Effects of Internalizing Air Emissions Externalities on Optimal Ride-Hailing Fleet Technology Composition and Operations

Matthew B. Bruchon,[1] Jeremy J. Michalek,[1,2,a] Inês L. Azevedo[3]

[1]Department of Engineering & Public Policy,
Carnegie Mellon University, Pittsburgh, PA 15213
[2]Department of Mechanical Engineering,
Carnegie Mellon University, Pittsburgh, PA 15213
[3]Department of Energy Resources Engineering,
Stanford University, Stanford, CA 94305

[a] Corresponding author. E-mail: jmichalek@cmu.edu

February 24, 2020

COMPETING INTERESTS STATEMENT: The authors declare no competing interests.

Abstract

Ride-hailing services from transportation network companies, such as Uber and Lyft, serve the fastest growing share of U.S. passenger travel demand. The high use-intensity of ride-hailing vehicles is economically attractive for electric vehicles, which typically have lower operating costs and higher capital costs than conventional vehicles. We optimize fleet technology composition (mix of conventional vehicles (CVs), hybrid electric vehicles (HEVs), and battery electric vehicles (BEVs)) and operations to satisfy exogeneous trip demand at minimum cost, and we compare results across scenarios. In nearly all cases, the optimal fleet includes a mix of technologies. In present and future scenarios for Austin, TX, Los Angeles, CA and New York, NY, HEVs and BEVs make up the largest portion of vehicle distance traveled in optimized fleets, and CVs are used primarily for periods of peak demand (if at all). Across a wide range of scenarios for the three cities, internalizing life cycle air pollution and greenhouse gas emission externality costs (via a Pigovian tax) leads to increased fleet electrification, a shift in charging toward periods when the grid is cleaner, and a reduction in emissions externalities of 12-31% in our base cases and 2-80% across our sensitivity scenarios. In all cases, the optimal fleet mix, its dispatch strategy, and resulting air emissions change substantively when air emissions externalities are internalized, suggesting a role for policy.

1 Background

Passenger cars produce the largest share of greenhouse gas emissions within the US transportation sector, which recently surpassed electric power as the country's highest-emitting economic sector [1]. Passenger cars also emit substantial conventional air pollution, and increased mortality from US air pollution (28% of which results from transportation) is comparable to the number of deaths attributable to automobile accidents, with an annual social cost of $886 billion [2].

Ride-hailing services are rapidly and dramatically changing the passenger car landscape: From 2009 to 2017, for-hire vehicles in the United States more than doubled their share of trips and their daily per capita usage, due primarily to the rapid growth of ride-hailing services [3], and by 2016, 15% of intra-urban trips in San Francisco were served by Uber and Lyft [4].

Vehicle electrification has the potential to drastically reduce ride-hailing emissions while perhaps also lowering operating costs. Electricity is often cleaner and cheaper than gasoline per vehicle distance traveled (VDT), and for intensively used vehicles lower fuel costs and operation emissions might offset their higher upfront costs and production emissions. The Intergovernmental Panel on Climate Change recently stated that electric modes of transportation would "need to displace fossil-fuel powered passenger vehicles by 2035-2050 to remain in line" with pathways to hold global warming to 1.5°C [5]. Recognizing TNC electrification's opportunity to reduce transportation emissions, the California Public Utilities Commission in 2018 released an initial overview of regulatory approaches that are worth further research as a means to encourage TNC electrification, including technology mandates, distance-based fees on combustion engine usage, and financial incentives [6]. Also in 2018, Uber announced a goal of an all-EV fleet within the city of London by 2025. This plan's stated motivation is to reduce pollution, and a per-mile "clean air fee" will fund driver financing programs [7]. Advances in vehicle electrification and automation may transform the way ride-hailing services operate [8].

However, the premise that full fleet electrification is a viable or desirable policy goal warrants further investigation. At the current cost of lithium-ion batteries, battery electric vehicles (BEVs, which plug in to charge and rely entirely on electricity stored in large battery packs) have a much higher upfront cost than conventional vehicles (CVs); battery manufacturing emissions are non-trivial [9, 10]; and, depending on region, timing, and vehicle design, electric vehicles do not always reduce air pollutant emissions or greenhouse gas emissions externalities compared to CVs (with lower-income census block groups more likely to face increased emissions externalities from BEVs) [11, 12, 13, 14]. Furthermore, the operations of BEVs suffer from logistical constraints of limited range and slower refueling (charging). BEVs cannot service demand while charging, so a larger fleet is required to satisfy a given level of demand. BEVs also must detour to recharge, increasing VDT. In contrast, gasoline hybrid electric vehicles (HEVs, which draw all net energy from gasoline but use a battery and electric motor to improve efficiency) have no additional range or refueling constraints, but they do burn gasoline and emit pollution from the tailpipe. In general, it may be that the lowest-cost or lowest-emission fleet does not use a single homogeneous technology but, rather, a mixture of technologies, with different duty cycles (peak versus off-peak) being served by different technologies.

We investigate the optimal technology mix and operations of a ride-hailing fleet whose operator has perfect knowledge of exogenous demand and total control over fleet acquisition and operations. Control of fleet acquisition may represent ride-hailing companies that own and lease vehicles in some locales, and centralized vehicle routing may become widespread as autonomous vehicle technology advances, whereas today's ride-hailing services only approximate centralized routing via human drivers responding to price signals. We assess the policy opportunity of electrification by comparing costs and emissions of pure CV, HEV, and BEV fleets with mixed fleets across a range of scenarios. By comparing cases that include or exclude emissions externality costs in fleet optimization, we assess the degree to which unpriced emissions externalities bias fleet outcomes away from socially optimal solutions and consider whether policy intervention may be therefore justified on economic efficiency grounds.

1.1 Background

A body of literature considers operations and outcomes of electrified vehicle fleets, but the question of electrification's role within a ride-hailing fleet's optimal technology mixture and its impact on resulting emissions is relatively unexplored.

Some studies use agent-based modeling (ABM) to explore operational impacts of homogeneous all-electric fleets. Bauer et al. (2018) estimate that such a fleet operating in Manhattan would reduce private costs and emissions relative to a homogeneous fleet composed of either CVs or HEVs; our fleet differs in its optimization of fleet mix under different objectives and its consideration of multiple cities. Chen et al. (2016) find that electrification can meet ride-hailing demand while barely increasing empty VDT, but only if the fleet size is increased [15].

Other studies use an ABM to model trips in combination with a second model. Sheppard et al. (2019) use ABM to generate simplified operational parameters for a national-scale optimal sizing of vehicles and infrastructure for an all-electric fleet, finding 12.5 million vehicles could replace the fleet of 276 million personally owned vehicles. Chen & Kockelman (2016b) incorporate a logit model choice model into an ABM to estimate that a shared, autonomous, all-electric vehicle fleet could capture 14% to 39% of all passenger trips within the Austin, Texas region, depending on pricing [16].

Studies employing ABM use simplifying assumptions or heuristics to model agents' behavior. Those heuristics' ability to achieve representative behavior or near-optimal behavior cannot easily be evaluated for each test case, so comparisons across scenarios can conflate effects of the scenarios with effects of the heuristics. Specifically, it is difficult to determine the degree to which differences in results across scenarios are due to differences in the scenarios themselves or due to differences in the performance of the heuristics across scenarios. Bertsimas et al. (2018) find that for vehicle routing problems, optimization coupled with well-designed heuristics increase fleet revenue results by as much as 9% relative to a heuristic alone and that heuristics perform unevenly across problem instances [17]; it is conceivable that this 9% gap widens when a fleet's technology mix is jointly optimized with its routing. Heuristics are necessary to address city-scale problems at manageable computational cost, but they introduce challenges for comparing across cases – such as comparing solutions when air emissions externality costs are or are not internalized. To address this limitation, we pair

heuristics with mathematical optimization to understand heuristic quality, to gain intuition on their biases, and to compare fairly across cases.

There is also a separate stream of methodologically-focused research applying optimization to the routing of range-limited electric vehicles. These are typically conducted at a very small scale (exact solutions for 100-200 trips or heuristic solutions for several hundred more), do not consider external costs, and rarely jointly optimize purchases and routing even in cases when a fixed mixture of powertrains is assumed [18, 19]. The SI describes some of these studies in greater detail. Optimizing fleet size and mix at any scale requires careful model formulation and development of problem-specific heuristics, which this study contributes for its problem (applied to an instance of 5,000 trips).

Additionally, in the grey literature a 2019 International Council on Clean Transportation report examined powertrain choice from the perspective of a TNC driver's personal investment decision [20]. It found that hybrids may be financially favorable and that battery vehicles may become favorable around 2023-2028, using assumptions for factors such as the total distance traveled per year, that, in practice, vary across vehicles in the fleet.

We contribute to the prior literature by (1) constructing a mixed-integer optimization model with heuristics that make meaningfully-sized problems tractable and provide near-optimal solutions for fair comparisons across scenarios for three cities and (2) applying the model to characterize how the optimal technology mix, operations, and life cycle air emissions externalities of a TNC fleet change across scenarios representing geographic and temporal variation, uncertainty, and the internalization of air emissions externalities (as a Pigovian tax passed through to the fleet operator). Our model is also unique in its treatment of vehicle costs, which incorporates into the optimization the effect vehicle usage has on period of use, resale value at end of use, and the resulting discounted future cash flow.

We include air emission externalities across the vehicle life cycle from greenhouse gas emissions (including carbon dioxide, methane, and nitrous oxide) and from criteria air pollutants (particulate matter, nitrogen and sulfur oxides, and secondary particulate matter from emissions of volatile organic compounds) using reduced complexity models that estimate health damages caused by emissions of air pollutants. We use TNC trip data from Austin, Texas to represent TNC demand, but to consider how findings vary from city to city, we also model Los Angeles and New York by changing parameters related to energy prices, health damage impacts, and marginal emissions from the electric grid to represent each location. In each scenario we find the fleet size, technology composition, and vehicle routing combination that satisfies TNC trip demand (matching origin-destination location and time) at minimum cost.

2 Results

Across a range of scenarios for three cities—Austin, Los Angeles, and New York City—we find the optimal fleet composition and operations for (1) minimizing private costs and (2) minimizing private costs plus air emissions externality costs,[1] and we compare resulting outcomes of policy interest. Each test case has the same total trip-miles, since demand is

[1] as though the firm faces a Pigovian tax on direct emissions as well as other life cycle emissions passed through suppliers to the fleet operator without inducing other changes in the economy

exogenous and must be met, and we present results per trip-mile with outcomes annualized and monetary values in 2018 USD. Costs labeled as "external" refer to air emissions externalities from vehicle manufacture and use (computed with a social discount rate of 3%), and costs labeled as "social" refer to the sum of private and external costs.

