OPERATIONS OF A SHARED AUTONOMOUS VEHICLE FLEET FOR THE AUSTIN, TEXAS MARKET

Daniel J. Fagnant
Assistant Professor
Department of Civil and Environmental Engineering
University of Utah
dan.fagnant@utah.edu
Phone: 801-585-2877

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

Prateek Bansal
The University of Texas at Austin
6.9E Cockrell Jr. Hall
Austin, TX 78712
prateekbansal@utexas.edu

The following paper is a pre-print, the final publication can be found in *Transportation Research Record,* No. 2536: 98-106, 2015.

ABSTRACT

The emergence of automated vehicles holds great promise for the future of transportation. While it commercial sales of fully self-driving vehicles will not commence for several more years, once this is possible, a new transportation mode for personal travel looks set to arrive. This new mode is the shared autonomous (or fully-automated) vehicle (SAV), combining features of short-term on-demand rentals with self-driving capabilities: in essence, a driverless taxi.

This investigation examines SAVs' potential implications at a low level of market penetration (1.3% of regional trips) by simulating a fleet of SAVs serving travelers in Austin, Texas' 12-mile by 24-mile regional core. The simulation uses a sample of trips from the region's planning model to generate demand across traffic analysis zones and a 32,272-link network. Trips call on the vehicles in 5-minute departure time windows, with link-level travel times varying by hour of day based on MATSim's dynamic traffic assignment simulation software.

Results show that each SAV is able to replace around 9 conventional vehicles within the 24 mi x 12 mi area while still maintaining a reasonable level of service (as proxied by user wait times, which average just 1.0 minutes). Additionally, approximately 8 percent more vehicle-miles traveled (VMT) may be generated, due to SAVs journeying unoccupied to the next traveler, or relocating to a more favorable position in anticipation of next-period demand.

INTRODUCTION

Vehicle automation appears poised to revolutionize the way in which we interface with the transportation system. Google expects to introduce a self-driving vehicle by 2017 (O'Brien 2012); and multiple auto manufacturers, including GM (LeBeau 2013), Mercedes Benz (Andersson 2013), Nissan (2013) and Volvo (Carter 2012), aim to sell vehicles with automated driving capabilities by 2020. While current regulations require a driver behind the wheel to take control in case of an emergency even if the vehicle is operating itself, it is likely that this requirement will fall away as further testing and demonstration proceeds apace, vehicle automation technology continues to mature, and the regulatory environment adjusts. Once this occurs, vehicles will be able to drive themselves even without a passenger in the car, opening the door to a new transportation mode, the Shared Autonomous Vehicle (SAV).

SAVs merge the paradigms of short-term car rentals (as used with car-sharing programs like Car2Go and ZipCar) and taxi services (hence, the alternative name of "aTaxis", as coined by Kornhauser et al. [2013]). The difference between the two frameworks is purely one of perception and semantics: are SAVs short-term rentals of vehicles that drive themselves, or are they taxis where the driver is the vehicle itself? The answer is both, and SAVs present a number of potential advantages over both existing non-automated frameworks.

In relation to car-sharing programs, SAVs have the capability to journey unoccupied to a waiting traveler, thus obviating the need for continuing the rental while at their destination, or worrying about whether a shared vehicle will be available when the traveler is ready to departing. Also, SAVs possess advantage over non-automated shared vehicles in that they can preemptively anticipate future demand and relocate in advance to better match vehicle supply and travel demand. While SAVs will cost more to acquire and rent than non-automated shared vehicles, relocation benefits are likely to eventually outweigh marginal technology costs.

When comparing an SAV framework to regular taxis, Burns et al. (2013) estimated that SAVs may be more cost effective on a per-mile basis than taxis operating in Manhattan, cutting average trip costs from $7.80 to $1 due to the automation of costly human labor, though these figures may be somewhat optimistic since their analysis assumed a low (marginal) cost of just $2,500 for self-driving automation capabilities. Even in the case of much higher SAV costs (of $70,000 per vehicle), Fagnant and Kockelman's (2014a) simulations show how SAV costs could cut taxi fares by around a third, while still delivering a 19% return on an the operator's investment. Additionally, SAVs are likely to operate under more a system-optimal and overall-profit-maximizing framework, rather than a taxi-driver-optimal one. That is, taxi drivers presumably seek to maximize their individual profits, even if the entire fleet can act cooperatively, to serve the same or greater demand, with lower wait times, and fewer passenger-less miles-traveled. Transportation network companies (TNCs) like Uber and Lyft occupy a cooperative space closer to SAVs due to their more centrally managed framework, but with a degree of routing, relocation (in anticipation of future demand), operation times, and other decision-making factors left to the driver. In contrast, SAVs may be 100% centrally-controlled and always available, enabling greater opportunities for a higher level of service at lower passenger-less costs.

All these factors indicate that SAV services may dramatically exceed current taxi and shared-vehicle market shares, quickly cutting into private-vehicle travel. While household vehicles should retain many distinct advantages (e.g., locked mobile storage, car seats for children, and freedom to leave a messy vehicle), SAVs will become more and more attractive, as costs fall and service improves with increasing market penetration.

Given the distinct advantages that this emerging mode may hold over taxis, TNCs, and shared vehicles, it is important to understand the possible implications and operation of SAVs, as they may become a potentially significant share of personal travel in urban areas. This investigation does exactly that, by modeling Austin, Texas travel patterns and anticipating SAV implications by serving tens of thousands of travelers each day, who had previously traveled using other modes (mostly private automobile). This investigation is also unique among SAV investigations to date (e.g., Fagnant and Kockelman [2014b], Kornhauser et al. [2013], Burns et al. [2013], and Pavone et al. [2011]), in that the analysis uses an actual transportation network, with link-level travel speeds that vary by time of day, to reflect variable levels of congestion.

THE AUSTIN NETWORK AND TRAVELER POPULATION

The Austin regional network, zone system, and trip tables were obtained from the Capital Area Metropolitan Planning Organization (CAMPO), and are used in CAMPO's regional travel demand modeling efforts. The original, six-county network is structured around 2,258 traffic analysis zones (TAZs) that define geospatial areas within the Austin metropolitan area. A centroid node is located at the geographic center of each TAZ, and all trips departing from or traveling to the TAZ are assumed to originate from or end at this centroid. A set of centroid connectors link these zone centroids to this rest of Austin's regional transportation network, which consists of 13,594 nodes and 32,272 links (including centroids and centroid connectors).

