with 10% fewer SAVs per person-trip. Relative to these congestion scenarios, the trip-generation scenarios served impacted average wait times to a much stronger degree, though even average wait times remained under a minute in the quarter-trips scenario (ST3). When trip rates doubled and fleet size rose 92% (ST1), average wait times were 51% lower than the base-case scenario, 5Is fell by 11% on a per-person basis, and the share of induced VMT fell 32%. With SAVs serving only a quarter of the base case (S0's) trips (i.e., under scenario ST3), average wait times rose by 89% and 5Is by 21%. Interestingly, the results suggest that rather than the degree of centralization (SC1 and SC2), the city's overall trip-making density is much more important for achieving beneficial outcomes, with the smaller service area scenario (SA2) showing favorable results, as the lower-demand zones were dropped/no longer served. However, readers should be cautioned that this can also mean fewer trips from the high-demand zones for travelers wishing to go beyond the service fence (SAV-service boundary), though this investigation did not suppress demand to model the effects of more limited destination choices.

The four vehicle-relocation strategies were also tested, alone and in combination, in order to determine each strategy's overall effectiveness. Strategy R1 (large-block relocation) is clearly the most effective, with better outcomes for trips per SAV, 5Is, and induced travel than all three other strategies combined. Part of R1's success may stem from the trip attraction rates, which pull vehicles to the urban core more strongly during the morning and more strongly to the periphery during the afternoon. When used alone, coarser resolution or "big picture" strategies appear much more effective for reducing wait times than do finer-resolution local strategies, though all contribute to lower wait times. Operators looking to reduce empty-vehicle travel (i.e., excess VMT) may wish to consider eliminating strategy R3, which should decrease induced travel by nearly 20%, though this will likely result in average wait time increases and to a lesser extent more 5Is. Though not modeled here, off-peak strategies may also be employed, to reduce relocation during times of low demand. As scenario SR0 shows, without any relocation strategy in place, travelers clearly suffer: 5Is increase almost ten-fold, and average wait times more than

double. However, the base-case scenario (S0) also illustrates the *costs* of having relocation strategies in place, with induced travel/added VMT rising from 4.9% to 10.7% of total VMT.

Finally, the effects of limiting the number of SAVs were tested, in order to determine how quickly level of service degrades (under scenarios SL1200 through SL1600). Figure 6 depicts the share of travelers waiting different time intervals as SAV supply falls (relative to the 1688 base-fleet size). Shares of travelers waiting at each of the specified time intervals increase very gradually, until certain thresholds are hit, at which point they rise quickly at near-linear rates, as SAV count falls.

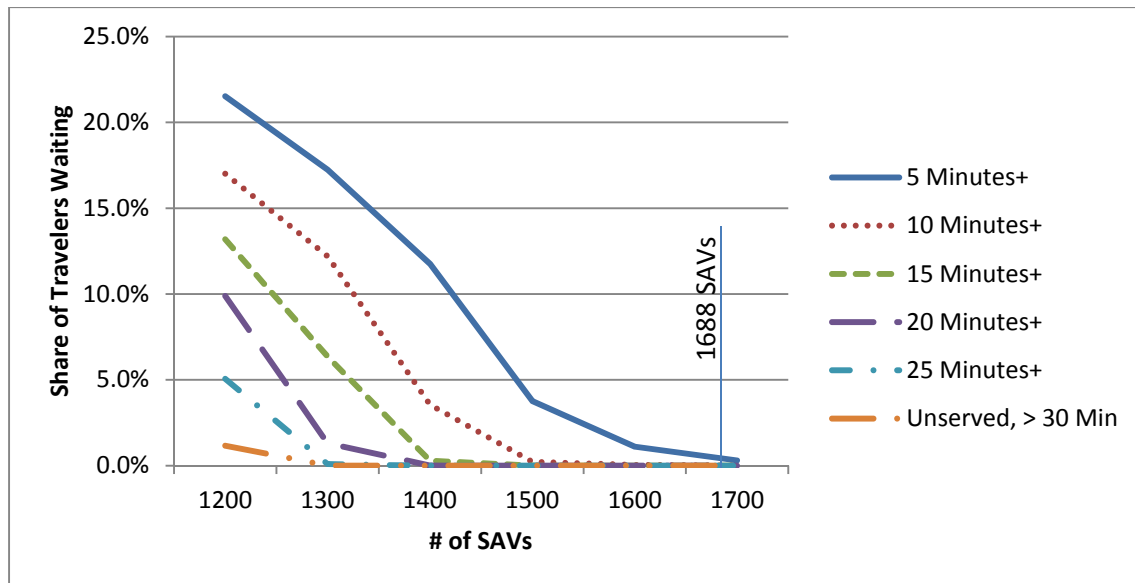


Figure 6: Traveler Shares Waiting for Longer Intervals as SAVs Become Limited (Base = 1688)