In the base test case for each city, we assume a 7% real discount rate used by the fleet operator, no labor costs, a BEV price of $33,950 (2019 Kia Soul), $50/ton CO_2 externality valuation, the AP3 model of conventional air pollution emission mortality effects, $9.41 million value of statistical life, and the Pope et al. (2019) air pollution concentration-response function. We use a trip dataset from Austin for all three cities, but private and external costs related to gasoline (at the tailpipe and refinery) and electricity vary across cities. These assumptions are discussed in Methods & Materials.

We first describe the impacts of a Pigovian tax on our results and assess the cost reductions possible through technology mixing. We then summarize key results from an extensive sensitivity analysis. In the SI we provide additional analysis of base case results and a range of sensitivity cases.

2.1 Impact of a Pigovian tax

Figure 1 summarizes private, external, and social cost results for each city when optimized with and without a Pigovian tax. Across cities in our base case, a Pigovian tax leads the fleet to purchase and dispatch its vehicles in a manner that raises private costs per trip-mile by 1% to 2% but reduces emissions externalities per trip-mile by 12% to 31%. Reduced external and social costs are greatest in Los Angeles, where fuel emissions externalities are high and electricity generation externalities are low relative to the other cities modeled. They are smallest in New York City, where the external costs of electricity generation are highest of the three cities.

To put these per trip-mile results in context, a recent Fehr & Peers consulting report estimated that Uber & Lyft drive 104 million monthly trip-miles in Los Angeles [21]. Multiplying those trip-miles by the 0.8¢/trip-mile increase in private costs and 2.1¢/trip-mile decrease in externalities, we can roughly estimate private cost increases of $11 million and external cost reductions of $27 million per year in Los Angeles.

The net effect of these cost changes is a reduction in overall social costs ranging from 1% to 2%. While this net effect is small, the distributional impacts are significant since the tax shifts the fleet's external costs away from the public, many of whom do not benefit from the fleet's services, and onto the fleet operator (and potentially its customers).

These external cost reductions are accomplished in each city by shifting away from gasoline usage (in CVs and HEVs) and increasing electricity usage (in BEVs). Figure 2 illustrates for each city, with and without the Pigovian tax, the vehicles purchased and the annual miles driven by cars of each powertrain type. For all three cities the Pigovian tax results in increased fleet electrification, both per vehicle and per mile, but the details of each city's private-optimal and socially-optimal fleets differ:

- In Austin, where gas prices are low relative to other modeled cities, a private-cost-minimizing fleet is composed of a majority of CVs, but those CVs are primarily used infrequently, during periods of relatively high demand, while HEVs serve as "baseload"

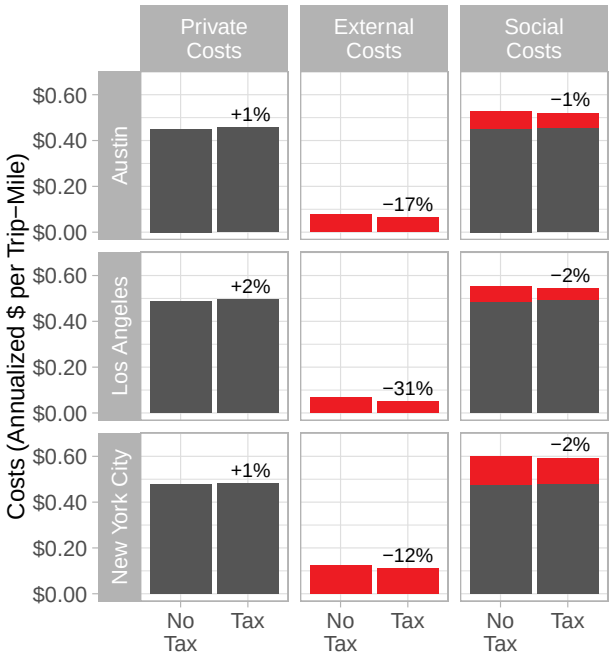


Figure 1: Average private costs (colored black), external costs from air emissions (colored red), and social costs (private + external) per trip-mile for TNC fleets in three cities optimally purchased, routed, and charged to minimize private costs while either excluding or including a Pigovian tax on air emissions externalities. The percentage change in each cost component caused by internalizing those externalities is annotated. Assumed private and external costs of energy inputs vary across cities as described in Sections 4.5 and 4.6. All cases use a 7% real private firm discount rate, no labor costs, $33,950 BEV price (2019 Kia Soul), $50/ton CO_2 externality price, the AP3 damage model, $9.41 million (2018) value of statistical life, and the Pope et al. (2019) concentration-response function.

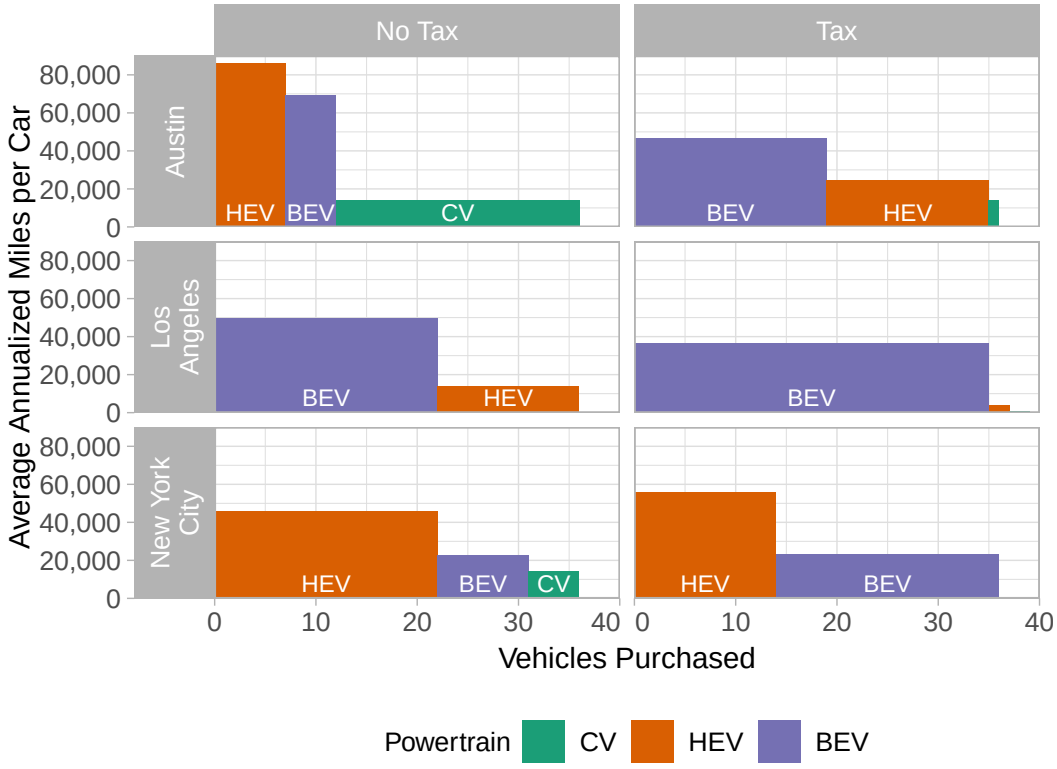
supply and are responsible for a plurality of total miles driven. When air emissions externalities are internalized, the fleet uses BEVs to serve baseload and HEVs for nearly all remaining trips, almost entirely eliminating CV usage.

- In Los Angeles, a private-cost-minimizing fleet uses no CVs due to higher gasoline prices. Instead, BEVs serve a majority of demand with HEVs used primarily in periods of high demand. Due to high gasoline externalities and low electricity externalities, a Pigovian tax results in a fleet that is almost entirely composed of BEVs.

- In New York City, where gasoline is more expensive than Austin but cheaper than Los Angeles, a private cost-minimizing fleet relies heavily on HEVs, using a mix of BEVs and CVs for high demand periods. A Pigovian tax eliminates CVs from the fleet and makes the fleet majority-BEV, but due in part to high externalities of electricity generation, HEVs are still used as baseload.

Across the three cities, the number of BEVs in the optimal fleet increases by 59% to 280% when a Pigovian tax is imposed on the fleet, and BEVs' total vehicle-distance traveled

Figure 2: Vehicle purchases (x-axis) and average utilization (y-axis) by powertrain type for cost-minimizing fleets when excluding (left) and including (right) a Pigovian tax on air emissions. Private and external costs of energy inputs vary across cities. All cases use a 7% real private firm discount rate, no labor costs, $33,950 BEV price (2019 Kia Soul), $50/ton CO_2 externality price, the AP3 air emissions damage model, $9.41 million (2018) value of statistical life, and the Pope et al. (2019) concentration-response function.

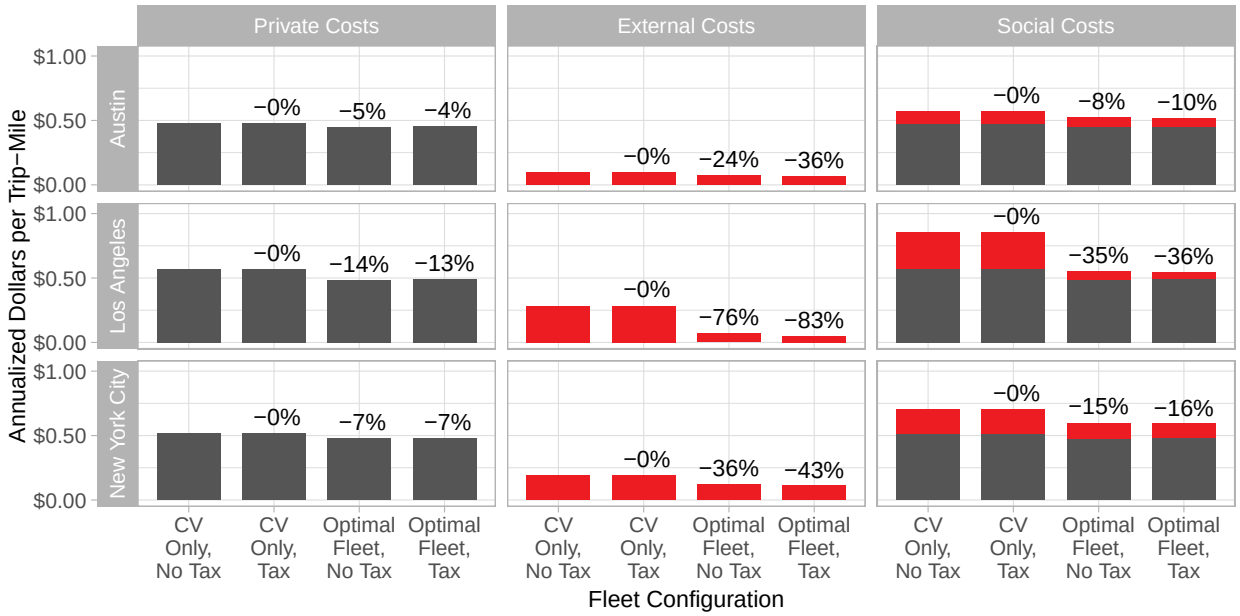


increases by between 17% to 155%. HEVs serve virtually all of the remaining demand in these three Pigovian tax cases, while CVs are at or near 0% of the fleet's purchases and distance traveled.

