DYNAMIC RIDE-SHARING AND FLEET SIZING FOR A SYSTEM OF SHARED AUTONOMOUS VEHICLES IN AUSTIN, TEXAS

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

Kara M. Kockelman
(Corresponding author)
E.P. Schoch Professor in Engineering
Department of Civil, Architectural and Environmental Engineering
The University of Texas at Austin
kkockelm@mail.utexas.edu
Phone: 512-471-0210

The following paper is a pre-print, the final publication can be found in *Transportation* 45: 1-16 (2016)

ABSTRACT

Shared autonomous (fully-automated) vehicles (SAVs) represent an emerging transportation mode for driverless and on-demand transport. Early actors include Google and Europe's CityMobil2, who are seeking early pilot deployments in low-speed settings. This work seeks to understand SAVs' potential for U.S. urban areas via multiple applications across the Austin, Texas, network. This work describes advances to existing agent- and network-based SAV simulations by enabling dynamic ride-sharing (DRS, to pool multiple travelers with similar origins, destinations and departure times in the same vehicle), optimizing fleet sizing, and anticipating profitability for operators in settings with no speed limitations on the vehicles and at adoption levels below 10 percent of all personal trip-making in the region.

Results suggest that DRS reduces total service times (wait times plus in-vehicle travel times) and travel costs for SAV users, even after accounting for extra passenger pick-ups, drop-offs and non-direct routings. While the base-case scenario (serving 56,324 person-trips per day, on average) showed that a fleet of SAVs allowing for DRS may result in vehicle-miles traveled that exceed person-trip miles demanded (due to anticipatory relocations of empty vehicles, between trip calls), it is possible to reduce overall VMT as trip-making intensity (SAV membership) rises and/or DRS users become more flexible in their trip timing and routing. Indeed, DRS appears critical to avoiding new congestion problems, since VMT may increase by over 8% without any ridesharing. Finally, these simulation results suggest that a private fleet operator paying $70,000 per new SAV could earn a 19% annual (long-term) return on investment while offering SAV services at $1.00 per mile of a non-shared trip (which is less than a third of Austin's average taxi cab fares).

INTRODUCTION

As vehicle automation continues to advance, one of the more promising opportunities is the concept of shared fully-automated vehicles (SAVs). This concept transforms the notion of travel in most developed countries from one that is largely by privately held personal vehicles to fleet services by driverless, demand-responsive vehicles, shared (or for hire) across a mix of users. Low-speed (25 mi/hr maximum) 12-passenger SAV deployments are underway in Europe, through the CityMobil2 project; and Google recently announced its intention of deploying a fleet of low-speed 2-passenger SAVs (Markoff 2014). While these pilot demonstrations are speed-limited, technological progress suggests they will ultimately travel anywhere a conventional non-automated vehicle can go.

This work builds on Fagnant and Kockelman's (2014, 2015), investigations of SAV operations using an agent-based simulation framework for an idealized city and then across Austin, Texas' coded network. Their latter work uses MATSim-estimated travel times to reflect the dynamic nature of congestion in the region, and mimics the region's highly heterogeneous travel patterns, to anticipate SAV system implications for various shares of travelers who had previously traveled using other modes (mostly private automobile).

The extended model and simulations used here allow for dynamic ride-sharing (DRS), and deliver a benefit-cost analysis for fleet operators, including optimal fleet sizing. DRS allows for on-demand carpooling, for travelers with similar or overlapping paths across both time and space. The new framework allows those willing to share rides to be linked in the same SAV, if their preference requirements are all met. Thus, SAVs can now both pick up multiple travelers at the same node if their destinations are in the same direction, or match travelers at new nodes while the SAV is en-route, as long as single-occupant travel times are not overly compromised.

While DRS has been examined previously as a type of automated taxi (aTaxi) paradigm, several salient features distinguish this work from past efforts. For example, Maciejewski and Nagel (2012) used multiple pick-up and drop-off locations, but their simulation was limited in scale, since they sought to evaluate nearly all service combinations. As a result, simulation times increased by a factor of 100 when moving from 100 customers with 1 depot to 1000 customers with 10 depots.With thousands of nodes and tens of thousands of customers, as needed in city-wide settings and as used here, their approach is not feasible for large-scale applications.

Kornhauser et al. (2013) took a different tack: after obtaining an occupant, each aTaxi simply waits a specified time before departing, to match person-trips with the same origin and nearly the same or directly-en-route destinations. While this approach enjoys operational simplicity, and may reduce vehicle diversion times (to pick up and/or drop off other travelers), much may be gained when serving other travelers along the way (and off the direct routing), particularly at already scheduled drop-off stops.

Jung et al. (2013) developed an innovative DRS scheme, using hybrid simulated annealing (SA), which assigns an initial state of vehicle matches (for example, nearest-vehicle dispatch) and then randomly perturbs vehicle-traveler match decisions to see if the solution can be improved. While this current work may be improved by incorporating the SA method, the approach used here (described below) enjoys certain advantages, predominantly in the area of anticipatory SAV relocation.

Agatz et al. (2011) examined DRS by seeking to minimize total (system-wide) VMT and allowing a substantial 20-minute departure-time window, dramatically improving ride-share matches. In contrast, the DRS methodology described here bins departure times into 5-minute intervals, for relatively inflexible desired departure times (according to the departing traveler's preference). As such, lower wait times take greater priority than system-wide VMT reductions.

THE SIMULATION SETTING

The Capital Area Metropolitan Planning Organization's (CAMPO) regional (6-county) coded roadway network and year-2010 trip tables were used to estimate SAV travel patterns and operational impacts in the Austin area. The network serves 2,258 traffic analysis zones (TAZs), across 5,300 square miles, with centroid nodes located at the center of each TAZ, from which all trips originate and end. Centroid connectors link these zone centroids to the rest of the region's coded network, comprised of 13,594 nodes and 32,272 links (including connectors).

A synthetic population of (one-way) trips was generated using the zone-based personal (non-commercial) trip tables, for four times of day: 6AM – 9AM for the morning peak, 9AM – 3:30PM for mid-day, 3:30PM – 6:30PM for an afternoon peak, and 6:30PM – 6AM for nighttime conditions. CAMPO's regional trip tables were used, and Seattle, Washington's 2006 household travel diaries (PSRC 2006) were for departure time distributions, to map to each of the four times of day. These origin-destination-departure time trip sets (containing 4.5 million trips) were then input into MATsim simulation software (Nagel and Axhausen 2013) to evaluate existing roadway travel conditions across a full (24-hour) weekday. MATSim operates by simulating each trip across the road network, using a dynamic traffic assignment methodology to route individual vehicles from origin to destination. These simulation results were used to estimate average travel speeds across the network, for every hour of the day.