To determine SAV travel demand, a synthetic population of (one-way) trips was generated from the region's zone-based trip tables, using 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 night conditions. Each of these time-of-day periods was used to identify different levels of trip generation and attraction between TAZs. Within each of these four broad periods, detailed trip departure time curves or distributions were estimated based on Seattle, Washington's year-2006 household travel diaries (PSRC 2006). This dataset was used since the Austin household travel survey data set's departure times did not make sense (e.g., the strongest demand during the PM peak was reported at 3PM, and other concerns arose regarding the representative nature of the local survey's departure time distribution), while the Seattle data exhibited a much smoother departure time distribution, with peak travel occurring at approximately 7:30 AM and 5 PM, as should be reasonably expected. Figure 1 shows the assumed departure time distribution for all trips.

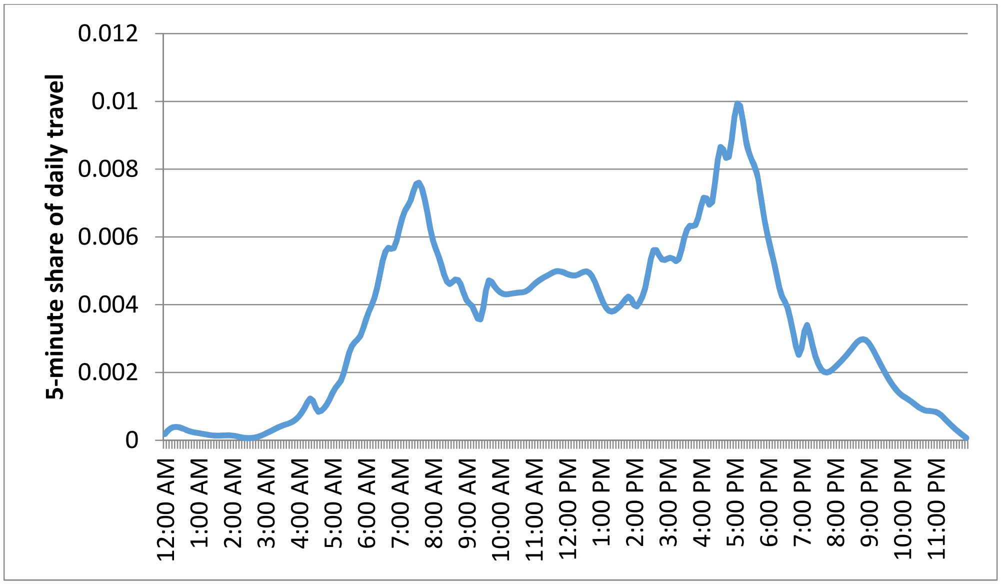

Figure 1: Share of Daily Person-level Departure Times, by Time of Day (Based on 5-Minute Bins) (PRSC 2006)

Once the trip population was generated, a full-weekday (24-hour) simulation of Austin's personal- and commercial-travel activities was conducted using the agent-based dynamic-traffic simulation software MATsim (Nagel and Axhausen, 2013). This evaluation assumed a typical weekday under current Austin conditions, using a base trip total of 4.5 million trips (per day), including commercial-vehicle trips, with 0.5 million of the total trips coming from and/or ending their travel outside the 6-county region. Due to MATSim's computational and memory limitations, 5% of the total 4.5 million trips were drawn at random, with corresponding adjustments to the link-level capacities. As such, each vehicle simulated in MATSim was assumed to represent 20 cars, on average. This is standard MATSim practice, suggested in MATSim's online tutorial (Nagel and Axhausen 2013). While this inevitably results in some loss of model fidelity, the overall congestion patterns that emerge should be relatively consistent with a larger or full simulation (if memory constraints are not an issue), since significant congestion typically occurs at several orders of magnitude beyond the base (20-vehicle) unit used here.

Outputs of the model run were generated, including link-level hourly average travel times for all 24 hours of the day. Next, a 100,000-trip subset of the person-trip population was selected using random draws, and the 57,161 travelers (1.3% of the total internal regional trips, originating from 734 TAZ centroids) falling within a centrally located 12-mile by 24-mile "geofence" were assumed to call on SAVs for their travel. This geofence area was chosen because it represents the area with the highest trip density, and would therefore be most suitable for SAV operation, in terms of both lower traveler wait times and less unoccupied SAV travel (as SAVs journey between one traveler drop-off to the next traveler pick-up). All trips originating from or traveling to destinations outside the geofence were assumed to rely on alternative travel modes (e.g., a rental car, privately owned car, bus, light-rail train, or taxi). Among trips with origins in the geofence area, 84% had destinations also inside the geofence. This indicates that most people residing within the geofence could typically meet most of their trip needs via an SAV system, though perhaps a couple times a week they may require other modes to access areas outside the geofence. Such a system may be better suited for centrally located residents or households giving up one or more vehicles, but retaining at least one. Figure 2a depicts the Austin regional network and modeled geofence location, Figure 2b shows the geofence area in greater detail, and Figure 2c shows the density of trip origins within the geofence, using half-mile-cell resolution within 2-mile (outlined) blocks, with darker areas representing higher trip-making intensities.