2.2 Value of optimally mixing technologies

Across cities, a fleet that optimally determines the mixture of powertrains to purchase and dispatch substantially reduces its private costs and the air emissions externalities it produces. Figure 3 illustrates each cost component for four fleet configurations: (1) an all-CV fleet optimized for private costs with no Pigovian tax on emissions externalities (arguably closest to the business-as-usual case of present-day fleets); (2) the same fleet facing a Pigovian tax; (3) a mixed fleet optimized for private costs; and (4) the same fleet facing a Pigovian tax. For an all-CV fleet, internalizing emissions externalities has virtually no ability to reduce them because routing decisions for CVs that minimize private fuel and capital costs also nearly minimize external costs (a very small reduction occurs because internalizing externalities alters tradeoffs between energy usage and distance-based acquisition costs).

Figure 3: Average private costs, external costs from air emissions, and social costs (private + external) per trip-mile for TNC fleets in three cities, considering an all-conventional vehicle fleet ("CV Only") and an optimally mixed fleet ("Optimal Fleet") with and without a Pigovian tax on air emissions externalities. The percentage cost reduction relative to the "CV Only, No Tax" case is annotated. Assumed private and external costs of energy inputs vary across cities. All cases use a 7% real private firm discount rate, no labor costs, $33,950 BEV price (2019 Kia Soul), $50/ton CO_2 externality price, the AP3 damage model, $9.41 million (2018) value of statistical life, and the Pope et al. (2019) concentration-response function.



Relative to an all-CV fleet, an optimally mixed fleet reduces private costs by 5% to 14% and reduces emissions externalities by 24% to 83%. In the SI, we also compare the optimal mixed fleet to optimal homogeneous fleets composed of either CVs HEVs, or BEVs. Across the three cities, the best homogeneous fleet does not depend on a Pigovian tax: it is all-HEV in Austin and New York City and all-BEV in Los Angeles regardless of whether a tax is included. Relative to the best homogeneous fleet, the mixed fleet optimized without a Pigovian tax reduces private costs by 1% to 4%, and the mixed fleet with a Pigovian tax reduces social costs by 2% to 4%.

Unlike all-CV and all-HEV fleets, internalizing an all-BEV fleet's emissions externalities can shift charging to lower-polluting times of day to reduce externalities, as shown in the SI. The magnitude of this reduction ranges from 2-4% depending on scenario.

2.3 Sensitivity to model inputs

Here, we briefly summarize findings across additional test cases. The SI contains more exhaustive results from a larger set of sensitivity cases.

2.3.1 Damage model and social cost of carbon

Our base case assumes a \$50/tonne CO_2 externality valuation and uses the AP3 reduced complexity model to estimate health damages from criteria air pollutants. Using a very high CO_2 externality valuation of \$300/tonne increases private-optimal externality estimates by a factor of roughly two to three across cities and leads the fleet to use nearly all BEVs, reducing externalities as much as 43% (in Austin).

Using either EASIUR or InMAP, two alternative reduced complexity models, rather than AP3 leads to similar or increased BEV uptake in Pigovian tax scenarios, and it leads to estimated externality reductions as high as 81% (in New York City using InMAP).

2.3.2 Fleet operator discount rate, car resale, and labor compensation

Our base case assumes that the fleet pays drivers no hourly wages (instead assuming either a flat percentage of fare or driverless cars) and makes its decisions using a 7% real discount rate for future operation costs and future resale value of its vehicles at the end of TNC use (resale value estimation described in Section 4.5, Eqn. 3).

Using a lower discount rate of 1%, the fleet places greater value on the future cash flow from reselling each car; this reduces the capital cost advantage of CVs and they are used negligibly even when minimizing private costs. With a higher discount rate of 13%, the capital cost advantage of CVs increases, but when a tax is introduced, they still serve less than 29% of VDT and BEVs still serve 31%-99% of VDT.

If we instead assume fleet vehicles have no resale value, outcomes shift slightly. This is because our resale value regression model estimates faster depreciation for BEVs than for HEVs, and faster depreciation for HEVs than for CVs. This means when resale value is removed, CVs' lower effective purchase costs increase by more than HEVs or BEVs in percentage terms, but the gap is narrow in absolute dollar terms. When minimizing private costs in Austin, for example, CVs fall from 26% in the base case (7% discount rate) to 23% of total VDT (no resale value).

If we assume the fleet pays its drivers an hourly wage of \$12, including when BEVs must go out of service to recharge, BEV VDT decreases by 17% (in Los Angeles) to 78% (in Austin).

2.3.3 Battery capacity increases and cost reductions

Our base case cases uses the 2018 Kia Soul with a retail price of \$33,950 and a 30 kilowatt-hour (kWh) battery. If we assume its battery cost were to fall by \$5,000 (a cost reduction of \$167/kWh) then the private cost-minimizing fleet would be majority-BEV in each city and a Pigovian tax would lead BEVs to serve 59%-99% of VDT.

If we instead used a 2020 Chevrolet Bolt as the reference vehicle, with a price of \$36,620 and a 66 kWh battery, a tax increases electrification slightly in Austin, where the higher sticker price makes BEVs less competitive as baseload, and increases it slightly in New York City, where each BEV can serve additional trips in high-demand periods. In Los Angeles, where BEVs already served a nearly all VDT, they serve roughly the same share of VDT but require fewer purchases to do so.

3 Discussion

Our results consistently suggest that internalizing air emissions externalities results in a greater degree of electrification (shift from CV to HEV and BEVs and shift from HEV to BEV) as well as operational changes that together reduce air emissions external costs (by 12% to 31% in the base case and 2% to 81% in sensitivity cases, depending on the city), and lower social costs (by 1% to 2% in the base case and 1% to 18% in sensitivity cases, depending on the city). This suggests a potential role for policy because when emissions externalities are unpriced, firms have incentives to improve economic efficiency by lowering private cost, increasing air emissions, implementing a lower degree of electrification, and charging BEVs when the grid is less clean than socially optimal. While the social cost improvement is fairly small across most of the scenarios examined, the change in who bears the cost can be significant. As one very rough estimation of magnitude, multiplying our results for percentage reduction in external costs per trip-mile by one consulting report's estimate of monthly TNC trip-miles in Los Angeles (our only analysis city covered by that study),[21]–the result would represent around $27 million dollars of annual environmental and health outcomes for that city.

Pigovian taxes offer efficiency and flexibility, but in the absence of such an option, other policies that encourage similar outcomes, such as policies encouraging increased electrification, could potentially improve economic efficiency. However, any such policy should be designed with care. A blunt instrument favoring one technology over others may not be desirable because (1) the optimal fleet is generally a mixed fleet; (2) beyond fleet composition, it is important how intensively each vehicle type is used; and (3) factors that vary with location and over time, like energy prices, vehicle cost, population density, and grid emission factors, can dramatically change the degree of electrification that is optimal.

These results should be interpreted in context. We model a single ride-hailing firm with perfect information and full control of fleet acquisition and operation that must meet all demand. In practice, current ride-hailing fleets in the US are staffed by workers who choose their own vehicles, which are often serve dual uses as personal vehicles, and choose when to work in response to price incentives. Ride-hailing fleets also have limited ability to anticipate future demand. Our model may approximate today's dispatch to the degree that accurate demand prediction is possible and to the degree that drivers respond to incentives about when to work, but we ignore the pricing mechanisms altogether, as well as the potential for dual-use vehicles. Our model may be a better approximation of a future automated fleet centrally owned and routed by the ride-hailing firm.

Ride-hailing services also need not meet all demand at the exact start and end time they were served in the RideAustin data. Flexibility on passenger pickup and dropoff time could improve system efficiency and potentially change optimal fleet composition and operations [22].

Despite these limitations, the ability to observe changes in optimal fleets under a variety of scenarios helps in developing intuition about fleet technology choices and operations as well as implications of failing to price externalities.

We discuss a range of additional cases and considerations in greater depth in the SI.

4 Materials and Methods

We construct an optimization model to choose fleet composition (mix of CVs, HEVs, and BEVs) and operations (vehicle routing and BEV charging) in order to minimize the cost of satisfying exogeneous demand (origin and destination location and time) under a range of scenarios. We first describe our model and the customized methods we develop to solve it at scale, and then we describe the data that we use to instantiate the model.