A 100,000-trip subset was then randomly drawn, with 57,161 of these travelers having both origins and destinations with a centrally located 12-mile by 24-mile "geofence". The geofence contains approximately 44% of the region's network links, with a network density of 49.6 links per square mile. This 57,161-trip sample represents just 1.3% of the 6-county region's internal trip-making, and seeks to represent a set of early SAV adopters across a core set of 734 TAZs (32.5% of the 6-region's total). Travelers originate from and journey to the region's TAZ centroids, meaning that each centroid effectively acts as an SAV pick-up and drop-off station. All trips with origins or destinations outside the geofence were assumed to rely on alternative travel modes. Figure 1a shows Austin's regional network and geofence, Figure 1b shows the geofence area in greater detail, and Figure 1c shows the density of those trip origins, at half-mile-cell resolution, within 2-mile (outlined) blocks, and with darker shades denoting higher trip intensities.

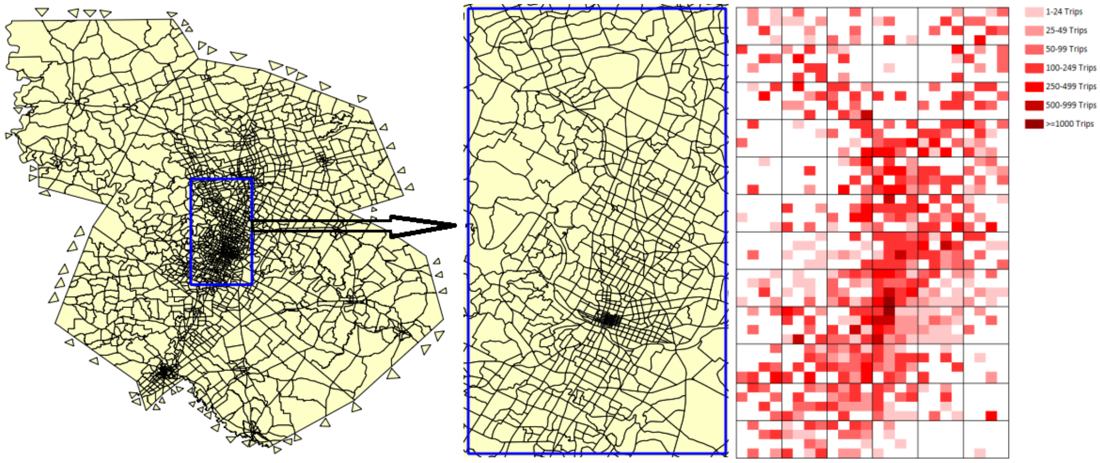

Figure 1: (a) Regional Transportation Network, (b) Nework within the 12 mi x 24 mi Geofence, (c) Distribution of Trip Origins (over 24-hour day, at ½-mile resolution)

MODEL SPECIFICATION AND OPERATIONS

Once the hourly travel times and trip patterns were in hand, an agent-based micro-simulation model was used to build an SAV fleet to ferry those trip-makers from their origins to destinations over the course of a 24-hour day. This model is coded in C++, and uses four primary (non-DRS) modules, including an SAV location and trip assignment module, SAV fleet generation module, SAV movement module, and SAV relocation module. In each of these modules, three sets of actors handle various aspects of the operation: travelers who place requests to a fleet manager and get on and off SAVs, the fleet manager which assigns traveler-SAV pairings and issues relocation commands to SAVs (in anticipation of waiting and future demand), and the individual SAVs that set their route paths and journey throughout the network serving the traveler population.

The first module acts by using the fleet manager to assign waiting travelers to the nearest SAV, with a first-in-first-out (FIFO) scheme to prioritize those who have been waiting longest. Travel demand or trips are grouped into 5-minute bins for vehicle assignment purposes, and each person looks 5-minutes out to see if they could find an available SAV. Travelers who wait 5 or more minutes to access an SAV must expand their search to a 10-minute radius. SAV paths are computed using a backward-modified Dijkstra's algorithm (Bell and Iida 1997) to determine the shortest time-dependent route for an SAV to reach each assigned traveler (and then his/her destination). This process serves as a heuristic for minimizing traveler wait times, with special emphasis on minimizing long waits, while providing an exact solution for minimized in-vehicle travel times.

An SAV "seed" day is run prior to all simulations in order to generate an adequately sized SAV fleet, to ensure that no traveler in the seed simulation will wait more than 10 minutes and still not find an available SAV within a 10-minute radius. At the end of the seed day, this starting fleet size is assumed fixed, and the vehicles' final locations are used for the start of the subsequent day.

The model tracks SAV movements by noting each vehicle's location, future path steps to reach the target destination(s), and distance to the next node for each SAV (if an SAV ends a given 5-minute period between nodes), along with all hour-dependent link-level travel times. During each 5-minute time step, SAVs move across the network, picking up and dropping off travelers (both of which incur a 1-minute time cost, to enable passenger baggage handling, seat belting, and so forth).

SAV relocations (between trip requests) are also often valuable, due to supply-demand imbalances over space and time. For example, SAVs may take more travelers from the geofence periphery to the central business district during the AM peak, resulting in longer wait times for new travelers originating in the outer areas, with excess SAVs lingering in the urban core. Thus, some advance relocation is handy. However, demand-anticipatory relocations can also result in more unoccupied (empty-SAV) VMT, so ideal relocation efforts strike a balance, between lower wait times and lower (empty) VMT.

To achieve this balance, the fleet manager uses a 2-mile by 2-mile block-based comparison of the share of currently waiting travelers plus soon expected travelers (in the next 5 minutes) versus the supply of unoccupied, stationary SAVs in each block. If a given block has 5% of the all free SAVs and 5% of expected demand, it is in perfect balance. If a block's supply exceeds its expected demand or vice versa, by 5+ SAVs, system rules push or pull unoccupied SAVs to or from adjacent blocks, prioritizing shifts to blocks exhibiting complementary imbalances. Additional details regarding these relocations, as well as the SAV user population, Austin network, geofence and model operations can be found in Fagnant and Kockelman (2015).