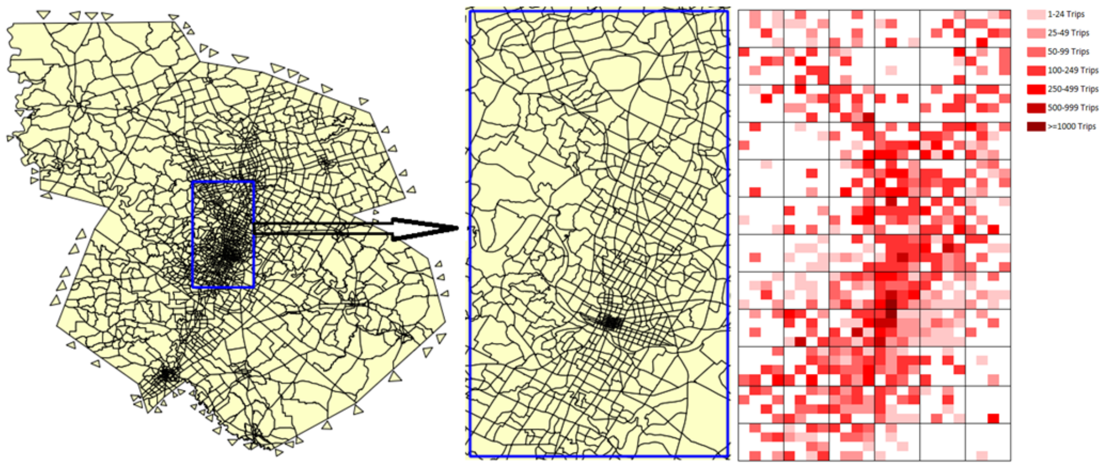

Figure 2: (a) Regional Transportation Network, (b) Nework within the 12 mi x 24 mi Geofence, (c) Distribution of Trip Origins (over 24-hour day, at ½-mile resolution)

MODEL SPECIFICATION AND OPERATIONS

The population of trips within the geofenced area, the transportation network, and hourly link-level travel times were then used to simulate how this subset of trips would be served by SAVs, rather than using personally-owned household vehicles. This simulation was conducted by loading network and trip characteristics into a new C++ coded program, and simulating the SAV fleet's travel operations over a 24-hour day. To accomplish this, four primary program sub-modules were developed, including SAV location and trip assignment, SAV fleet generation, SAV movement, and SAV relocation.

SAV Location and Trip Assignment

The SAV location module operates by determining which available SAVs are closest to waiting travelers (prioritizing those who have been waiting longest), and then assigning available SAVs to those trips. For each new traveler waiting for an SAV, the closest SAV is sought using a backward-modified Dijkstra's algorithm (Bell and Iida 1997). This ensures that the chosen SAV has a shorter travel time to the waiting traveler than any other SAV that is not currently occupied. A base maximum path time is set equal to 5 minutes, and, if an SAVs is located within the desired time constraint, it will be assigned to the trip. Once an SAV has been assigned to a traveler, a path is generated for the SAV, from its current location to the waiting traveler (if the SAV and traveler are on different nodes) and then to the traveler's destination. This is conducted using a time-dependent version of Dijkstra's algorithm, by tracking future arrival times at individual nodes and corresponding link speeds emanating from those nodes at the arrival time.

Persons unable to find an available SAV within a 5-minute travel time are placed on a wait list. These waiting persons expand their maximum SAV search radius to 10 minutes. The program prioritizes those who have been waiting the longest, serving these individuals first before looking for SAVs for travelers who have been waiting a shorter time, or who have just placed a pick-up request. As such, an SAV may be assigned to a traveler who has been waiting 10 minutes and is 8 minutes away from a free SAV over another traveler who has been waiting 5 minutes and is just 2 minutes away from the same vehicle (provided that there are no closer SAVs to the first traveler).

Another feature of the SAV search is a process by which the search area expands. First, travelers look for free SAVs at their immediate node, then a distance of one minute away, then two minutes, and so forth, until the maximum search distance is reached or a free SAV is located. This is conducted to help ensure that vehicles will be assigned to the *closest* traveler, rather than simply to the *first* traveler who looks within a given 5-minute interval.

SAV Fleet Generation

In order to assign an SAV to a trip, an SAV fleet must first exist. The fleet size is determined by running an SAV "seed" simulation run, in which new SAVs are generated when any traveler has waited for 10 minutes and is still unable to locate an available SAV that is 10 minutes away or less. In other words, if nearby vehicle does not free up in the next 5 minutes (when the traveler will conduct another search), the traveler must wait at least 20 minutes. In these instances, a new SAV is generated for the waiting traveler at his/her current location and the SAV remains in the system for the rest of the day. At the end of the seed day, the entire SAV fleet is assumed to be in existence, and no new SAVs are created for the next full day, for which the outcome results are measured and reported. All SAVs begin the following day at the location in which they ended the seed day, reflecting the phenomenon that each individual SAV will not always end up at or near the place where it began its day.

SAV Movement

Once an SAV is assigned to a traveler or given relocation instructions, it begins traveling on the network. During this time the SAV follows the series of previously planned (shortest-path) steps, tracking its position within the network, until 5 minutes of travel have elapsed or the SAV has reached its final destination. Link-level travel speeds vary every hour, thanks to the MATsim simulation results (using 5 percent of the original trip table, on a 5-percent capacity network, to reduce computing burdens in this advanced, dynamic micro-simulation model). SAVs also track the time to the next node on their path, so an SAV's partial progress on a link is saved at the end of the 5-minute time interval, to be continued at the start of the next time interval. If an SAV arrives at a traveler, a pick-up time cost of one minute is incurred before the SAV continues on its path. Similarly, a one-minute time cost is incurred for drop-offs, with SAVs able to both drop off a current passenger and pick up a new, waiting traveler in the same 5-minute interval, if time allows.

SAV Relocation

While the SAV location, assignment, generation, and movement framework described above is sufficient for basic operation of an SAV system, an SAV's ability to relocate in response to waiting travelers and the next (5-minute) period's anticipated demand is important for improving the overall system's level of service. It is important to note that this involves a critical tradeoff: as SAVs pre-emptively move in order to better serve current unserved and future anticipated demand (thus reducing traveler wait times), the total amount of unoccupied (empty-vehicle) VMT grows. That is, more relocation results in lower wait times but also higher VMT. As such, it is advantageous to strike a balance in order to achieve relatively low wait times without overly increasing VMT. Further investigations into these relocation strategies could explicitly state a tradeoff thorough use of an objective function, for example minimizing traveler wait time (or wait time squared, if excessive wait times are deemed particularly important) plus unoccupied VMT, across travelers and SAVs. Those wait times and VMT can be converted to dollars using factors of roughly $23 per hour[1] and $0.50 per mile (AAA 2012), for example.