4.1 Optimization Model

Figure 4 illustrates our modeling framework with an example. Vehicle purchase choices determine the vehicles available to dispatch (left). Routing options, jointly optimized with purchases, are represented using a graph, where each vertex (dot) represents a specific place and time, and the arcs connecting them include available options for:

- Trip arcs: passenger trip requests that must be served
- Charging arcs: spending time parked at a charging location for recharging a battery or waiting for the next trip
- Dispatch arcs: routing of a vehicle from its home base to the first served trip
- Return arcs: routing of a vehicle from its final served trip back to its home base
- Relocation arcs: routing of a vehicle from the end of one passenger trip to the beginning of a next passenger trip or between passenger trips and recharge locations

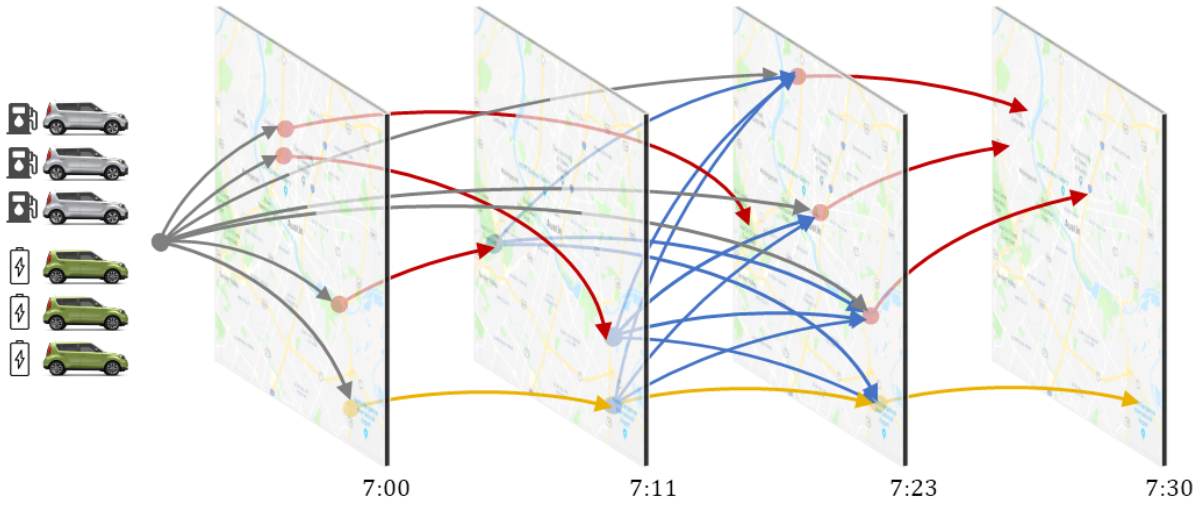


Figure 4: Illustration of time-space graph showing passenger trips (red), charging arcs (yellow), dispatch arcs (gray), and relocation arcs (blue). Some arcs are omitted for simplicity.

In describing our model, we first define the full mixed-integer linear programming (MILP) model `FullMILP` used to represent this problem then describe a set of heuristics that we use to improve scalability.

4.2 MILP formulation

Our `FullMILP` formulation, shown in Table 1, finds the cost-minimizing fleet technology mix and assignment of vehicles to trip arcs where the set of decision variables $\mathcal{X}$ includes the number of vehicles n_k of each powertrain type k purchased, assignments $a_{k,i,j}$ of vehicles k to arcs (i,j), charge level $q_{k,t}$ and energy charged from the grid $\Delta q_{k,t}^{\text{CHG}}$ for each vehicle k at each discrete time point t, and total annualized capital cost κ_k for each vehicle type (determined by vehicle utilization levels) for all vehicle types $k \in \mathcal{K}$, arcs $(i,j) \in \mathcal{A}$, and times $t \in \mathcal{T}$. The full set of notation is shown in Table 2.

In all test cases, the objective function, Eq.(1a), sums the relevant vehicle purchase costs κ_k, gasoline and per-mile maintenance costs $c_{k,i,j}$, and time-varying battery charging costs c_t. In cases where air emissions externalities are internalized, $\tau = 1$, so the fleet also considers a Pigovian tax on externalities from manufacturing emissions δ_k, tailpipe and fuel refining emissions $d_{k,i,j}$, and grid emissions d_t.

At the core of `FullMILP` are equations that are standard for many vehicle routing problems. Constraint 1b ensures preservation of flow for each vehicle through the network (forcing vehicles to return to the depot after serving trips), Constraint 1c requires that all passenger trips be satisfied, and Constraint 1d requires that a vehicle must be purchased to be dispatched. The remainder of the formulation is customized for our case.

Constraints 1e–1f model capital costs and manufacturing damages. For all vehicles, Constraint 1e uses a set of linear constraints Ω to define a convex piecewise linear cost floor representing the sum of annualized vehicle costs, including salvage value, which is a function of vehicle assignment. We discuss this aspect of our formulation in more detail in the SI.

Constraints 1h-1i manage BEV charge level. Constraint 1g applies to timesteps at which regular interval charging timesteps begin (default 15-minute intervals), defines charger usage, and tracks charge level changes. Constraint 1h applies to all other timesteps, at which there is no charging option, so that charge level is fully determined by traversed arcs' energy requirements. Constraint 1i enforces bounds of BEV charge levels. The implied amount of electricity purchased from the grid is quantified for the objective function in the "where" statement as the change in charge unexplained by travel.

The set of vehicle types $k \in \mathcal{K}$ indexes individual vehicles for BEVs (each with binary purchase and routing decisions) but groups vehicles into types for CVs and HEVs (with integer purchase and routing decisions) for computational efficiency. This grouping means that `FullMILP` assumes refueling time and routing of CVs and HEVs are negligible, such that individually tracking fuel level is unnecessary and `FullMILP` need not separately index each car. Aside from these refueling implications, CV and HEV dispatch are otherwise representative of a fleet of discrete vehicles.

Table 1: Formulation of the `FullMILP` optimization problem.

$$\underset{\mathcal{X}}{\text{minimize}} \sum_{k\in\mathcal{K}} \kappa_k + \sum_{(i,j)\in\mathcal{A}} c_{k,i,j} a_{k,i,j} + \sum_{k\in\mathcal{K}} \sum_{t\in\mathcal{T}_+} c_t \Delta q^{\text{CHG}}_{k,t} + \tau \left[\sum_{k\in\mathcal{K}} \delta_k + \sum_{(i,j)\in\mathcal{A}} d_{k,i,j} a_{k,i,j} + \sum_{k\in\mathcal{K}} \sum_{t\in\mathcal{T}_+} d_t \Delta q^{\text{CHG}}_{k,t} \right]$$

Minimize private and external car, fuel, and electricity costs (1a)

subject to

$$\sum_{i\in\mathcal{V}} a_{k,i,j} = \sum_{i\in\mathcal{V}} a_{k,j,i} \qquad \forall k \in \mathcal{K}, j \in \mathcal{V}\backslash\{r,s\}$$

Flow is preserved across nodes except source and sink (1b)

$$\sum_{k\in\mathcal{K}} a_{k,i,j} = n_{i,j} \qquad \forall (i,j) \in \{\mathcal{A} : n_{i,j} > 0\}$$

Demand is satisfied (1c)

$$\sum_{j\in\mathcal{V}\backslash r} a_{k,r,j} = n_k \qquad \forall k \in \mathcal{K}$$

Dispatched cars are purchased (1d)

$$\kappa_k \geq \alpha^{\text{COSTS}}_{\omega,k} n_k + \beta^{\text{COSTS}}_{\omega,k} \sum_{(i,j)\in\mathcal{A}} m_{i,j} a_{k,i,j} \qquad \forall k \in \mathcal{K}, \omega \in \Omega_k$$

Per-vehicle private capital costs vary with usage (1e)

$$\delta_k \geq \alpha^{\text{DAMAGES}}_{\omega,k} n_k + \beta^{\text{DAMAGES}}_{\omega,k} \sum_{(i,j)\in\mathcal{A}} m_{i,j} a_{k,i,j} \qquad \forall k \in \mathcal{K}, \omega \in \Omega_k$$

Per-vehicle manufacturing external costs vary with usage (1f)

$$q_{k,t+1} \leq q_{k,t} + \sum_{(i,j)\in\{\mathcal{A}:t_i=t\}} a_{k,i,j} \Delta q^{\text{MAX}}_{k,i,j} \qquad \forall k \in \mathcal{K}_{\text{B}}, t \in \mathcal{T}_{\text{Q}}$$

BEV charge level is tracked (times with a charger starting timeslot) (1g)

$$q_{k,t+1} = q_{k,t} + \sum_{(i,j)\in\{\mathcal{A}:t_i=t\}} a_{k,i,j} \Delta q^{\text{MAX}}_{k,i,j} \qquad \forall k \in \mathcal{K}_{\text{B}}, t \in \mathcal{T}\backslash\mathcal{T}_{\text{Q}}$$

BEV charge level is tracked (times with no charger starting timeslot) (1h)

$$0 \leq q_{k,t} \leq q^{\text{MAX}}_k \qquad \forall k \in \mathcal{K}_{\text{B}}, t \in \mathcal{T}$$

Charge level does not exceed battery capacity (1i)

$$a_{k,i,j} \in \{0,1\},\ n_k \in \{0,1\} \qquad \forall k \in \mathcal{K}_{\text{B}}, (i,j) \in \mathcal{A}$$

BEV routing and purchase decisions are binary (1j)

$$a_{k,i,j} \in Z_+,\ n_k \in Z_+ \qquad \forall k \in \mathcal{K}\backslash\mathcal{K}_{\text{B}}, (i,j) \in \mathcal{A}$$

CV & HEV routing and purchase decisions are integral (1k)

$$q_{k,t} \in R_+ \qquad \forall k \in \mathcal{K}_{\text{B}}, t \in \mathcal{T}$$

BEV charge level is always nonnegative (1l)

$$\kappa_k \in R_+ \qquad \forall k \in \mathcal{K}$$

Per-vehicle private capital costs are nonnegative (1m)

$$\delta_k \in R_+ \qquad \forall k \in \mathcal{K}$$

Per-vehicle manufacturing external costs are nonnegative (1n)

where

$$\Delta q^{\text{CHG}}_{k,t} = q_{k,t+1} - q_{k,t} + \sum_{(i,j)\in\{\mathcal{A}:t_i=t,\Delta q^{\text{MAX}}_{k,i,j}<0\}} a_{k,i,j} \Delta q^{\text{MAX}}_{k,i,j} \qquad \forall k \in \mathcal{K}_{\text{B}}, t \in \mathcal{T}_{\text{Q}}$$

Charging of each BEV at each timestep equals charge change minus routing losses (1o)