DYNAMIC RIDE-SHARING

To improve the model's capabilities, DRS opportunities were introduced, allowing two or more independent travelers to share a single SAV, provided that neither traveler is overly inconvenienced. DRS has significant potential for SAVs applications (vs. carpooling with household-owned vehicles). Travelers can rely on a fleet manager to handle the burden of traveler matching, and SAV per-mile cost savings will likely be greater, since the vehicle's capital costs can be incorporated into SAV pricing, but are considered sunk costs if using a household-owned car.

The SAV search process was modified to allow travelers to access SAVs that are currently occupied or claimed by other trip-makers. Potential "handoffs" were also evaluated, to see whether any occupied SAVs could drop off current passengers and then pick up the waiting traveler sooner than other (presently empty) SAVs. These handoffs were not considered true shared rides, which were prioritized if a valid match was found. If the claimed or occupied SAV is the nearest SAV to the new traveler, a series of conditions are checked to determine whether the ride should/will be shared:

1. Current passengers' trip duration increases ≤ 20% (total trip duration with ride-sharing vs. without ride-sharing); *and*
2. Current passengers' remaining trip time increases ≤ 40%; *and*

3. New traveler's total trip time increase grows by ≤ Max(20% total trip without ride-sharing, or 3 minutes); *and*
4. New travelers will be picked up at least within the next 5 minutes; *and*
5. Total planned trip time to serve all passengers ≤ remaining time to serve the current trips + time to serve the new trip + 1 minute drop-off time, if not pooled.

While some of these conditions appear to overlap, each is important in its own right. For example, Condition 1 is the base setting, ensuring that travelers currently in SAVs are not overly burdened with added travel time. In other words, this condition ensures that their decision to share a ride is not excessively costly. Condition 2 prevents travelers who are nearly at their destination from suddenly diverting relatively far out of their way to serve another traveler. Condition 3 takes the new traveler's perspective, to ensure that this particular SAV is worth claiming. Condition 4 deals with the dynamic nature of travel: after 5 minutes many SAVs, if not most, will have moved from their current location and another one may be preferred. Finally, Condition 5 ensures that the trip should be matched from a system perspective. It prevents a short trip from being matched to a longer trip in an opposing direction trip that may satisfy the first four conditions. For example, consider a 40-minute northbound trip paired with a 3-minute southbound trip, both departing from the same node. If the southbound trip is served first, it will add 7 minutes to the northbound trip (including drop-off), would be an unwise ride-sharing decision, but nonetheless be matched without Condition 5.

All combinations of pick-ups and drop-offs for potential trip matches are tested in this way, though not all combinations are considered valid. Same node pick-ups and drop-offs must be concurrent in time, and each traveler must be picked up before he/she can be dropped off. Multiple travelers may simultaneously exit and/or enter an SAV at a given node. If multiple pick-up/drop-off combination orderings are valid for a shared ride, the earliest final drop-off time combination is chosen.

A DAY IN THE LIFE OF AN SAV

To better understand the model operation, an example SAV was tracked throughout an entire 24-hour day, with Figure 2 illustrating its operation in three parts. The first diagram (Figure 2a, upper left) illustrates pick-up and drop-off locations and their ordering, as the SAV travels from one location to the next. Line-weights depict the SAV's occupancy, with the thinnest line-type denoting no occupants, the medium depicting one occupant, and the heaviest holding two persons. Figure 2b (upper right image) shows the actual network links used to travel between locations, and Figure 2c (lower bar chart) depicts the SAV's 24-hour utilization timeline, showing 5-minute periods for when it was moving, picking up, and/or dropping off. Numbers corresponding to visited nodes (i.e., ordered locations) are also shown on the timeline, to better illustrate this SAV's spatial and temporal path over the course of a day.