Using a similar grid-based model, four different SAV relocation strategies were tested in Fagnant and Kockelman (2014b), alone, in combination, and in comparison to a no-relocation strategy. Their results showed how the most effective of the four strategies evaluated the relative imbalance in waiting travelers and expected demand for trip-making across 2-mile by 2-mile blocks, and then pulled SAVs from adjacent blocks if local-block supply was too low in relation to expected demand, or pushed SAVs into neighboring blocks if local supply greatly exceeded expected (next-period) demand. This resulted in dramatic improvements in wait times, with the share of 5-minute wait intervals (incurred with every 5-minute period a traveler waits for an SAV) falling by 82 percent (from 2422 to 433) when using this strategy (versus no relocation strategy in place), even with a slightly smaller SAV fleet serving the same travel demand. Since demand throughout the geofenced Austin area is relatively high and centralized, when aggregated into 2-mile by 2-mile blocks, this relocation heuristic strategy should function well. Readers should be cautioned, however, that this strategy's effectiveness may be limited when two or more high-demand areas are separated by a wide, low-demand area (for example, between two or more cities). In such instances, a more efficient relocation approach would be to shift vehicles within each high-demand area rather independently, and relocate vehicles across the areas only as overall imbalances become more significant.

This same block balancing strategy was implemented in this investigation, using the following steps:

1. Calculate block balances for each 2-mile by 2-mile block, comparing the share of available SAVs in the block against the share of total waiting and expected block demand.
2. Identify the block with the greatest block imbalance above a given threshold (i.e., too many or too few SAVs, relative to expected demand), and adjacent blocks from which to pull or push SAVs.
3. Determine which SAVs to push into adjacent blocks, if the block balance is high, and which SAVs to pull from adjacent blocks, if the block balance is low.
4. Recalculate block balances, based on scheduled relocation actions.

[1] Litman (2013) notes that wait times may be valued at 70% of the wage rate, which is just over $23 per hour for the Austin area, as of May 2013 (BLS 2014). This implies that for every minute each traveler spends waiting, a 38.4 cent cost is incurred.

5. Return to Step 2, until all blocks have either been rebalanced, or have block imbalances below the threshold value.

Step 1 calculates a block balance for each 2-mile by 2-mile block, using Eq. 1:

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

This formula compares the share of SAVs within a given block to share of (expected, next-period) total demand within the same block, normalizing by the total number of SAVs (or fleet size). Therefore, the total block balance represents the excess or deficit number of SAVs within the block in relation to system-wide SAV supply and expected travel demand. Expected travel demand is calculated as waiting trips plus the expected number of new travelers that are likely to request pick-up and departure in the next five-minute interval. The number of new travelers is estimated based by segmenting system-wide trips into one-hour bins, and obtaining average 5-minute trip rates for each block. Any agency or firm operating a fleet of SAVs could probably use historical demand data to inform their fleet's relocation decisions.

Once block balances are assessed, the block with the greatest imbalance is chosen in Step 2 (i.e., the greatest absolute value of Equation 1's result). Those with balance values less than -5 will attempt to pull available SAVs from neighboring blocks, first seeking to pull SAVs (if present) from the surrounding blocks with the highest (positive) balance scores. If a block has a positive balance above +5, it will similarly attempt to push SAVs into neighboring blocks with the lowest balance scores. In both cases, the balance difference between blocks must be greater than 1 in order to justify relocation.

After directions are assigned, the next task (Step 3) is to determine which individual SAVs to push or pull into the neighboring blocks. This is done by conducting path searches to determine which SAVs are closest to the node that is located nearest to the center of the block that the SAV will be moving into. If a pushed SAV is closest to the central nodes in two or more blocks (for example, 5.5 minutes to the block immediately north and 7.4 minutes to the block immediately west), it will be assigned to travel in the direction with the shortest path. These SAV paths are created from their current locations to the central node in the destination block. Each path is then trimmed after 5 minutes of relocation travel, such that the SAV can reassess its position and potentially be assigned to pick up an actual traveler at the start of the next 5-minute interval. If it has entered the new block and has traveled at least 2 minutes while in the new block in the direction of the central node, it will be held at that position for a coming assignment; this halt on relocation towards the new block's central node helps ensure that too many pushed SAVs do not all end up at the central node.

At this point, the block balances are updated (Step 4) and block balancing actions are complete for the given block. Step 5 concludes the algorithm by choosing the block with the next greatest imbalance, and continuing this process until all blocks have either been rebalanced during the current time interval, or their (absolute) block balance scores are no greater than the threshold limit, which is set to 5 in this investigation. Figure 3 depicts an example of the block balancing relocation process, showing balances before relocation assignment, SAV assignment directions by block, and balances after relocation. Integer values are shown here for readability, though actual balance figures are typically fractional.

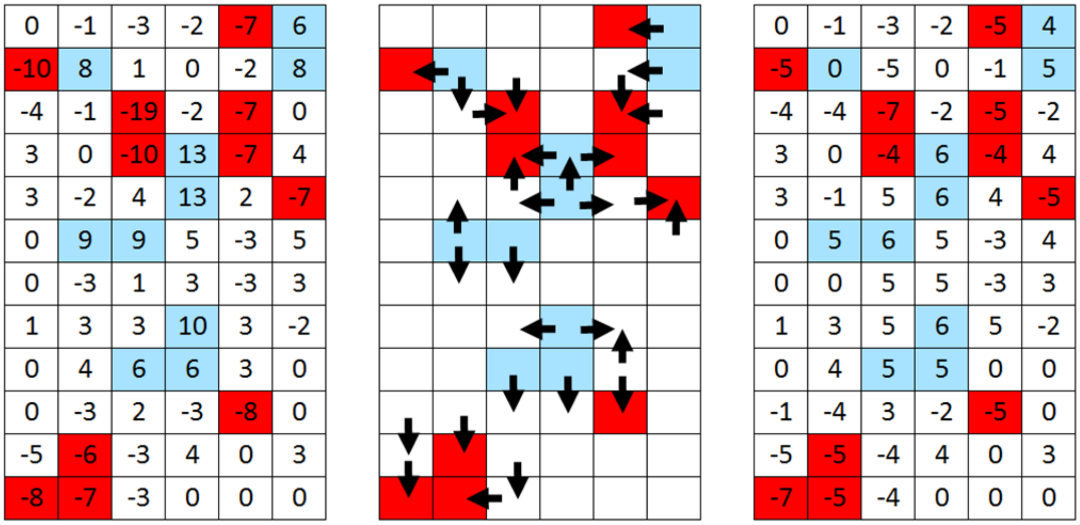

Figure 3: Example SAV Relocations to Improve Balance in 2-mile Square Blocks (a) Initial Expected Imbalances, (b) Directional SAV Block Shifts, and (c) Resulting Imbalances