Table 2: Sets, decision variables, and input parameters

Label	Type	Description
$\mathcal{X}$	Set	All decision variables
$\mathcal{V}$	Set	Vertices representing points in space-time
$\mathcal{A}$	Set	Arcs connecting feasible pairs of vertices in $\mathcal{V}$
$\mathcal{K}$	Set	Vehicles or vehicle types (BEVs are represented individually, whereas CVs and HEVs are each tracked as a group)
$\mathcal{K}_{\mathrm{B}}$	Set	Battery electric vehicles (subset of $\mathcal{K}$, indexed individually)
$\mathcal{T}$	Set	All unique arc start and end times
$\mathcal{T}_{\mathrm{Q}}$	Set	All unique charging arc start times (subset of $\mathcal{T}$)
Ω_k	Set	Linear constraints that make up the piecewise linear convex cost floor for capital cost κ_k for vehicle type k
n_k	Variable	Number of vehicle k purchased (BEVs are tracked individually, whereas CVs and HEVs are tracked as a group)
$a_{k,i,j}$	Variable	Assignment of vehicle k to arc (i,j)
$q_{k,t}$	Variable	Charge level of vehicle k at time t
$\Delta q_{k,t}^{\mathrm{CHG}}$	Variable	Energy charged to vehicle k from the grid at time t
κ_k	Variable	Private acquisition cost for vehicle k
δ_k	Variable	Manufacturing damages for vehicle k
τ	Parameter	Flag controlling whether air emissions externalities are included as a tax
r	Parameter	Source vertex from which all routes originate
s	Parameter	Sink vertex at which all routes terminate
t_i	Parameter	Time of vertex i
$n_{i,j}$	Parameter	Number of trips requested along arc (i,j)
$m_{i,j}$	Parameter	Travel distance along arc (i,j) (annualized)
m_{MAX}	Parameter	Maximum lifetime travel distance of a vehicle
q_k^{MAX}	Parameter	Energy capacity of vehicle k (∞ for CVs and HEVs)
$c_{k,i,j}$	Parameter	Private cost for vehicle k to traverse arc (i,j)
$d_{k,i,j}$	Parameter	External cost from vehicle k traversing arc (i,j)
c_t	Parameter	Private cost per kWh of electricity from the grid at time t
d_t	Parameter	External cost per kWh of electricity from the grid at time t
$\Delta q_{k,i,j}^{\mathrm{MAX}}$	Parameter	Maximum energy change for car k induced by travel on arc (i,j) (positive for charging arcs, negative for all others)
$\alpha_{\omega,k}^{\mathrm{COSTS}}$	Parameter	Intercept term for line ω, representing a portion of the convex piecewise linear VDT-dependent capital costs
$\beta_{\omega,k}^{\mathrm{COSTS}}$	Parameter	Slope term for line ω, representing a portion of the convex piecewise linear VDT-dependent capital costs
$\alpha_{\omega,k}^{\mathrm{DAMAGES}}$	Parameter	Intercept term for line ω, representing a portion of the convex piecewise linear VDT-dependent manufacturing external costs
$\beta_{\omega,k}^{\mathrm{DAMAGES}}$	Parameter	Slope term for line ω, representing a portion of the convex piecewise linear VDT-dependent manufacturing external costs

4.3 Heuristics

Solving the `FullMILP` problem with a standard commercial solver is impractical for city-scale problems with thousands of trips, particularly due to BEV charge constraints. To improve scalability, we introduce a set of customized heuristics that reduce problem size and tend to discover solutions quickly, allowing us to find near-optimal solutions to a sample of 5,000 trips. This is a larger instance than commercial tools can solve for many vehicle routing problem variants in reasonable time and larger than the optimization state of art for exact solutions (200 trips) described in the SI. We solve `FullMILP` first via a sequence of optimizations and heuristics:

1. A novel `MCF_VaryingFleetSize` heuristic reduces problem size by taking all feasible relocations from each trip to potential next trips and eliminating relocations that are likely to be higher-cost and therefore unused in optimal routing solutions. It adapts prior work [17] and uses `MCF_CarLimit`, a customization of the widely-known minimum-cost network flow problem [23].

2. A novel `ShrinkingBattery` heuristic builds an initial feasible solution from an aggregated simplification of the electric subset of the vehicle fleet, iteratively making the aggregation more realistic.

3. A customized variant of a widely-used `RuinAndRecreate` heuristic randomly selects pieces of the solution to re-optimize, improving the `ShrinkingBattery` solution.

4. The full `FullMILP` formulation is executed, taking the best solution found by Steps 1-3 as a starting point and upper bound on cost. It measures solution quality relative to a lower bound on cost defined by `FullMILP`'s linear relaxation, which is iteratively tightened. In many of the cases we test, this step simply verifies that the upper bound found by Steps 1-3 is within tolerance of the solution, but in some cases this step also improves the solution.

The `MCF_VaryingFleetSize`, `ShrinkingBattery`, and `RuinAndRecreate` heuristics constitute a substantial portion of this study's contribution—and this research question would be unanswerable at meaningful scale without them—but because they are all tools to help solve the `FullMILP` formulation, we present their underlying intuition and algorithmic steps in the SI.

4.4 Passenger trip data and driver relocations

We instantiate the model using a dataset of 1.5 million passenger trips from June 6, 2016 to April 13, 2017, released in 2017 by RideAustin, a nonprofit ride-hailing service in Austin, Texas. We extract passenger trip origin and destination, starting and ending timestamps, and distance traveled to define trip arcs $n_{i,j}$. We sample down to 5,000 trips using the weekday-season categories shown in SI Figure S11, plus a separate category for the high-demand days of the South By Southwest Festival. The number of trips sampled from each category is proportional to average daily demand (which increased season to season as RideAustin

became more popular), and costs and distance values (which affect capital costs of each vehicle) are scaled up to annual quantities based on the number of days per year represented by each category. For tractability, we use k-means clustering to group locations into 25 clusters and round times to the nearest 3 minutes. Because efficiency varies with driving conditions [12], we estimate each trip's efficiency for each powertrain type using average speed, computed from the known distance and duration, and interpolating efficiency (gallons or kWh per mile) between standard test city and highway drive cycles (EPA drive cycles with average speeds of 21.2 and 48.3 miles per hour, respectively) [24].

RideAustin data does not include travel between passenger trips. For every potential relocation from each trip to each subsequent trip (or charging node), we estimate the required distance traveled and duration using k-nearest neighbors regression [25] on the RideAustin trips. This method and its implications are described further in the SI.

Relocations from prior trip to next trip were disallowed if the actual time gap (from the first trip's end to the second trip's start) were shorter than the estimated relocation duration or longer than 30 minutes. For tractability, vehicles may chain trips more than 30 minutes apart but must park at the central charge station depot between trips. When the estimated duration is shorter than the time gap between trips, we assume the vehicle travels at the estimated speed for the estimated trip duration, then idles for the remainder of excess time (assuming the combustion engine, if applicable, is shut off using a start-stop system). For relocations between passenger trips and the charging station or the source/sink nodes, we instead assume the vehicle parks immediately at the station and departs the station as late as possible.

4.5 Vehicle and charger technology

We model a typical present-day ride-hailing vehicle with otherwise-identical CV, HEV, and BEV counterparts. For model year 2018 in the United States market, there are five light-duty passenger vehicles with BEV and CV variants. Of those, the Kia Soul is best-suited for ride-hailing due to sufficient backseat space, so we adopt it for this study. SI Figure S19 shows that its efficiency and range are representative of present-day BEVs excluding Teslas (too expensive for mass-market TNCs), the Chevrolet Bolt, and the BYD e6.

We assume one charging station (also the depot from which all vehicles must begin and end trip chains) and place it at the centroid of all trip origin-destinations. Because our results do not show substantial increase in VDT from BEVs routing to and from charging stations, we do not consider sensitivity cases with more charging stations. There is no capacity constraint for charging or parking at this location. The charger is the fast-charger specification (CHAdeMO) that is compatible with the Soul, which can charge its 30 kWh battery to 90% in 46 minutes (linearized to a rate of 35.2 kWh/hour for simplicity). The BEV has an MSRP of $33,950, a city efficiency of 27.3 kWh/100 miles, and a highway efficiency of 36.1 kWh/100 miles. The middle-trim version of the Soul CV is used, with an MSRP of $20,500, city efficiency of 26.1 miles per gallon (mpg), and highway efficiency of 30.9 mpg. The hypothetical hybrid version of the Soul's parameters are estimated using differences in cost and efficiency between the similarly-sized Kia Optima sedan's gasoline and hybrid variants, resulting in a cost of $25,000, a city efficiency of 40.7 mpg, and a highway efficiency of 39.7 mpg. CV and BEV variant efficiencies are taken from the fueleconomy.gov

[26]; their MSRPs were accessed from the manufacturer's product websites.

In the base case of present-day Austin, energy prices come from EIA-estimated 2017 Austin Energy annual averages for transportation sector retail electricity prices (10.90¢/kWh) and gasoline prices for 87-octane gasoline ($2.20). [27, 28].

To annualize vehicle purchase costs, the MSRP minus a discounted future cash flow from resale of the vehicle (whether due to high mileage or age) is multiplied by a capital recovery factor F, shown in Equation 2:

$$f_{\mathrm{CR}}\left(r,N\right)=\frac{r\left(1+r\right)^{N-1}}{\left(1+r\right)^{N}-1} \tag{2}$$

where N is the age, in years, of the vehicle at which it ceases fleet operation and is sold in the used car market (N may be non-integer) and r is the discount rate. Note that this capital recovery factor is for equivalent annual payments from years 0 to $N-1$ (rather than years 1 to N).

We assume that vehicles are retired from the fleet and sold in the used market after N_{MAX} years or d_{MAX} miles, whichever happens first. Given a private firm discount rate r, a vehicle purchase price p, and vehicle resale value function $v(N,d)$ that depends on age N and annual distance traveled d, the private costs of each vehicle investment are:

$$\kappa=\left(1-\frac{v(N,d)}{\left(1+r\right)^{N}}\right)p\times f_{\mathrm{CR}}(r,N) \tag{3}$$

where $N=\min(N_{\mathrm{MAX}},\frac{d_{\mathrm{MAX}}}{d})$ and d is defined for each vehicle k as $\sum_{(i,j)\in\mathcal{A}}m_{i,j}a_{k,i,j}$. Here we use the symbol κ for capital cost loosely because the MILP model treats κ as a decision variable bound below by a set of constraints that represent a piecewise linear convex function approximating Eq(3). We describe this in more detail in the SI.

We assume a private firm real discount rate of 7% (with annual inflation of 2%), a maximum vehicle age of $N_{\mathrm{MAX}}=12$ years, and a maximum VDT of $d_{\mathrm{MAX}}=170,000$ miles, based on Argonne National Laboratory's Greenhouse gases, Regulated Emissions, and Energy use in Transportation Model (GREET) [29]. For each powertrain type, a separate regression (described in the SI) estimated the relationship between age, miles driven, and resale value using resale values queried from Kelley Blue Book.

4.6 Air emission externality costs

In scenarios where external costs of emissions are considered, air emissions externalities from the manufacturing stage are added as a Pigovian tax on vehicle investments:

$$T_{\mathrm{MFG}}=\sum_{i\in\mathcal{P}}\gamma_{i}^{\mathrm{MFG}}c_{i}^{\mathrm{P}} \tag{4}$$

where $\mathcal{P}$ is the set of pollutants considered, $\gamma_{i}^{\mathrm{MFG}}$ is the quantity of pollutant i produced during manufacturing, and c_{i}^{P} is the external cost per unit of pollutant i emitted. We consider greenhouse gas emissions from CO_2, methane, and N_2O; we consider health damages from PM2.5, SO_X, and NO_X, and VOC.