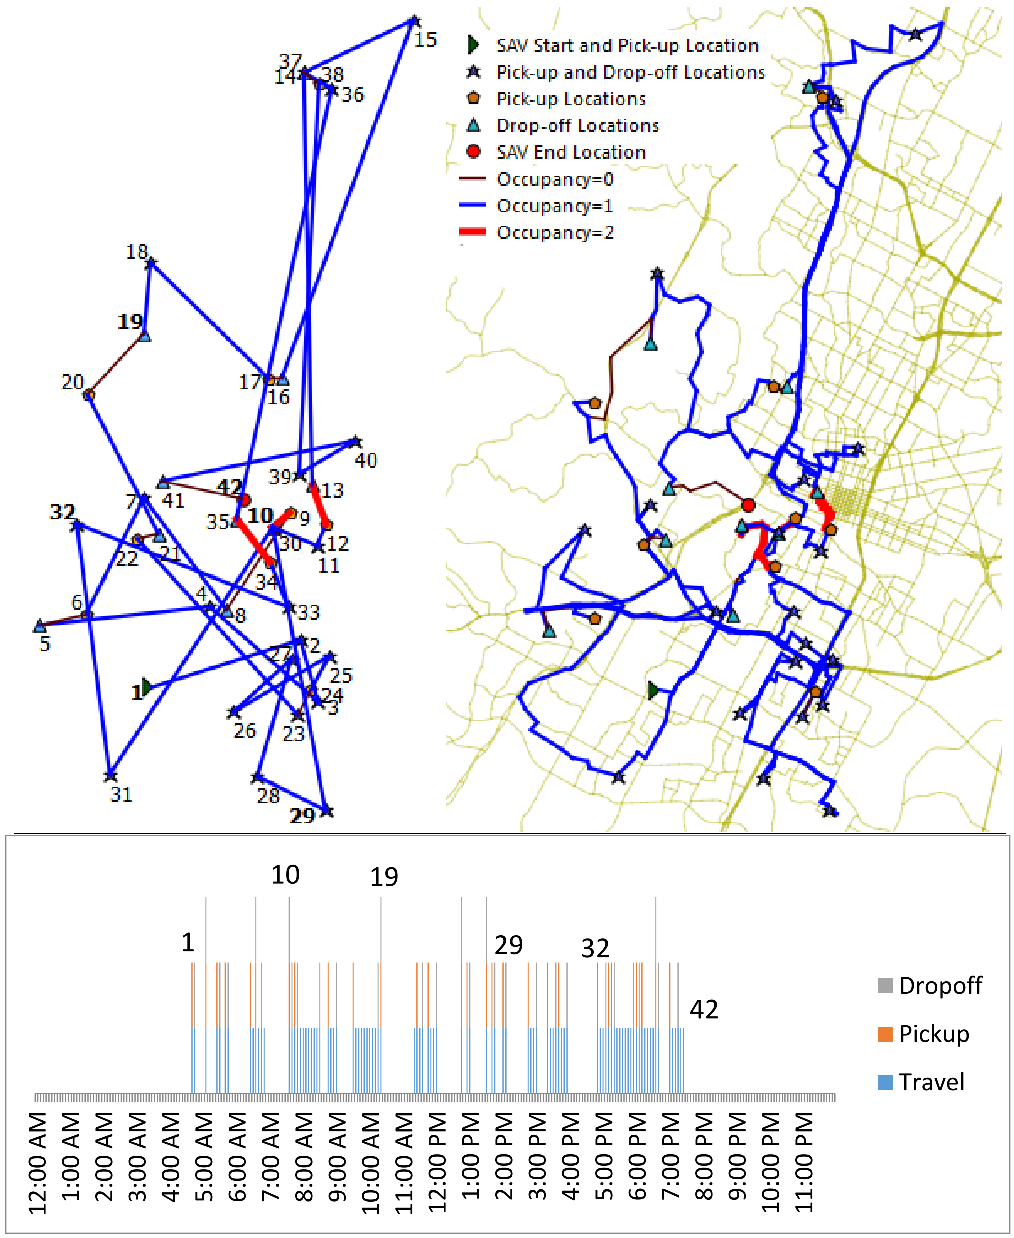

Figure 2: Sample SAV 24-Hour Travel Pattern (a) Node Origin and Destination Ordering, (b) Network Link Utilization and Traveler Origin and Destination Locations, and (c) SAV Travel Timeline

This particular SAV began its operation at 4:40 AM and ended by 7:40 PM. It served 31 person-trips and was "in use" for approximately 8.08 hours of the day[1]. During this time the SAV was either carrying passengers (for about 6.71 hours), relocating itself (about 0.33 hours), or spending one minute picking up and one minute dropping off each traveler it carried (for 1.03 hours total). While there were still a number of trips to be served after this SAV completed its

[1] This SAV was used during 97 of the 24-hour day's 288 5-minute intervals, or for 8 hours and 5 minutes. It was also stationary for a portion of some of these 97 intervals, when travelers were dropped off early in the interval, but the SAV had not yet been assigned to another traveler.

day (around 8% of the daily total), the fleet size (1,715 SAVs to serve 56,324 person-trips) was large enough that this SAV was not needed.

Among the 31 total trips served by this example SAV, trip durations varied from 5 minutes to 50 minutes, and averaged 16 minutes (including pick-up and drop-off times of 1 minute each). Just three trips were shared: two between 7 and 8 AM, and one between 5 and 6 PM, with shared times lasting less than 10 minutes per trip. Two "rebalancing" relocations occurred, including the final trip movement and one just before 7 AM. Finally, of the 31 person-trips, five involved minor unoccupied relocations, to move the vehicle from the SAVs' previous drop-off location to a new pick-up location. It was able to remain in place for the other 26 pickups.

MODEL APPLICATION AND RESULTS

A total fleet size of 1,715 SAVs was generated during the seed day in order to serve the 56,324 person-trips. Assuming an average of 3.02 person-trips per day and 0.99 licensed drivers per conventional vehicle, as shown in the U.S.'s National Household Travel Survey (NHTS) of 2009 (FHWA 2009), each SAV in this (range-limited/geofenced) scenario could potentially replace around 10.77 conventional vehicles, assuming similar demand patterns before SAVs are introduced. Wait times averaged just 1.18 minutes (beyond the average 2.5 minutes associated with the clustering of incoming trip requests to 5-minute intervals), with 98.6% of travelers waiting 10 minutes or less, and average wait times of 4.49 minutes during the peak hour (5PM – 6PM).

While this paradigm appears socially beneficial in terms of replacing many conventional vehicles with a much smaller fleet of SAVs, it comes with some costs in terms of extra (i.e., empty-) VMT, even with DRS enabled. Total added VMT[2] remains positive at 4.5%, with just 6,152 ride-sharing matches out of 56,324 trips occurring on this low-trip-share simulation (and with just 4.83% of total VMT having 2 or more occupants). Almost all shared trips occurred between two persons, with 15,623 VMT (per day) covered by two-person-occupied vehicles, versus 393 VMT covered by 3-person occupancies and 9 VMT occurring via 4-person ride-shares (per day, on average). As SAV fleets capture greater market share (e.g., 10%, 20%, or even 90% of trip-making in the served region/geofence, versus the 1.3% modeled here), presumably much more opportunity will exist for shared rides (thanks to more frequent match-making). Of course, there is also excess driving beyond simple origin-to-destination travel associated with non-shared vehicles. Many drivers incur extra travel searching for parking, and/or park a block or two from their intended destinations (see, e.g., Shoup 2007).

Higher per-mile shared-vehicle marginal costs (as compared to per-mile marginal costs for household-owned vehicles) may also reduce overall VMT. In a privately-owned household-vehicle setting, ownership costs are paid up front. In contrast, ownership costs are embedded in an SAV's rental price, raising marginal per-mile travel costs, and thus potentially reducing demand. On the other hand, the added ease of motorized travel may push overall demand upwards, undercutting transit, high-occupancy (privately-owned) vehicles, and non-motorized

[2] Added VMT reflects extra (unoccupied) travel by SAVs, and reflects travel reductions due to DRS. Total added VMT is calculated by comparing the amount of travel in a given scenario to the amount of travel for the exact same population, if every person were driving a personal vehicle directly from his/her origin to his/her destination.

mode choices. Roadway pricing or other demand-management policies may well be needed, to avoid excessive AV use and worsened roadway congestion.

SCENARIO VARIATIONS

Following the base model's simulation run, a series of alternative scenarios were simulated, testing the implications of various fleet sizes, DRS implementations, and travel demand settings. Three major scenarios types were tested, including a same-sized non-DRS SAV fleet of 1715 vehicles (for direct comparison with the DRS-enabled fleet), allowing a maximum of 30% or 40% total increased travel time for the first and third DRS conditions noted above (up from the base case assumption of 20%), and varying total trip-making demands. Table 1 shows results for fleet size limitations and higher allowable DRS travel time scenarios.