The other three relocation strategies noted in Fagnant and Kockelman (2014b) are not used here. These include a similar block-balance strategy, using 1-mile square blocks, relocation of extra SAVs to quarter-mile grid cells with zero SAVs in them and surrounding them (and thus half-mile travel distance away), and a stockpile-shifting strategy that relocates SAVs a quarter mile (1 grid cell away) if too many SAVs are present at a given location relative to the immediately surrounding cells (i.e., local imbalances of 3 of more in available SAVs). While these other strategies were somewhat helpful in reducing delays, their overall impact was less than that of the 2-mile-block rebalancing strategy, even when all three were combined. Moreover, the latter two strategies (involving very local or myopic shifts) may not be as effective in the more realistic network setting modeled here, since not every cell is a potential trip generator here, and differences in nearby trip-generation rates can vary dramatically across adjacent Austin cells. In this Austin setting, only one of the 72 two-mile by two-mile blocks had no simulated SAV demand, and 43.7 percent of the half-mile by half-mile cells had demand (with demand originating from an average of 1.46 centroids per non-zero-demand cell). Among the 503 half-mile cells exhibiting some demand, their cumulative trip generation may exceed demand in adjacent cells by a factor of 10 (e.g., 50 trips might be expected in one cell within a 5-minute time period and just 5 trips in the adjacent cell).

MODEL APPLICATION AND RESULTS

From the 4.5 million trips in the Austin regional (6-county) trip table, an initial subset of 100,000 trips was randomly selected, to represent a small share of Austin's total regional trips to be served by SAVs. Among these 100,000 person-trips, 56 percent had both their origins and destinations within the 12 mile x 24 mile geofence modeled here. Their departure times were designed to mimic a natural 24-hour cycle of trips, as described earlier and as shown in Figure 1, with the spatial pattern of trip origins shown (earlier) in Figure 2c. This single ("seed") day was then simulated to first generate a fleet of SAVs, to ensure all (seed-day) wait times lie below 10 minutes. Then, a different day was simulated using the same starting trip population (of 4.5 million trips, from which 100,000 are drawn) to examine the travel implications of this pre-determined SAV fleet size, in terms of vehicle occupancies, unoccupied travel, wait times, and other metrics. While just a single day of travel (in addition to the seed day) was conducted in this simulation noted here, Fagnant and Kockleman's (2014a) results[2] indicate that these outcomes/results should be relatively stable after accounting for day-to-day variations in demand, over an entire year.

All SAVs begin the following day at the location in which they ended the seed day, reflecting the phenomenon that each individual SAV will not always end up at or near the place where it began at the start of the day. These results show how approximately 1,977 SAVs are needed to serve the sample of trips. This means that each SAV serves an average of 28.5 person-trips on the single simulated day. Assuming an average of 3.02 person-trips per day per licensed driver (i.e., someone who could elect to drive his/her own vehicle) and 0.99 licensed drivers per conventional vehicle, an SAV in this scenario could reasonably be expected to replace around 9.34 conventional vehicles, if travel demands remain very similar to demand patterns before SAVs are introduced and assuming one can ignore all travel to (and from) locations outside the geofence.

This figure is biased-high, since it assumes all substituted trips are personal-vehicle trips. While taxi or TNC trips can constitute a share of these replacements, their share is likely be small in this scenario[3]. Also, trips made by persons living inside the geofence (who are more likely to give up a private vehicle) to destinations outside of it ("external trips") will need to be served by other modes, with trip distances often longer than trips within the geofence. Conversely, the first household vehicles to be shed will likely be those that are under-utilized, with other households forgoing purchases of a vehicle that will only be marginally used. For example, Martin and Shaheen (2011) estimate that current effects on vehicle ownership are 9 to 13 vehicles replaced for every non-automated shared vehicle. As such, a likely scenario is a multi-vehicle household shedding one or more vehicles, but retaining at least one to ensure ease of external travel. Therefore, one might expect the first SAVs to replace many household vehicles at first, with falling household vehicle replacement rates as market penetration grows. To fully understand the vehicle replacement implications, mode choice and vehicle ownership models are needed, as well as a greater examination of travel outside the geofence.

This SAV fleet size offers an excellent level of service: Average wait times throughout the day are modeled at 1.00 minutes, with 94.3% of travelers waiting less than 5 minutes, 98.8% of travelers waiting under 10 minutes, and just 0.10% of travelers waiting 15-29 minutes. The longest average wait times occurred during the 5PM – 6PM hour, when demand was highest and speeds slowest/congestion worst, with average wait times of 3.85 minutes. These numeric results assume that all travelers request their trips exactly on 5-minute intervals, since that is when

[2] Fagnant and Kockelman (2014a) simulated SAV operation for 7 representative travel days, spanning from the bottom 5th percentile of personal VMT in Texas' NHTS 2009 data to the top 95th percentile, with results suggesting that average operational results across all days were similar to those found on the median travel day.

[3] With approximately 2.3% of travel within the geofence operating by SAV and 0.12% of current household travel using taxis (NHTS 2009), around 5% of SAV travel may come from former taxi occupants.

vehicle assignment decisions are made; in reality, many will call between 5-minute time points, adding (on average) another 2.5 minutes to the expected wait times (following an SAV trip request). Of course, some travelers will elect to call many minutes or hours in advance of needing an SAV, though these results suggest that such reservations may not be too helpful, except perhaps in lower-density and/or harder-to-reach locations. Moreover, advance vehicle assignments could make the system operate worse, especially if the person who placed the call is not ready and the SAV could be serving another traveler, particularly during high-demand periods of the day.