To compute external costs per unit of greenhouse gas emissions, we adopt the social cost of carbon $50 per ton of CO_2 equivalent estimated by the Interagency Working Group on the Social Cost of Carbon [30]. For conventional air pollutants, external costs depend on emission location and we use the AP3 model [31] to compute and monetize estimated health damages associated with these emissions. AP3 is one of several reduced complexity models that estimate health impacts resulting from air pollution. In contrast to the estimates generated by complex chemical transport-based air pollution models, reduced complexity models generate estimates at an acceptable level of accuracy while enabling estimates to be found for large numbers of scenarios quickly.

We adopt estimates of emissions from manufacturing each vehicle technology from GREET [29], adjusting inputs to the modeled vehicles' curb weight and battery weight, and we assume that manufacturing emissions from each production step occur in U.S. counties where similar economic activity occurs. When air emission externalities are included, p in Eq.(3) is the vehicle's MSRP + T_{MFG}. When externalities are excluded, p is simply the vehicle's MSRP. The SI includes further details and input values.

Air emission externalities associated with vehicle operations were estimated in a similar manner.

$$T_{\mathrm{OP}} = \sum_{i \in \mathcal{P}} \gamma_i^{\mathrm{OP}} c_i^{\mathrm{P}} \tag{5}$$

As with manufacturing emissions, we use the social cost of carbon and AP3 to estimate damages per unit of pollutant emitted from vehicle operations. We adopt GREET tailpipe emissions estimates for each technology, compute emission rates per gallon of gasoline consumed, and adjust rates based on the fuel consumption rate of each vehicle technology on each route arc. We assign tailpipe emissions to each scenario's relevant county (Travis County, TX in the base case of Austin). For upstream emissions associated with BEV charging, we use estimates of emissions damages from power plants that act at the margin in the eGrid subregion containing the city for each scenario (ERCOT in the case of Austin) based on updated data from Siler-Evans et al. (2012) [32, 33], averaged by season and hour of day in 2017. For upstream emissions associated with feedstock production and transportation for making gasoline and fuel for power plants, we use emissions quantities from GREET and AP3 damage factors of refinery counties in the city's region (described further in SI).

All non-grid damages reflect AP3 damages from a Center for Air, Climate, and Energy Solutions database [31, 34]. For emissions damages from all sources, we use the AP3 damage model, a value of statistical life of $9.41M (2018), a carbon price of $50/ton, and the Pope et al. (2019) concentration-response relationship [35].

References

[1] EPA OARCCDUS (2018) Inventory of U.S. greenhouse gas emissions and sinks, Technical report.

[2] Goodkind AL, Tessum CW, Coggins JS, Hill JD, Marshall JD (2019) Fine-scale damage estimates of particulate matter air pollution reveal opportunities for location-specific mitigation of emissions. *Proceedings of the National Academy of Sciences of the United States of America* 116(18):8775–8780.

[3] Conway M, Salon D, King D (2018) Trends in Taxi Use and the Advent of Ridehailing, 19952017: Evidence from the US National Household Travel Survey. *Urban Science* 2(3):79.

[4] San Francisco County Transportation Authority (2017) TNCs Today: A Profile of San Francisco Transportation Network Company Activity.

[5] IPCC (2018) Chapter 4: Strengthening and Implementing the Global Response, Technical Report June 2018.

[6] George SR, Zafar M (2018) Electrifying the Ride-Sourcing Sector in California.

[7] Uber (2018) The Clean Air PlanHelping London go electric.

[8] Fulton L, et al. (2017) Three Revolutions in Urban Transportation.

[9] Tessum CW, Hill JD, Marshall JD (2014) Life cycle air quality impacts of conventional and alternative light-duty transportation in the United States. *Proceedings of the National Academy of Sciences of the United States of America* 111(52):18490–5.

[10] Michalek JJ, et al. (2011) Valuation of plug-in vehicle life-cycle air emissions and oil displacement benefits. *Proceedings of the National Academy of Sciences of the United States of America* 108(40):16554–8.

[11] Weis A, Jaramillo P, Michalek J (2016) Consequential life cycle air emissions externalities for plug-in electric vehicles in the PJM interconnection. *Environmental Research Letters* 11(2):024009.

[12] Yuksel T, Tamayao MAM, Hendrickson C, Azevedo IML, Michalek JJ (2016) Effect of regional grid mix, driving patterns and climate on the comparative carbon footprint of gasoline and plug-in electric vehicles in the United States. *Environmental Research Letters* 11(4):044007.

[13] Holland SP, Mansur ET, Muller NZ, Yates AJ (2019) Distributional effects of air pollution from electric vehicle adoption. *Journal of the Association of Environmental and Resource Economists* 6(S1):S65–S94.

[14] Holland SP, Mansur ET, Muller NZ, Yates AJ (2016) Are there environmental benefits from driving electric vehicles? the importance of local factors. *American Economic Review* 106(12):3700–3729.

[15] Chen TD, Kockelman KM, Hanna JP (2016) Operations of a shared, autonomous, electric vehicle fleet: Implications of vehicle & charging infrastructure decisions. *Transportation Research Part A: Policy and Practice* 94:243–254.

[16] Chen TD, Kockelman KM (2016) Management of a Shared Autonomous Electric Vehicle Fleet. *Transportation Research Record: Journal of the Transportation Research Board* 2572.

[17] Bertsimas D, Jaillet P, Martin S (2017) Online Vehicle Routing : The Edge of Optimization in Large-Scale Applications. *Operations Research* pp. 1–48.

[18] Pourazarm S, Cassandras CG, Wang T (2016) Optimal routing and charging of energy-limited vehicles in traffic networks. *International Journal of Robust and Nonlinear Control* 26(6):1325–1350.

[19] Sassi O, Cherif WR, Oulamara A (2015) Vehicle Routing Problem with Mixed fleet of conventional and heterogenous electric vehicles and time dependent charging costs. *International Journal of Mathematical, Computational, Physical, Electrical and Computer Engineering* 9(3):167–177.

[20] Pavlenko N, Slowik P, Lutsey N (2019) When does electrifying shared mobility make economic sense?, (ICCT), Technical Report January.

[21] Balding M, Whinery T, Leshner E, Womeldorff E (2019) Estimated TNC Share of VMT in Six US Metropolitan Regions.

[22] Alonso-mora J, Samaranayake S, Wallar A, Frazzoli E, Rus D (2017) On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment. *Proceedings of the National Academy of Sciences of the United States of America* 114(3):462–467.

[23] Klein M (1967) A Primal Method for Minimal Cost Flows with Applications to the Assignment and Transportation Problems. *Management Science* 14(3):205–220.

[24] US Department of Energy Office of Energy Efficiency and US Environmental Protection Agency (2018) Detailed Test Information.

[25] Wang H, Kuo YH, Kifer D, Li Z (2016) A simple baseline for travel time estimation using large-scale trip data in *Proceedings of the 24th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems - GIS '16*. (ACM Press, New York, New York, USA), pp. 1–4.

[26] US Environmental Protection Agency (2018) Download Fuel Economy Data: 2018 Datafile.

[27] US Energy Information Administration (2018) Retail Gasoline and Diesel Prices: Regular - Conventional Areas, Annual Period.

[28] US Energy Information Administration (2018) Average retail price of electricity, Texas, annual.

[29] Argonne National Laboratory (2019) Greenhouse gases, Regulated Emissions, and Energy use in Transportation Model.

[30] United States Government (2016) Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866. *Interagency Working Group on Social Cost of Greenhouse Gases* (August):1–21.

[31] Muller NZ, Mendelsohn R (2006) The Air Pollution Emission Experiments and Policy analysis model (APEEP) Technical Appendix. *Yale University2* 1.

[32] Siler-Evans K, Azevedo IL, Morgan MG (2012) Marginal Emissions Factors for the U.S. Electricity System. *Environmental Science & Technology* 46(9):4742–4748.

[33] Azevedo IL, Horner N, Siler-evans K, Vaishnav P (2017) Electricity Marginal Factors Estimates. *Center for Climate and Energy Decision Making*.

[34] Center for Air C, Solutions E (2018) CACES Reduced Complexity Models Data Download.

[35] Pope CA, et al. (2019) Mortality Risk and Fine Particulate Air Pollution in a Large, Representative Cohort of U.S. Adults. *Environmental health perspectives* 127(7):77007.

Supplementary Information for

Effects of Internalizing Air Emissions Externalities on Optimal Ride-Hailing Fleet Technology Composition and Operations

Matthew B. Bruchon, Jeremy J. Michalek, and Inês L. Azevedo

Corresponding Author: Jeremy J. Michalek
E-mail: jmichalek@cmu.edu

This document includes:

Supplementary text
Figs. S1 to S20
Tables S1 to S17
SI References

Contents

1 Additional discussion of results **3**
A Base case results analysis . . . 3
B Model limitations . . . 3

2 Sensitivity cases **4**
A TNC operating models . . . 4
A.1 Fleet configuration . . . 4
A.2 Discount rates . . . 4
A.3 Resale value . . . 4
A.4 Labor costs . . . 4
B Battery technology . . . 4
B.1 Battery capacity . . . 4
B.2 Battery price . . . 4
C External cost inputs . . . 4
C.1 Alternate damage models . . . 5
C.2 Social cost of carbon . . . 5

3 Data and Modeling Assumptions **5**
A Trip and relocation data . . . 5
A.1 Trip data preprocessing . . . 5
A.2 Trip sampling method . . . 5
A.3 Clustering of trips . . . 5
A.4 Addition of the home base and charge station . . . 6
A.5 Characterization of driving conditions . . . 6
A.6 Characterization of relocation behavior . . . 7
A.7 Graph reduction . . . 7
B Vehicle technology parameters . . . 8
C Cost data . . . 8
C.1 Capital cost modeling . . . 8
C.2 Gasoline and electricity prices . . . 8
C.3 Maintenance costs . . . 8
D Damages data . . . 8
D.1 Damage models . . . 8
D.2 Emissions from gasoline combustion and upstream processes . . . 8
D.3 Grid emissions . . . 8
D.4 Manufacturing, recycling, and disposal damages . . . 8

4 Optimization **8**
A Relevant optimization literature . . . 8
B Detailed optimization methods . . . 9
B.1 Formulation . . . 9
B.2 Heuristics . . . 9

Supporting Information Text

Section 1 provides further analysis of the base case results that were described in the main text and lists some limitations of the modeling results described in this study. Section 2 contains additional result sets for a battery of additional test cases, varying input parameters or modeling assumptions. Section 3 summarizes the data and modeling assumptions that were used to determine inputs for each analysis region's base case. Section 4 describes prior optimization literature relevant to this study's optimization model and describes in detail the mathematical formulation and heuristics used.