Table 1: Austin Network-Based Model Results across Various Scenarios
(serving 56,324 person-trips over 24 hours)

Measure	With DRS	Without DRS	+ 30% DRS trav. time	+ 40% DRS trav. time
# SAVs	1715	1715	1643	1601
Vehicle replacement rate	10.77	10.77	11.24	11.53
Extra VMT	4.49%	8.68%	2.67%	1.52%
Avg. wait time (min.)	1.18	1.87	1.27	1.37
Avg. PM peak wait (min.)	4.49	8.96	4.82	4.99
Avg. total service (min.)	14.71	14.97	15.20	15.69
% Waiting ≥ 10 min	1.45%	5.65%	1.71%	1.90%
% Waiting ≥ 15 min	0.22%	2.08%	0.27%	0.43%
# Shared trips	6151	0	9233	11,723
% Shared miles	4.83%	0.00%	8.32%	11.20%

A fourth scenario type was also conducted, using mixed shares of DRS-willing and non-DRS-willing travelers, with results suggesting that outcomes (in terms of shared rides, system-wide VMT, wait times, etc.) are roughly quadratic in the share of travelers willing to use DRS. That is, each DRS-willing traveler must be able to find another DRS-willing traveler in order to share a ride, and this becomes increasingly easy as the proportion of DRS-willing travelers grows. However, with substantial market penetration growth, some saturation point may be eventually be reached, potentially resulting in falling DRS matching rates on a per-traveler basis, though the absolute number of shared rides would presumably continue to grow. Additional results regarding these scenarios can be found in Fagnant (2014).

Same-Sized Fleets for DRS and Non-DRS Scenarios

In comparing the DRS vs. non-DRS scenarios, it is apparent that system operation improves when 11% of trips (but less than 5% of VMT) are shared. Fleet-wide added travel (compared to the same number of trips served by privately-held, household vehicles) can be cut by 43%. Wait times also fall (including the share of longer wait periods), though total service time (from pick-up request to final trip drop-off time) increase only slightly, from 14.71 to 14.97 minutes per person-trip. This implies that in-vehicle travel time is likely being substituted for out-of-vehicle wait time at a ratio of approximately 0.6:1 when using DRS.

Higher DRS Travel Time Tolerances

Two other scenarios examined the impacts of adjusting ride-matching parameter settings. The added maximum amount of time that any ride-sharing traveler would have to spend (from initial SAV request, to his/her final drop-off at destination - under DRS conditions 1 and 3) in the base-case scenario was 20%. This parameter was increased to 30% and then 40%, to appreciate its operational effects. Results suggest that changing the maximum from 20% to 30% yielded significant benefits at relatively low cost, in terms of total service times (wait time plus travel time), while the change from 30% to 40% (extra travel time) produced only minor benefits, at much higher cost. For example, the first increase (from 20% to 30%) reduced the amount of extra or empty-SAV VMT by 4.4 miles (per new/added shared-trip) at a cost of 8.9 minutes of added total service time per new shared-trip[3], while also shrinking the SAV fleet size by 72 vehicles, or 4.2 percent. A fleet operator may find this trade-off of lower fleet size and VMT for higher passenger total travel times reasonable, and wish to use a 30% assumption. When increasing the maximum extra travel time ride-sharers are willing to wait by another 10%, to 40% total, VMT was reduced by 2.4 miles at a cost of 11.1 minutes of added service time per new shared-trip, and fleet size fell by just 42 SAVs, indicating that this setting is likely too high to be worthwhile.

Increasing Travel Demand

The final scenario variations tested the impact of scaling the fleet to serve greater demand. Assuming that such services prove successful in one or more cities and regions, demand for SAVs and DRS may grow, along with fleet sizes. As noted above, with just 1.3% of trips served (and 2.3% within the geofence), less than 5% of all SAV VMT resulted in ride-sharing. Increasing trip demands over the same geofenced area may generate economies of density in trip matching, reducing overall VMT and the share of empty VMT.

To these ends, the total base travel demand was grown by factors of roughly 2 and 5, to represent approximately 2.47% and 6.01% of total regional trips, or 4.6% and 11.1% of all geofenced trips. The conventional vehicle replacement rate per SAV was assumed constant, at 10:1, in order to determine travel implications outside of fleet sizing shifts, with scenario outcomes shown in Table 2.

Table 2: SAV Operational Metrics When Serving Larger Trip Shares

% Trips Served within Geofence	**2.3%**	**4.6%**	**11.1%**
# SAVs in fleet	1,846	3,640	9,037
# shared rides per day	5,755	12,933	35,053
% of shared VMT	4.5%	5.3%	5.9%
% extra travel	4.9%	1.8%	-0.2%

[3] New shared-trips are the rise in the number of trips shared over the average simulated day, not whole new person-trips.

Average service time per person-trip (min.)	14.47	14.09	13.93
% travelers waiting ≥ 10 min.	0.77%	0.09%	0.02%

These results are consistent with those shown in Fagnant and Kockelman's (2014) grid-based scenarios. With increased market share, conventional-vehicle replacement should improve, as well as wait times and total service times. Moreover, a higher share of the served population will find ride-sharing matches, resulting in greater VMT reductions (as compared to a non-SAV fleet), even after accounting for unoccupied- (empty-) vehicle relocations. With an even greater market share or more flexible ride-sharing travelers, total fleet VMT may be reduced even further below that evident in today's conventionally-owned vehicle systems. Higher shares were not tested due to computer memory issues, though these may be attempted via code changes in future work.

RECOGNIZING DAY-TO-DAY DEMAND VARIATION

To better appreciate the fleet operator's financial perspective, and the year-long customer's experience, it is important to simulate day-to-day variations in travel demand. To approximate a year's variability, day-to-day variations in personal trips no longer than 50 miles were obtained from the 2009 NHTS (FHWA 2009), over the course of an entire year. The nation's records yielded an average of 1953 person-trips per day, while the state of Texas offered 294 person-trips per day (on average) and the Dallas-Ft. Worth (DFW) metroplex offered 52. These trip records are provided by different persons, every day; so there is great variability in the nature of the trips, that goes beyond inter-regional variations (due to climate and local events, for example) and inter-day variation (from Monday to Friday, and April to November, for example).

The Texas statewide data set was ultimately chosen since it likely represents the closest variation one can expect in sizing central Austin's SAV fleet. As described in Fagnant (2014), based on comparison with Salt Lake City traffic count data (which were available for a series of 365 calendar days), the DFW-only NHTS sample was too small (and thus too variable) to represent the day-to-day variability in *total* demand by tens of thousands of year-long (day-to-day stable) SAV fleet members, even if some regional travel variations across Texas may offset one another. (For example, low demand during a Saturday storm in Houston could partly offset relatively high demand accompanying a football game in Dallas-Ft. Worth on that same day.)