Other system simulation results showed that 24-hour travel-*distance*-weighted speeds averaged 43.6 mph. However, when taking a time-weighted system perspective, using total travel distance divided by total travel miles (VMT/VHT), average system speeds are 26.1 mph. This reflects the phenomenon that, if an SAV travels 5 miles at 5 mph and 5 miles at 50 mph, it will take 1.1 hours to travel the 10 miles resulting in an effective system speed of 9.1 mph, rather than a travel-distance weighted speed of 27.5 mph. Moreover, 19.4% of total SAV VMT was at speeds of 20 mph or less, likely on local roads and/or during congested times, while 41.4% of total SAV VMT occurred at speeds over 50 mph, typically during off-peak times and on freeways.

A comparison with New York City's taxi fleet casts this Austin-based SAV system in a very favorable light. The NYC's Taxi and Limousine Commission's (2014) Factbook notes that the city's 13,437 yellow taxis serve an average of 36 trips per day, somewhat more than the 28 trips served by SAVs here. However, these simulations indicate that as total demand goes up, more trips can be served per SAV. 90.3 percent of trips that the NYC taxi fleet serves are on the island of Manhattan, a 22.7 square-mile land area (though the entire city is 469 square miles), in contrast to the 288 square miles served here. While the modeled Austin-traveler trips averaged 5.2 miles, yellow taxi trips in NYC average just 2.6 miles, so each yellow taxi travels, on average, 70,000 miles annually, with a stunning 51.5% unoccupied share of VMT (versus the 8.0 percentage simulated here). While NYC taxi demands and service are distinctive (e.g., an extensive subway system can serve many longer trips), such comparisons draw attention to the dramatic service improvements that SAVs may bring communities.

Electric Vehicle Use Implications

One intriguing question to ask is whether SAV fleets could be served by electric vehicles. Electric SAVs may provide a number of advantages over gasoline-powered SAVs, including, for example, fewer emissions for communities, greater energy security for a nation, and perhaps even cost advantages -- if the price of electric vehicle batteries continues to fall. Some AV technology providers see this as a promising future, with Induct demonstrating a fully driverless and electric low-speed passenger transport shuttle in January 2014 in Las Vegas, Nevada, at the Consumer Electronics Show (Induct 2014).

Simulations are valuable for assessing the potential charging implications of an electric SAV fleet, as recently investigated (for cost comparisons, but not battery-charging implications) by Burns et al. (2013). Here, occupied plus unoccupied vehicle distances per vehicle-trip average 6.09 miles, and the SAV fleet was traveling, picking up, dropping off, or otherwise active for 7.14 hours of the day, with SAVs averaging 2.91 stationary/non-moving intervals of at least one hour (when no travelers were being served and no relocations were being pursued) each day, and another 0.80 intervals between 30 minutes and 59.9 minutes (of stationary/sitting time) each day. Such long wait intervals could be productively used for vehicle battery charging, if desired by fleet operators, and if charging stations are reasonably close by. However, daily travel distances averaged 174 miles per SAV, with mileage distributions shown in Figure 4. These distances are much longer than the range of most battery-electric (non-hybrid, electric-power-only) vehicles (BEVs).

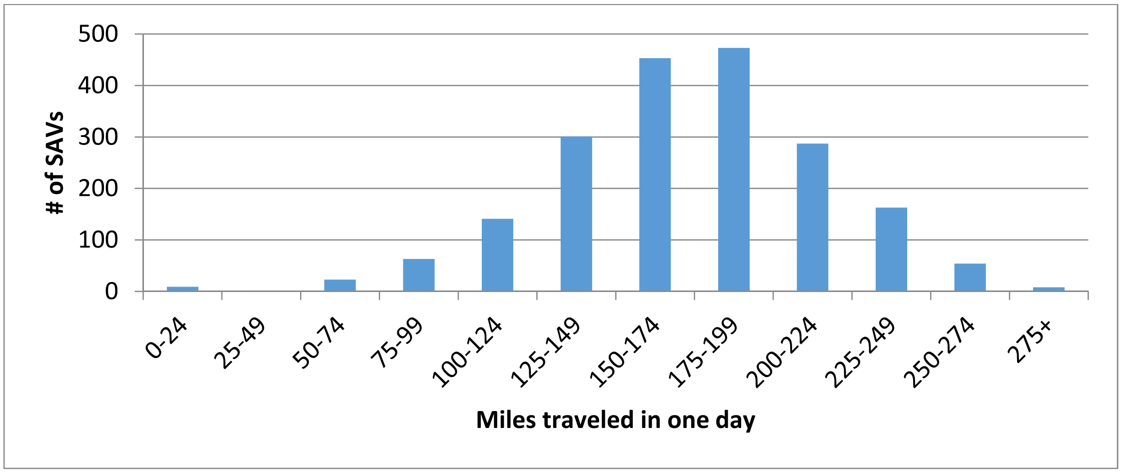

Figure 4: Daily Travel Distance per SAV in Austin Network-Based Setting

Most currently available BEVs for sale in the U.S. have all-electric ranges between 60 and 100 miles (e.g., the Chevrolet Spark, Ford Focus, Honda Fit, Mitsubishi i-MiEV, and Nissan Leaf). For these, the U.S. EPA (2014) estimates typical charge times (to fully restore a depleted battery) to vary between 4 and 7 hours on Level 2 (240 volt) charging devices. This could pose a serious issue for all-electric BEVs in an SAV fleet, but not much of an issue for the Tesla Model S (which enjoys a 208- to 265-mile range and a charge time of under 5 hours when using a Level 2 dual charger [EPA 2014]) or plug-in hybrid EVs (PHEVs), like the Chevrolet Volt, Honda Accord Plug-in, Ford C-MAX Energi, Ford Fusion Energi, and Prius Plug-in Hybrid. Furthermore, fast-charging Level 3 (480-volt) systems can charge large batteries in under an hour, so SAVs that need more frequent daytime charging may need to rely on these devices. Of course, some time is required to develop the automation technology and legal frameworks needed to successfully deploy SAVs. In the meantime, battery charging times, BEV ranges and costs will improve, along with deployment of fast-charging facilities and remote inductive charging devices (allowing SAVs to self-charge wirelessly [MacKenzie 2013]).