While limiting the system’s total number of SAVs has some definite advantages (lower service cost, fewer cold starts, and more trips served per SAV), the traveler delay costs are significant. Nevertheless, a SAV service provider may opt to limit the number of SAVs below the base-case conditions, perhaps charging different rates for quicker versus more delayed vehicle provision to travelers. For example, 1600 SAVs the share of persons waiting more than 5 minutes was just 1.1% and this rose to 3.8% at 1500 SAVs - with 0.2% of travelers waiting 10 minutes or more and around 3 travelers per day waiting 15 minutes plus. Some providers may deem one of these lower service levels acceptable and opt for a smaller fleet size. However, at just 1400 SAVs, service substantially degraded, as the shares of travelers waiting for 5, 10 and 15 minutes or more rose to 11.8%, 3.6%, and 0.3%, respectively. Further vehicle limitations result in even worse wait times. As a counter-argument to reducing SAV fleet sizes, suppliers may wish to set targets for their highest-demand days (e.g., Fridays in August, rather than average days). Even relatively poor service would likely gain many adherents, especially if suppliers lowered rates for poor service (e.g., offering a discount for travelers waiting more than 10 minutes). Related to this, Barth and Todd (2001) conducted a survey among UC Riverside staff, faculty and graduate students regarding shared (but not autonomous) electric-vehicle use. 39% of respondents indicated that they would use likely the service, and over three quarters of those respondents stated that they would wait up to 10 minutes or more. However, the program in question loaned

vehicles free of charge for the first hour of use; paying SAV customers will likely be more demanding.

MODEL LIMITATIONS AND FUTURE WORK

While this model and these results offer a broad and new understanding of the travel benefits and environmental implications of SAVs, there are several opportunities for improvement. For example, rather than a symmetric urban area with omnipresent gridded network, assumed trip generation and attraction rates, and vehicle-relocation heuristics, future modeling efforts should examine actual urban areas, with heterogeneous land use and travel patterns and a realistic (and congestible) network. Mode, destination, and time-of-day choice models may also be incorporated, to give travelers some flexibility, along with demand variations day by day, to reflect weekends, seasonality, and special events. Multiple SAV types may be also modeled, along with differing travel-party sizes. SAV charges could fall with delays experienced, and vehicles could be called and released in continuous time (rather than simulated in 5-minute intervals).

The authors also seek to model driverless dynamic ridesharing (casual carpooling, on-demand and served in real time) as an alternative, similar to some of the efforts undertaken by Kornhauser et al. (2013), though with limited fleet sizes, SAV relocation efforts, and other enhancements. This new paradigm will facilitate additional VMT reductions and environmental benefits. Moreover, routines may be developed to move vehicles out of higher-priced parking areas during times of day with low demand. Such improvements will add realism while enhancing scenario evaluations, enabling more robust analysis of our world's SAV future and related environmental impacts.

CONCLUSIONS

This agent-based investigation of an urban SAV paradigm simulates the movement of travelers and their shared vehicles around a city, throughout the day. The model offers a basic framework for characterizing the environmental and travel implications of an SAV fleet, with estimates of how different contexts and vehicle relocation strategies affect customer wait times, VMT, and cold vehicle starts. Case-study results indicate that a system of SAVs may well save members ten times the number of cars they would need for self-owned personal-vehicle travel, but would incur about 11% more travel (to reach the “next in line” traveler). The overall emissions savings are expected to be sizable, for most species. Quicker vehicle fleet turnover may generate even more benefits, as older and more polluting vehicles are replaced with newer, cleaner ones.