Average increases in household travel from the NHTS data for the top 5% of days in the survey year are 76% in the DFW region alone, 28% looking across the State of Texas, and 14% across the entire U.S., while the average decreases for the bottom 5% of days are -72%, -33% and -23% in those same regions, respectively. In comparing the traffic count variations to those in the NHTS, the within-Texas variations appear reasonable, while DFW's day-to-day variations are too extreme to represent a single region's actual demand variations (Fagnant 2014).

Thus, NHTS travel data from the state of Texas were used to estimate seven distinctive demand days. Accurately assessing this day-to-day variation is crucial in order to ensure that the fleet is properly sized for the entire year, ensuring that services on particularly high-demand days do not collapse as they struggle to keep up with demand. Two of the days are designed to reflect the 18 highest- and 18 lowest-demand days in the year (i.e., the top and bottom 5 percent of days), while the other five days rely on the average VMT within the five inner quintiles of the rest of the year (i.e., the other 90 percent of days). Figure 3 shows how these representative days compare to the cumulative distribution of the 365 days data available in the 2009 NHTS's Texas sample.

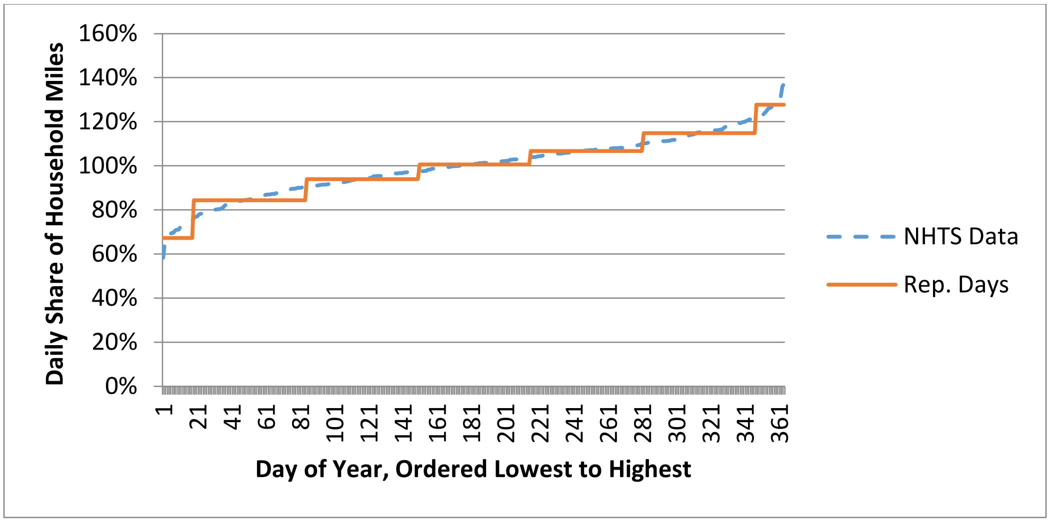

Figure 3: Daily Household Travel in Texas, as a Share of Daily Average

OPTIMAL SAV FLEET SIZING

The above discussions, of fleet operations and travel demand variations, are key to operator costs and system profitability. Fleet sizing can also be varied, with important consequences for costs and customer experience. As shown in Fagnant and Kockelman (2014, 2015), SAV fleet size has direct implications for conventional vehicle replacement rates, as well as system-wide VMT, traveler wait times, and life-cycle environmental impacts. Moreover, operators will wish to size their fleets to maximize profits, while offering users a relatively high level of service (to avoid demand losses and thereby revenue penalties).

With this motivation, a new framework was developed to determine an optimal fleet size. \$70,000 per-SAV purchase costs were assumed (representing \$50,000 costs for AV technology and another \$20,000 for vehicle costs[4], with an additional \$0.50 per-mile operating costs (AAA 2012). Per-SAV capital costs were annualized using the formula:

$$A = \frac{P \cdot i}{1-(1+i)^{-N}} \tag{3-3}$$

where *A* is the annualized SAV capital cost, *P* is the SAV purchase price, *N* is the expected number of service years, and *i* is the discount rate (Newnan and Lavelle 1998). SAVs were assumed to have a 250,000 mile service life, consistent with the expected 7-year service life of Toronto, Canada taxis (which travel over 248,000 miles in the average lifetime [Stevens and

[4] Boesler (2012) notes the U.S.'s top 27 selling vehicles sold for between \$16,000 and \$27,000. SAVs are assumed here to be relatively compact cars or mid-size cars, so a \$20,000 base price assumption was made here.

Marams 2009]), though SAVs may be serviceable longer, thanks to smoother automated driving loads.

Wait times were assessed a penalty, at 70% of the average wage rate (Litman 2013), which is just over $23 per hour for the Austin area, as of May 2013 (BLS 2014). This implies that for every minute the average traveler spends waiting, a 38.4 cent cost is incurred (by the traveler directly, and by the SAV provider indirectly, as assumed here). While these wait penalties do not directly reflect discounted fares that fleet operators may offer to travelers (unless, perhaps, the wait is excessive), wait time is implicitly linked to demand. That is, with lower wait times, more travelers may opt to use SAVs, thus strengthening overall demand; conversely, if wait times are often long, demand may diminish. Therefore, for this analysis, fleet sizing was conducted as if real wait costs are felt by the fleet provider, though they were removed when reporting the final return on investment once the fleet size is determined.

TaxiFareFinder.com estimates Austin taxi travel to cost approximately $2.65 per trip, as a flat or fixed fee, plus another $2.70 per mile, and then a 15% tip on top of those base costs. Assuming an average person- trip distance of 5.64 miles (from the SAV-served trips desired of the population here, internal to the geofence), this works out to an average of $20.56 for a one-way trip, or $3.65 per mile. Since SAVs may replace taxis with a more efficient and cost-effective system, an average $1 per trip-mile fare is assumed here, or $5.64 in operator revenue for the average trip.