SAV Emissions Implications and Grid-Based Comparisons

SAV emissions implications were also evaluated, using that the same method described by Fagnant and Kockelman (2014b). This method applies life-cycle energy usage and emissions rates associated with vehicle manufacture, per-mile running operations, cold-vehicle starts, and parking infrastructure provision, all using rates estimated by Chester and Horvath (2009). The current U.S. light-duty vehicle fleet distribution (BTS 2012) was used, split between passenger cars (sedans), SUVs, pick-up trucks and vans, for comparison with an SAV fleet consisting entirely of passenger cars. It is possible that SAVs will include other vehicle types, but many may be built as smaller cars, perhaps even two-seaters like those Car2Go is currently using for in its shared vehicle fleet, and as Google plans for its SAV fleet (Markoff 2014). Thus, fleet purchase decisions could result in even more favorable (or lower) emissions and energy savings than estimated here, though smaller vehicles potentially limit ride-sharing (to fewer persons) and cargo-carrying opportunities.

Table 1 shows anticipated emissions outcomes, as well as estimates generated by Fagnant and Kockelman (2014b) using a grid-based SAV model for an idealized city and network. This comparison contrasts results between those shown here (in a realistic 12-mile by 24-mile travel-demand setting) with Fagnant and Kockelman's (2014b) grid-based evaluation results (in an idealized 10-mile by 10-mile setting).

Table 1: Anticipated SAV Life-Cycle Emissions Outcomes Using the Austin Network-Based Scenario (Per SAV Introduced)

Environmental Impact	**US Vehicle Fleet vs. SAV Comparison (over SAV lifetime)**					
	US Vehicle Fleet Avg.	**% Pass. Car Running Emissions**	**% Pass. Car Starting Emissions**	**SAVs**	**% Change**	**Grid-Based Estimates**
Energy use (GJ)	1230	88.6%	0.0%	1064	-14%	-12%
GHG (metric tons)	90.1	87.7%	0.0%	83.2	-7.6%	-5.6%
SO_2 (kg)	30.6	14.2%	0.0%	24.6	-20%	-19%
CO (kg)	3,833	58.1%	38.7%	2590	-32%	-34%
NOx (kg)	243	73.3%	14.7%	198	-18%	-18%
VOC (kg)	180	39.0%	43.7%	95.2	-47%	-49%
PM_{10} (kg)	30.2	65.8%	6.6%	27.9	-7.6%	-6.5%

Emissions and environmental outcomes using SAVs are clearly preferable to the current U.S. vehicle fleet. These anticipated environmental outcomes are quite similar to the grid-based results, thanks to similar vehicle replacement rates, trip service levels, and cold-start trip shares. Emissions outcomes disfavored the network-based scenario for species that had high shares of life-cycle emissions stemming from cold-starting emissions (since the network-based scenario resulted in 85% vs. 92% reductions in cold-starts) while the network-based scenario was favored for species where the life-cycle share of running emissions were high (since the network-based scenario resulted in 8.0% vs. 10.7% increases in VMT). Thus, while outcomes in both scenarios were quite similar, the network-based scenario performed slightly better for energy use, GHG, SO_2, and PM_{10}, but slightly less well for CO and VOC.

Other differences between the network-based and grid-based evaluations are similarly illuminating. The latter, pure-grid scenario, with quarter-mile cells and smooth (idealized) demand profiles, out-performs the much more realistic, actual-network-based Austin scenario, for conventional-vehicle replacement and wait times, but with more unoccupied travel. This grid-based evaluation suggested that each SAV could replace two to three more conventional vehicles than this more realistic setting (i.e., it yielded a replacement rate of 11.76 to 1 rather than 9.34 to 1), while cutting average wait times nearly 70% (from 1.00 to 0.30 minutes), with 32% more unoccupied (empty-SAV) VMT (10.7% added VMT in the gridded case vs. 8.0% in the Austin-network setting). The differences in these two settings' results come from a host of very different supporting assumptions. However, neither permits all trips to be taken: both have geofences that cut off trips with destinations beyond fence boundaries.

First, the travel demand profile differed significantly between the two evaluations. The grid-based evaluation assumed a smaller service area and higher trip density, with 60,551 trips per day across a 100 square-mile area, versus 56,324 trips per day across a 288 square-mile area. Average trip-end intensities also varied quite smoothly across quarter-mile cells in the grid-based application (with near-linear changes in travel demand rates between the city center and outer zones), whereas the Austin setting exhibits much greater spatial variation in trip-making intensities (as evident in Figure 2c). The simulated, grid-based setting also added more fleet vehicles based on initial simulations, to keep wait times lower than would probably be optimal for real fleet managers; this Austin fleet sizing is less generous, and presumably more realistic, but traveler wait times remain reasonably low.

Another key distinction between the grid-based and Austin network evaluations emerges in average speeds and average trip distances. Here, travel-weighted 24-hour running speeds average 26.1 mph, whereas constant speeds of 21 mph and 33 mph were assumed in the simulated context, and the 21 mph speed only applied during a 1-hour AM peak and 2.5-hour PM peak period (with 33 mph SAV travel speeds at all other times). Trip distances were constrained to 15 miles in length in the prior application, while this application permits a much wider range of travel behaviors. Finally, this setting allows for a real network – sometimes dense, but often sparse, adding circuity to travel routes; in contrast, the simulated setting assumed a tightly spaced (quarter-mile) grid of north-south and east-west streets throughout the region. Circuity in accessing travelers and then their destinations is harder to serve, especially at lower average speeds, across a wider range of trip-making intensities.

It is interesting how well the Austin fleet still serves its travelers, given the series of disadvantages that exist in this more realistic simulation. Lower trip densities mean that SAVs must travel farther on average to pick up travelers, and slower speeds mean that SAVs will be occupied for a longer duration during the journey, tying them up and preventing them from serving other travelers, and potentially hampering relocation efficiency. Also, while shorter trips lessen travel times, it also means that relocation and unoccupied travel will comprise a greater share of the total. All of these factors suggest that a larger fleet will be needed to achieve an equivalent level of service. But the vehicle-replacement rates remain very strong, at 9.3 conventional vehicles per SAV[4].