1. Additional discussion of results

A. Base case results analysis. To accompany the base case results presented in the main text, Figures S1 and S2 show the breakdown of private and external costs, respectively, into subcomponents. These subcomponent values each depend largely on the powertrains that are in use; BEVs and HEVs tend to cost more to purchase but have lower maintenance and energy costs.

Figure S1 shows that vehicles make up the largest portion of private costs, an effect made more pronounced when a Pigovian tax induces the fleet to purchase HEVs and BEVs despite their higher sticker price. In cities where a fleet shifts very heavily towards BEVs, such as Los Angeles under a Pigovian tax, private gasoline cost reductions are not offset by equally large private electricity cost increases.

Figure S2 shows that when there is no Pigovian tax, tailpipe emissions are the largest component of external costs (except in Los Angeles, where BEVs make up the majority of the fleet). Across cities, whether manufacture, refinery, or grid emissions make up a larger portion of emissions depends on which powertrain dominates the fleet mix. A Pigovian tax, by consistently encouraging the purchase of additional BEVs, leads to higher manufacturing emissions; however, those are consistently smaller than both refining and tailpipe emissions reductions.

B. Model limitations. This study's results should be interpreted within the context of its limitations, which result from a combination of modeling choices and computational challenges. Some key limitations include:

- Our implementation of a Pigovian tax is entirely applied to one sub-sector of the economy. This means that our model does not take into consideration economy-wide effects that a Pigovian tax on air emissions externalities would have (such as shifting the grid toward a cleaner mix, reducing damages from charging BEVs).
- We only consider externalities arising from air emissions and ignore the costs of congestion and accidents, assuming they are comparable across vehicle technologies. However, if we were to include them, it may change the balance of capital versus operating costs and could potentially have an effect on the optimal technology mix and fleet operations.
- Our approach to building representative daily samples of demand relies on random sampling from groupings of similar days (e.g., winter weekdays). Because of this sampling approach, our model may tend to "smooth out" bursts of demand that may happen on specific days, e.g. when there is unusually high congestion.
- The assumption of perfect demand information across the modeling period metrics of interest, such as deadheading ratio, overly optimistic. Similarly, assuming perfect demand and cost information enables the fleet to schedule BEV charging in a lower-cost manner than would be possible for a fleet that needed to be robust to uncertainty.
- This model does not consider how differences in climate across analysis regions (or across seasons within a given analysis region) may impact each powertrain's temperature-dependent efficiency or how much air conditioning load the vehicle is burdened with. Prior studies have shown that regional climate differences can influence which powertrain is optimal in a given area. (1)
- Due to computational challenges in optimizing the operations of plug-in hybrid electric vehicles (PHEVs), which can operate in either charge-depleting or charge-maintaining mode, PHEVs were excluded from this model. Given their flexibility to only use the internal combustion engine when the battery is too low, it is easy to imagine that PHEVs may play a role in the optimal fleet mix.
- For reasons of computational tractability, this model includes only one battery charging station, which also serves as the depot where vehicles must originate from and end their tours. This limitation would bias results away from BEV usage if demand were uncertain or if charging stations had finite capacity. However, the fact that all-BEV fleets show only a very slight increase in total VDT implies that this limitation likely has little impact on this study's findings, given its other limitations. (This issue may be a more severe limitation when using trip data that is spread across a more geographically diffuse area.)
- As alluded to above, this model does not consider capacity constraints at the single charging station. This limitation biases results towards BEVs, since charge scheduling can simply consider when a charging station is on the way to the car's next trip or (in Pigovian tax cases) whether the electric grid's marginal emissions are relatively low.
- We assume that vehicles may idle or park between passenger trips, unless there is a gap greater than 30 minutes. If fleets would otherwise need to cruise a neighborhood while waiting for ride requests, then this assumption may bias results towards CVs (which are less efficient during low-speed, stop-and-go driving than HEVs or BEVs).

- Our method for estimating resale values relies on historical data. This is unlikely to capture the fact that BEV technologies (and customer perceptions) may be evolving at a faster rate than CVs and HEvs, and thus may bias findings against BEVs.

2. Sensitivity cases

A. TNC operating models.

A.1. Fleet configuration. To initially assess how much affect powertrain choice may have on outcomes of interest, we ran three homogeneous fleet cases (all-BEV, all-CV, and all-HEV) for each city alongside the base case (optimal selection of a technology mix). Figure S3 tabulates the results from those test cases. In Austin and New York, an all-BEV fleet is the most expensive option; however, in Los Angeles, due to different fuel and electricity prices, an all-CV fleet is most expensive. In each city, under a Pigovian tax, a mixed fleet's external costs are as low as (or, in the case of Los Angeles, lower than) an all-BEV fleet, but with notably lower private costs.

Across all cities, an all-CV or all-HEV fleet does not alter emissions in a meaningful way under a Pigovian tax and an all-BEV fleet is able to reduce its emissions by 2%-4%.

A.2. Discount rates. In addition to the base case, in which the TNC uses a real discount rate of 7% to discount the future cash flow of vehicle resale, we considered a lower value of 1% and an upper value of 13%. Figure S4 tabulates the results from those test cases. In each city, higher discount rates shift the optimal fleet mix towards those powertrains with lower capital costs (e.g., from BEV to HEV and/or CV). Several trends are notable across these cases: (a) in no case does the fleet use CVs for more than 30% of all VMT, even when they make up the majority of the fleet; (b) some HEVs are used in every case; (c) in only one case (New York City with a 13% discount rate) are no BEVs used.

A.3. Resale value. In addition to the base case, which considered resale value as a discounted cash flow occurring either when the vehicle "aged out" (at 12 years) or reached relatively high mileage (at 170,000 miles), we considered a set of test cases in which the same retirement rules were used but resale value was set to zero. Figure S5 tabulates the results from those test cases. There is relatively little change in outcomes due to this, largely because of the regression functions we use to characterize resale value versus annualized vehicle-distance traveled. Using those regression functions, CVs have the lowest capital costs, but tend to have higher resale values for a given usage level, meaning that removing resale from the model increases their cost more than HEVs (and HEVs more than BEVs for similar reasons). There is a general trend towards using CVs less heavily when resale value is removed, but the effect is small; other effects vary by city.

A.4. Labor costs. In addition to the base case, in which we assume labor costs are unaffected by powertrain choice because they are either zero (i.e., the fleet is automated) or computed as a flat percentage of ride fees, we considered an hourly rate of $12/hour, the current statewide minimum wage in California. We treated this hourly rate as being paid to the driver of each car the entire time the car is being driven or while it is charging, unless the charging activity occurs overnight. These labor costs may be expected to penalize BEVs due to the additional costs paid to drivers when they must make detours and wait for the car to charge. S6 tabulates the results from those test cases. As anticipated, labor costs reduce utilization of BEVs, but the effect varies widely by city: BEV VDT is reduced by 16% (Los Angeles) to 72% (Austin).

B. Battery technology.

B.1. Battery capacity. The 2019 Kia Soul is used as the base case BEV reference model. As a sensitivity case, we considered the 2020 Chevrolet Bolt, which has a significantly larger Lithium-ion battery (66 kWh for the Bolt versus 30 kWh for the Soul) but with a relatively small increase in sticker price ($36,260 for the Bolt versus $33,950 for the Soul). Figure S7 tabulates the results from those test cases. When minimizing private costs, the fleet mix changes little when the BEV model is changed; in Los Angeles and New York, BEVs become slightly more optimal when the Bolt is used instead of the Soul, but in Austin, the effect is reversed (perhaps because there is less opportunity for the differential between gasoline and electricity prices to overcome the Bolt's slightly higher capital costs).

Considering an all-BEV fleet in Austin, the Bolt's larger battery allows the fleet size to be reduced by 12-16% depending on whether a Pigovian tax is imposed. (Only Austin was tested because, in general across most test cases we ran, the optimal fleet mix is nearly as small as possible given the proportion of BEVs used).

B.2. Battery price. To consider the effect technological learning and economies of scale may have on BEV battery prices in the future, we considered a set of test cases in which the purchase price of BEVs was reduced by $5,000. Figure S8 tabulates the results from those test cases.

Aside a price reduction encouraging more BEVs to be used, a notable finding is that when BEVs become cheaper, the impact of a Pigovian tax is reduced. The private cost-minimizing fleet is induced to use more BEVs when they are cheaper, but there are generally some peak load hours in which it still makes sense in most test cases to purchase CVs or HEVs. Because of this, the external cost reductions a tax achieves are lessened from 24% to 7% in Austin, from 34% to 18% in Los Angeles, and from 12% to 2% in New York City.

C. External cost inputs.

C.1. Alternate damage models. In addition to the AP3 reduced complexity model of spatially resolved air emissions damages, we considered two other reduced complexity models: EASIUR and InMAP. Figure S9 tabulates the results from those test cases. In Austin and Los Angeles, differences between damage models are not large enough to drastically change the fleet's technology mix or operations. However, because damage valuations differ, the external cost reductions caused by a Pigovian tax also differ. The ordering of damage models, in terms of damage valuation, varies from city to city.

In New York City, InMAP estimates much higher damages from internal combustion than AP3 or EASIUR do. Because of this, a Pigovian tax would be much higher there according to InMAP (19.5 cents using InMAP versus 12.4 cents using AP3 and 9.5 cents using EASIUR). This higher tax induces a much more drastic shift, in which the fleet converts nearly fully to BEVs; this results in an externality reduction of 81%.

C.2. Social cost of carbon. In addition to the base case value of \$50/tonne of CO_2-equivalent, we considered a case in which those emissions were priced at \$300. S10 tabulates the results from those test cases. In LA, which already uses BEVs heavily, relatively little change occurs due to a higher social cost of carbon. In Austin and New York City, the fleet electrifies substantially more under a higher social cost of carbon, leading to larger externality reductions. In Austin, those externality reductions increase from 17% to 43% under a higher social cost of carbon; in New York City, they increase from 12% to 31%.