A series of simulations were thus run, with varying fleet sizes, using a Golden Section Search optimization procedure (Shao and Chang 2008). This procedure assumes functional concavity (i.e., monotonically increasing until the maximum is reached, and then monotonically decreasing for the remainder of the interval) and works as follows:

1. Boundary conditions for SAV fleet size (x_1, x_2) are first established (here x_1 = 1500 and x_2= 2200 or 2500 SAVs) and evaluated to determine the expected profits ($f(x_1)$, $f(x_2)$) of each.
2. Two points are chosen (x_3, x_4) between these two extreme/boundary values and evaluated ($f(x_3)$, $f(x_4)$). To proceed, at least one of these new $f(x_i)$ values must be greater than both $f(x_1)$ and $f(x_2)$.
3. If $f(x_3) > f(x_4)$, the fleet size corresponding to the maximum profit must lie on the interval between (x_1, x_4), so (x_1, x_4) is established as the new boundary, with known value $f(x_3)$ falling within this interval. Otherwise, if $f(x_4) > f(x_3)$, the new interval will be (x_3, x_2), with value $f(x_4)$ lying inside.
4. A new fleet size value (x_5) between the new boundary conditions is chosen, and evaluated $f(x_5)$; and the process continues until an optimal fleet size is identified within ±5 SAVs.

See Fagnant (2014) for more details on this methodology and application.

Applying this method, an optimal fleet size of 2118 SAVs was estimated, suggesting an 8.7 conventional vehicles per 1 SAV replacement rate, and the average SAV serving 26.6 person-trips per day within this 12 mi x 24 mi section of Austin. A secondary scenario was also tested with (marginal) operating costs halved, to $0.25 per mile (to reflect possible reductions in fuel usage and reduced vehicle wear due to smoother operation). This significantly improved profits (from an IRR of 13.4% to 42.9%), and resulted in a much smaller fleet size, of just 1704 SAVs, equivalent to a 10.8 vehicle replacement rate. Figure 4 shows how total (expected) annual return on investment for an SAV fleet operator varies with fleet size in these two scenarios, before removing traveler wait costs (since the operator likely will not pay these directly).

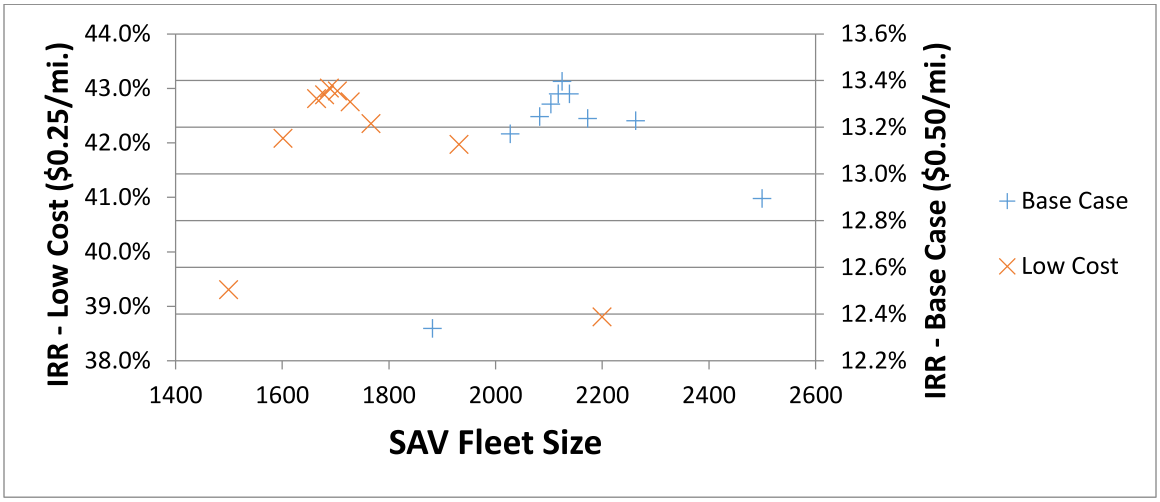

Figure 4: Estimated Annual Internal Rates of Return (Including Wait Costs) across Variable SAV Fleet Sizes

It is also informative to note that total return on investment remained relatively stable in this process, lying between 12.3% and 13.4% in the base case ($0.50/mi.) scenario across almost all fleet sizes[5], and between 38.8% and 43.0% in the low-cost ($0.25/mi.) scenario, even with substantial variations in fleet size (33% and 47%, respectively). Table 3 shows base scenario component costs for the boundary fleet values and the optimal 2118 SAV fleet size, to further illuminate fleet sizing implications.

Table 3: Per-Trip SAV Costs, Revenues and Profits

Fleet Size	Mileage Costs	Capital Costs (at 7%)	Wait Costs	Revenue per Trip	Profit per Trip (w/ wait costs)	Profit per Trip (no wait costs)
1882	$3.001	$1.979	$0.421	$5.640	$0.240	$0.661
2118	$2.995	$2.007	$0.320	$5.640	$0.319	$0.639
2500	$2.988	$2.054	$0.252	$5.640	$0.346	$0.598

These results indicate that all fleet size scenarios result in similar outcomes due to very similar per-trip mileage, high annual mileage (resulting in a high retirement/turnover rate of vehicles), and relatively low wait times. Since mileage cost differences across fleet size values are minimal (decreasing slightly with larger fleet size, due to fewer unoccupied relocations), the main tradeoff becomes capital costs versus wait costs. As the IRR grows larger, the disparity between capital costs in the various scenarios grows; so a smaller fleet is preferred for the low-cost scenario, while a larger fleet is best for the base-case scenario. If wait time costs are removed

[5] Wait costs were excessive with a fleet of just 1500 SAVs, eliminating almost all profit in the base-case scenario.

from the equation to reflect actual costs to be paid by the operator, return on investment for the base-case scenario optimal fleet size rises from 13.4% to 19.4%. As noted earlier, while smaller fleet sizes may increase profits further, they may also result in lower demand levels, so an optimal fleet size of 2118 SAVs is recommended here, for the base-case conditions.

Many factors may change these results, as shown in the lower-operating-costs scenario. Since mileage costs do not change substantially with fleet size, smaller optimal fleet sizes may be achieved by increasing fares, assuming constant demand. As such, neither the 8.7 nor the 10.7 replacement rate should be taken as a fixed optimal value. Rather, operators should understand that an optimal SAV-conventional household vehicle replacement rate in this type of context should be around 10-to-1 (though possibly somewhat lower, since trips to destinations outside the geofence will likely have longer distances, on average), and a methodology like the one used here may be employed to determine specific fleet sizes, given a proper understanding of the underlying context. Other questions also arise that are not directly answered here, like how competitive SAVs may be with household vehicle ownership?
In addition to changing demand and fares, these contexts may vary by potentially limiting SAV speeds, expanding the geofence into low trip intensity areas, or widening the service area in general, which would result in longer average trips. In essence, these results suggest that sizing the SAV fleet for an average day works relatively well for the rest of the year, and sizable returns on investment are quite possible (or lower consumer prices with enough competition), even when accounting for variations between high-demand and low-demand days and higher per-SAV purchase costs.