CONCLUSIONS

These Austin-based simulation results suggest that a fleet of SAVs could serve many if not all intra-urban trips with replacement rates of around 1 SAV per 9.3 conventional vehicles, assuming other modes are available for travel outside the geofence (e.g., non-shed household

[4] The replacement rate estimated here is 9.3, when accounting for a pure-trip substitution, but should likely be lower, since trips with destinations outside the geofence are unreachable with SAVs under this proposed framework, these trips should likely be longer on average than trips with internal geofence destinations, and mode shifts may also stem from other sources than private-vehicle travel.

vehicles), and a direct 1:1 substitution of household vehicle trips for SAV trips within the geofence. However, in the process SAVs may generate around 8.0% new unoccupied/empty-vehicle travel that would not exist if travelers were driving their own vehicles. Prior, results by Fagnant and Kockelman (2014b) indicated that, as demand intensity (over space) for SAV travel increases, the number of conventional vehicles that each SAV can replace grows, wait times fall, and unoccupied/empty-vehicle travel distances fall. All this points to a higher cost per SAV in the early stages of deployment (in terms of new VMT), though such costs should fall in the long term, as larger SAV fleet sizes lead to greater efficiency.

Moreover, these results have substantial implications for parking and emissions. For example, if an SAV fleet is sized to replace 9 conventional vehicles for every SAV, total parking demand will fall by around 8 vehicle spaces per SAV (or possibly more, since the vehicles are largely in use during the daytime). These spaces would free up parking supply for privately held vehicles or other land uses. In this way, the land and costs of parking provision could shift to better uses, like parks and retail establishments, offices, wider sidewalks, bus parking, and bike lanes.

With regards to vehicle emissions and air quality, many benefits may exist, even in the face of 8.0 percent higher VMTs, as was demonstrated here. For example, SAVs may be purpose-built as a fleet of passenger cars, replacing many current, heavier vehicles with higher emissions rates (like pickup trucks, SUVs and passenger vans). SAVs will also be traveling much more frequently throughout the day than conventional vehicles (averaging 26 trips per day rather than 3, and in use 8 hours each day, rather than 1 hour), so they will have many fewer cold starts than the vehicles they are replacing. Cold-start emissions are much higher than after a vehicle’s catalytic converter has warmed up, and these results suggest 85% fewer cold starts (defined as rest periods greater than 1 hour), when replacing conventional, privately held vehicles with SAVs.

Finally, SAVs hold great promise for harnessing vehicle automation technology, offering higher utilization rates and faster fleet turnover. By using SAVs intensely (estimated here to be 174 miles per SAV per day or 63,335 miles per year), they will presumably wear out and need replacement every three to five years. Since vehicle automation technology is evolving rapidly, this cycling will allow fleet operators to consistently provide SAVs with the latest sensors, actuation controls, and other automation hardware, which tend to be much more difficult to provide than simple SAV system firmware and software updates.

In summary, while the future remains uncertain, these results indicate that SAVs may become a very attractive option for personal travel. Each SAV has the potential to replace many conventional vehicles, freeing up parking and leading to more efficient household personal vehicle ownership choices. Though extra VMT through unoccupied travel is a potential downside, vehicle fleet changes, a reduction in cold-starts, and dynamic ride sharing may be able to counteract these negative impacts and lead to net beneficial environmental outcomes.

REFERENCES

American Automobile Association (2012). Your Driving Costs: How Much are you Really Paying to Drive? Heathrow, FL. http://newsroom.aaa.com/wp-content/uploads/2012/04/YourDrivingCosts2012.PDF

Andersson, Leif Hans Daniel (2013). Autonomous Vehicles from Mercedes-Benz, Google, Nissan by 2020. *The Dish Daily*. November 22.

Bell, M. G. and Iida, Y. (1997). *Transportation Network Analysis*. John Wiley & Sons. New York.

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.

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. http://www.rita.dot.gov/bts/sites/rita.dot.gov.bts/files/publications/national_transportation_statistics/index.html

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

Carter, Marc (2012). Volvo Developing Accident-Avoiding Self-Driving Cars for the Year 2020. *Inhabitat*. December 5.

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.

Fagnant, Daniel and Kockelman, Kara (2014a). Dynamic Ride-Sharing and Optimal Fleet Sizing for a System of Shared Autonomous Vehicles. Proceedings of the 94th Annual Meeting of the Transportation Research Board, and under review for publication in *Transportation*.

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

Induct (2014). Navia Named "Product of the Future at CES". Induct's online Newsroom: http://induct-technology.com/en/category/news.

LeBeau, Philip (2013). General Motors on Track to Sell Self-Driving Cars. 7 October, CNBC.

Litman, Todd (2013). Transportation Cost and Benefit Analysis II – Travel Time Costs. Victoria Transport Policy Institute. http://www.vtpi.org/tdm/index.php.

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.

MacKenzie, Angus (2013). Bosch and Evatran Partner to Bring EV Wireless Charging System to the US. *Gizmag*. June 19. http://www.gizmag.com/bosch-evatran-inductive-charging-system-ev/27971/

Markoff, John (2014). Google's Next Phase in Driverless Cars: No Steering Wheel or Brake Pedals. *New York Times*. May 27. http://www.nytimes.com/2014/05/28/technology/googles-next-phase-in-driverless-cars-no-brakes-or-steering-wheel.html?_r=0

Martin, E. and S. Shaheen, (2011) The Impact of Carsharing on Household Vehicle Ownership. *Access.* 38: 22-27.

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

Nissan Motor Company (2013). Nissan Announces Unprecedented Autonomous Drive Benchmarks [Press Release]. http://nissannews.com/en-US/nissan/usa/releases/nissan-announces-unprecedented-autonomous-drive-benchmarks.

O'Brien, Chris (2012). Sergey Brin Hopes People will be Driving Google Robot Cars in "Several Years". *Silicon Beat*. http://www.siliconbeat.com/2012/09/25/sergey-brin-hopes-people-will-be-driving-google-robot-cars-in-several-years/

Pavone, M., S. Smith, E. Frazzoli, and D. Rus (2011). Load Balancing for Mobility-on-Demand Systems. *Robotics: Science and Systems Online Proceedings* 7.

Puget Sound Regional Council (2006). 2006 Household Activity Survey. Seattle, Washington. http://psrc.org/data/surveys/2006-household/.

United States Environmental Protection Agency (2014). All-Electric Vehicles: Compare Side-By-Side. http://www.fueleconomy.gov/feg/evsbs.shtml.