Different trip generation and distribution rates, travel behaviors, degrees of trip-making centralization, service area sizes, and levels of congestion were also modeled here, with results indicating that having many SAV users within a concentrated area, and low congestion levels are key factors for reducing extra/induced VMT and lowering average and extreme wait times (particularly the longer wait times). Four SAV relocation strategies were applied to most scenarios, in tandem and independently, with results suggesting that relocation methods using global (block to block) frameworks perform better than methods with more localized (zone to zone) outlooks. It is also worth noting that the model's initial vehicle-generation process likely

added more SAVs than needed to provide a reasonable level of service, at lower cost to travelers. Less extensive / less vehicle-rich SAV programs may be able to serve the same number of trips as twelve or thirteen conventional vehicles, with any travelers experiencing unusual delays getting their trips discounted or free.

As suggested via these simulations, there are a great many savings to look forward to in SAV settings, if lower-cost travel does not spur excessive vehicle use on space-limited infrastructure. Many barriers remain for implementing such a system in most areas; these include technological barriers (since fully automated vehicles, able to operate in all urban environments safely and without backup drivers, are not yet on the road), regulatory barriers (since vehicle-licensing standards are needed, and commercial barriers to taxi-like services must be overcome), and cost (with affordability a key issue). As such issues are tackled, this model provides a forward-looking perspective on a world of shared mobility, enabled by vehicle automation. It is not yet clear what the future will bring, but the possibilities are enticing.

REFERENCES

Barth, Matthew and Michael Todd (2001) User Behavior Evaluation of an Intelligent Shared Electric Vehicle System. *Transportation Research Record* No. 1760: 145-152.

Bell and Iida (1997) *Transportation Network Analysis*. New York: John Wiley & Sons.

Bureau of Transportation Statistics (2012) Period Sales, Market Shares, and Sales-Weighted Fuel Economies of New Domestic and Imported Automobiles. U.S. Department of Transportation, Washington, D.C.

Burns, Lawrence, William Jordan, and Bonnie Scarborough (2013) Transforming Personal Mobility. The Earth Institute – Columbia University. New York.

Center for Information and Society (2013) Automated Driving: Legislative and Regulatory Action. Stanford University.

Chester, Mikhail and Arpad Horvath (2009) Life-cycle Energy and Emissions Inventories for Motorcycles, Diesel Automobiles, School Buses, Electric Buses, Chicago Rail, and New York City Rail. UC Berkeley Center for Future Urban Transport.

Department for Transport (2005) Making Car Sharing and Car Clubs Work: A Good Practice Guide. London, U.K.

The Economist (2013) How Does a Self-Driving Car Work? April 29.

Federal Highway Administration (2009) National Household Travel Survey. U.S. Department of Transportation. Washington, D.C.

Fisher, Adam (2013) Inside Google's Quest To Popularize Self-Driving Cars. *Popular Science*, September 18.

Ford, Hillary Jeanette (2012) Shared Autonomous Taxis: Implementing an Efficient Alternative to Automotive Dependency. Bachelors Thesis in Science and Engineering. Princeton University.

Kang, Jee E. and W. W. Recker (2009) An Activity-Based Assessment of the Potential Impacts of Plug-in Hybrid Electric Vehicles on Energy and Emissions using 1-Day Travel Data. *Transportation Research Part D*, 14 (8): 541-556.

KPMG and CAR (2012) Self-Driving Cars: The Next Revolution.

Martin, Elliot, S. Shaheen (2011). The Impact of Car-sharing on Public Transit and Non-Motorized Travel: An Exploration of North American Car-sharing Survey Data. *Energies* 4, 2094-2114.

Puget Sound Regional Council (2006) 2006 Household Activity Survey. Seattle, WA.

Santos, A., N. McGuckin, H.Y. Nakamoto, D. Gray, and S. Liss (2011) Summary of Travel Trends: 2009 National Household Travel Survey. Federal Highway Administration Report #FHWA-PL-11-022. Washington, D.C.

Schrank, David, Bill Eisele and Tim Lomax (2012) 2012 Urban Mobility Report. Texas A&M Transportation Institute. <http://mobility.tamu.edu/ums/report/>

Shaheen, Susan and Adam Cohen (2013) Innovative Mobility Carsharing Outlook. Transportation Sustainability Research Center, University of California at Berkeley.

Shoup, Donald (2007) Cruising for Parking. *Access* 30, 16-22.