CONCLUDING REMARKS

Rising degrees of vehicle automation are expected to eventually have profound impacts on our transportation systems, opening the way for a novel transportation mode, the SAV. The results of this work suggest that DRS applications may be critical in limiting excess VMT stemming from unoccupied vehicle relocations, by simultaneously pooling multiple person-trips in the same vehicle. Under base conditions for 1.3% of Austin trip making within a 24 mi x 12 mi geofence, with conservative DRS parameters, excess VMT may be cut from 8.7% to 4.5%; and, as trip-making intensity rises and DRS parameters are loosened, greater ride-sharing and less relocation may actually reduce net VMT. DRS may also greatly reduce wait times, particularly during the heaviest peak hour (from 9.0 to 4.5 minutes, as simulated here). Average total service (wait, plus in-vehicle) time may also be improved via DRS (from 15.0 to 14.7 minutes, as modeled here), even after non-direct routing time costs and time spent picking up or dropping off other passengers is added. This investigation also demonstrates how SAVs could be quite profitable: Assuming SAV purchase prices of $70,000 and travel fares of $1 per trip-mile (less than a third of what Austin taxis charge), and no competition, a fleet operator is simulated to achieve a substantial 19% return on his/her investment.

Ultimately, VMT impacts, conventional-vehicle replacement ratios, operator profits, and many other outcomes depend heavily on implementation details. Market penetration, relocation strategies, DRS assumptions, trip pricing decisions, geofence service areas, and maximum SAV occupancies will probably have important impacts on all these outcomes. This investigation points towards some clear broad outcomes that hold great relevance for future planning and policy-making efforts, regardless of implementation details. An SAV system on the scale envisioned here should lead to lower household vehicle ownership rates, lower parking requirements, traveler cost savings, and significant operator profit opportunities. Additionally, if cities and regions are to avoid some of the excess VMT scenarios that can emerge under SAV (much like taxi) operations, DRS opportunities must be appropriately incentivized.

This work provides a series of case study applications, simulation techniques, and evaluation methods to anticipate and appreciate the potential impacts of AV adoption, SAV applications, and DRS opportunities – and the relative influence of key variables in such systems. The methods used and scenario outcomes discussed provide guideposts for both innovators (who seek to implement a large-scale SAV fleet), as well as transportation planners and policy makers (who must plan for their arrival).

REFERENCES

Agatz, N., A. Erera, M. Savelsbergh, and X. Wang (2011) Dynamic Ride-Sharing: A Simulation Study in Metro Atlanta. *Transportation Research Part B.* 45: 1450-1464.

American Automobile Association (2012). Your Driving Costs: How Much are you Really Paying to Drive? Heathrow, FL.

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

Boesler, Matthew (2012). The 27 Best Selling Vehicles in America. *Business Insider*, August 1.

Bureau of Labor Statistics (2014). May 2013 Metropolitan and Nonmetropolitan Area Occupational Employment and Wage Estimates: Austin-Round Rock-San Marcos, TX. Washington, D.C.

Fagnant, Daniel (2014) The Future of Fully Automated Vehicles: Opportunities for Vehicle- and Ride-Sharing with Cost and Emissions Savings. Doctoral Dissertation in Civil Engineering. University of Texas at Austin.

Fagnant, Daniel and Kara Kockelman (2014) Environmental Implications for Autonomous Shared Vehicles, Using Agent-Based Model Scenarios. *Transportation Research Part C* 40: 1-13.

Fagnant, Daniel and Kara Kockelman (2015) Development and Application of a Network-Based Shared Fully-Automated Vehicle Model in Austin, Texas. Forthcoming in *Transportation Research Record*. Washington, D.C.

Federal Highway Administration (2009). National Household Travel Survey. U.S. Department of Transportation. Washington, D.C. Retrieved at http://nhts.ornl.gov/index.shtml.

Jung, Jaeyoung, R. Jayakrishnan, and Ji Young Park (2012). Design and Modeling of Real-time Shared-Taxi Dispatch Algorithms. Transportation Research Board 92nd Annual Meeting Compendium of Papers. Report 13-1798.

Kornhauser, A., Chang A., Clark C., Gao J., Korac D., Lebowitz B., Swoboda A. (2013). Uncongested Mobility for All: New Jersey's Area-wide aTaxi System. Princeton University. Princeton, New Jersey. Retrieved at: http://orfe.princeton.edu/~alaink/NJ_aTaxiOrf467F12/ORF467F12aTaxiFinalReport_Draft.pdf.

Litman, Todd (2013). Transportation Cost and Benefit Analysis II – Travel Time Costs. Victoria Transport Policy Institute. Retrieved at: http://www.vtpi.org/tca/tca0502.pdf.

Maciejewski, Michal and Kai Nagel (2012). Towards Multi-Agent Simulation of the Dynamic Vehicle Routing Problem in MATSim. PPAM 2011, Part II, LNCS 7204, 551-560.

Markoff, John (2014) Google's Next Phase in Driverless Cars: No Steering Wheel or Brake Pedals. *New York Times*. May 27.

Nagel, Kai and Axhausen, Kay (2013). MATSim: Multi-Agent Transport Simulation. Version 5.0. http://www.matsim.org.

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

Shao, Riming and Liuchen Chang (2008) A New Maximum Power Point Tracking Method for Photovoltaic Arrays Using Golden Section Search Algorithm. *Proceedings of the 2008 Canadian Conference on Electrical and Computer Engineering*. Niagra Falls, ON.

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

Stevens, Matthew and Ben Marans (2009). Toronto Hybrid Taxi Pilot. Toronto Atmospheric Fund. Toronto, Ontario. Retrieved at: http://www.fleetwise.ca/taxi.pdf.

TaxiFareFinder.com (2014). Taxi Fare Finder – Austin.

Utah Department of Transportation (2014). 2013 Hourly Traffic Volume Reports & ATR Maps. Retrieved at: http://www.udot.utah.gov/main/f?p=100:pg:0::::V,T:,2313.