3. Data and Modeling Assumptions

A. Trip and relocation data.

A.1. Trip data preprocessing. The RideAustin trip dataset includes some observations that seem highly atypical or even impossible. Certain observations may represent actual unusual trips that do not fall within the scope of our model (e.g., intercity-length trips or trips abandoned by the passenger or driver immediately after their start), and others may represent data collection artifacts of some kind. All of those anomalous trips were removed from the dataset. Only trips meeting all of the following criteria were retained: no missing data for start or end latitude or longitude, distance traveled below 50 miles, duration exclusively between 1.5 and 120 minutes, and speed (as implied by distance traveled and duration) exclusively between 1 and 100 miles per hour. These filters reduced the dataset from 1.49 million to 1.46 million trips.

A.2. Trip sampling method. The RideAustin dataset was sampled down to a set of a little more than 5,000 representative trips to make the joint optimization of a full year's purchase and routing decisions tractable. Manual inspection showed that typical trip volumes by hour of day tended to follow somewhat similar patterns across weekdays, with different patterns on weekends and with peak demand usually occurring on Friday and Saturday nights. Based on these trends, and on seasonally variable input data for marginal grid emissions, sample segments were defined using:

- Season of year: Summer (May-September), Winter (November-March), and Summer/Fall (April and October).
- Weekday (Sunday-Thursday) versus Weekend (Friday-Saturday).
- South by Southwest (March 10-18, 2017): due to drastically increased demand during the South by Southwest music festival, it was separately sampled as one weekend cluster and one weekday cluster.

To ensure that each representative days' sampled volumes and hourly demand profiles were similar to those in the full dataset, each combination of segment and hour of day was sampled proportionally to its average volume across all available days in the dataset. A sample size of 5,000 trips was targeted, but since the number to sample from each segment-hour combination was rounded up to the nearest trip, the actual sample size was 5,089 trips. These segments were placed one after the other in a combined trip dataset sample used in the optimization, with each segment running for 24 hours beginning at 6 A.M. (the lowest-demand hour of the day on average). These segments and all of their demand must be met, which influences the fleet's size and mix, but to compute annualized costs and vehicle-distance traveled, each arc's costs and distances are weighted by the relevant category's number of days per average calendar year. Those multiplication factors are provided in Table S6.

Over the course of the RideAustin dataset, daily volumes show a long-term trend of increasing demand (a trend that is likely shared with larger TNCs' long-term demand). As illustrated in Figure S11, this means that segments such as Summer Weekday and Summer Weekend have fewer trips in the optimization model than other seasons of the year. Rather than attempt to distinguish this long-term growth trend from seasonality and remove it, we treat the RideAustin dataset as a steady-state year.

A.3. Clustering of trips. As shown in Figure S12, the majority of RideAustin trips begin or end in the downtown core, with a handful of other locations (e.g., Austin-Bergstrom International Airport and the University of Texas-Austin campus) having concentrations of high demand. Most trips are ultimately represented in our model exactly as found in the data, but we group certain similar trips together—largely during high-demand periods in time-space (for which vehicle routing choices are are more difficult to optimize)—to reduce the number of unique arcs needed to model trip requests. A variant of Ward's method of hierarchical agglomerative clustering on non-squared distances (implemented via the `hclust` command in the `stats` R package with option "`method=ward.D`") groups all starting and ending locations into a smaller set of neighborhoods. Figure S13 shows the resulting neighborhood clusters and the volume of trips beginning or ending in each one. Since RideAustin data report trip start/end times to the nearest second, geographic clustering must be combined with rounding of timestamps in order to reduce the size of the trip network being optimized over. Trip start and end times were rounded to the nearest three minutes.

After this rounding of time and location, when individual passenger trip requests do not share a combination of starting timestamp, ending timestamp, starting cluster, and ending cluster with other trips, their actual start/end location and time were preserved at full fidelity. For passenger trip requests that were grouped into a consolidated trip arc (most of which begin or end in the geographically compact downtown clusters during a period of high demand), the consolidated trip arc uses the average start/end location and time of the grouped trips, resulting in some loss of fidelity to actual trip starting and/or ending points. The demand along that arc is the number of trip requests combined along it. Despite being combined into one arc in the optimization model, these trip requests cannot be pooled into one car; rather, multiple cars must traverse that trip arc to satisfy demand.

A value of 25 was chosen for the number of neighborhood clusters to compute. For computational tractability, clusters were computed from a random sample of 20,000 trips, and as a second step, k-nearest neighbors assigned each unique latitude-longitude pair in the RideAustin dataset to one of the 25 clusters. Because k-nearest neighbors uses "votes" of multiple nearby neighbors to classify locations, no latitude-longitude pairs were ever assigned to two extremely low-volume outlying clusters, reducing the trip dataset to 23 clusters.

The number of clusters was chosen to achieve a desired reduction in optimization problem size, and it was then validated on a random sample of 5,000 trips (for computational tractability of validation methods) using three different methods:

- Figure S14 uses the elbow method to validate number of clusters. It shows (on a logarithmic scale) that the within-cluster sum of squares, a measure of similarity within clusters, continues to decline as the number of clusters increases, but appears to reach diminishing returns prior to the chosen number of clusters. Unless plotted on a logarithmic scale, this metric appears relatively flat for values above 5 clusters.

- Figure S15 uses the average silhouette method. It shows that average cluster silhouette width, a measure of how well-clustered observations are using the average distance between clusters, is relatively flat for values above 10 clusters.

- Figure S16 uses the gap statistic method. It shows that the gap statistic, a measure of how different the clustered data's structure is from a random distribution of points, is relatively flat for values above 20 clusters.

Many trips were not consolidated into grouped arcs with other trips; Table S7 shows the degree to which clustering and rounding of timestamps consolidated trips. These steps preserve full geographic and temporal fidelity for 4,370 out of 5,089 trips in the sample used for optimization model runs.

A.4. Addition of the home base and charge station. We assume one home base from which vehicles depart to serve trips and to which they must return. This home base also serves as the fleet's single charging station, with unlimited capacity, and we use k-means clustering with k=1 to site the station at the centroid of all unique trip origin-destination pairs. This assumption is optimistic for battery vehicles in that the charge station has infinite capacity; on the other hand, it is pessimistic for them in that there are not multiple charge sites optimally scattered along high-volume roadways. Across virtually all scenarios, usage of battery vehicles leads to a negligible increase in total vehicle-distance traveled, suggesting that multiple charge sites would not significantly change our metrics of interest (at least not for a full-information fleet serving RideAustin trips).

Modeling a home base and charge station requires adding relocation arcs to travel to/from the base and trips. The model's flow preservation constraints mean that each vehicle's routes can originate and terminate only at the base (and must do so). The model uses 15-minute timeslots during which each vehicle can linger at the base (optionally charging their batteries). We add to the optimization an optional relocation arc from the base to each trip's start point, and from each trip's end point back to the base. These arcs assume vehicles would stay at the base as long as possible, leaving exactly when needed to arrive on time to the next trip and returning once each trip concludes. They depart from the latest viable 15-minute timeslot and connect to the earliest viable one. When BEVs would, in actuality, arrive a few minutes prior to the start of that 15-minute timeslot (or could leave a few minutes after its end), they are allowed the option of using the additional time to charge.

A.5. Characterization of driving conditions. Each powertrain's efficiency varies based on whether the driving style more closely resembles highway or city driving. While it is impossible to know from the available fields in RideAustin data what the predominant driving style of any given trip was, we use the available fields to create a rough estimation of driving style that, while imperfect for any given trip, allows the model to consider how driving style may influence tradeoffs among technologies.

Using data fields for trip duration and distance traveled, we compute the average speed in miles per hour for each trip. We compare this average speed to the average speeds of the EPA city and highway test cycles (48.3 miles per hour for highway driving and 21.2 miles per hour for city driving). The efficiency on a given trip, e_{T}, is estimated for each powertrain based on the trip's average speed, s_{T}, using linear interpolation:

$$e_{trip} = \left\{ \begin{array}{ll} e^{\mathrm{CITY}} & \text{for } s_{\mathrm{T}} < 21.2 \\ \frac{(s_{\mathrm{T}}-21.2)}{48.3-21.2} e^{\mathrm{HWY}} + (1 - \frac{(s_{\mathrm{T}}-21.2)}{48.3-21.2}) e^{\mathrm{CITY}} & \text{for } 21.2 \leq s_{\mathrm{T}} \leq 48.3 \\ e^{\mathrm{HWY}} & \text{for } s_{\mathrm{T}} > 48.3 \end{array} \right\} \quad [1]$$

where e^{CITY} and e^{HWY} are the powertrain's city and highway efficiency, respectively, in gallons per mile. We truncate efficiencies at EPA city and highway ratings to avoid extrapolation. Few trips (1.7% of those sampled from) have an average speed above the highway-cycle average speed, and there is little easily applicable data for how the assumed efficiency should be modified for very high-speed driving. Most trips (54.6% of those sampled) have an average speed below the city-cycle average speed

of 21.2 miles per hour. For those trips, this approach can be interpreted as assuming that idling in stop-and-go traffic (one reason trips may have low average speeds) incurs no additional fuel consumption, which would be most consistent with a fleet whose combustion engine vehicles include the start-stop technologies that reduce idling in many new cars. To the extent that slower-speed trips penalize efficiency of internal combustion engine vehicles moreso than hybrid and electric vehicles in practice, this factor may bias our results somewhat in favor of conventional vehicles.

A.6. Characterization of relocation behavior. RideAustin data does not include data on how vehicles are driven while waiting for their next passenger trip request, nor is there data to describe what route they take or how efficiently they are driven to the next trip's starting point. These shortcomings apply not only to trip-to-trip relocations, but also to relocations to/from a trip and the home base where cars park and batteries recharge. To estimate vehicle efficiency for relocation arcs, we use similar trips to estimate time needed to drive from the prior trip's end to the next trip's start (as the product of estimated driving distance and speed) and assume any additional time is spent idling.

We use k-nearest neighbors regression to estimate the driving speed and distance from each trip's ending point in space-time to the next trip's starting point. Since typical traffic conditions (and optimal routes) vary over time, we subset RideAustin data by day of week and hour of day, resulting in 168 subsets of data. For each subset of data, separate k-nearest neighbors regressions for speed and distance estimate each possible trip-to-trip relocation arc's distance and speed as the average of the nearest k trips. As a distance metric to determine which k trips are "nearest" to each relocation, we use the Cartesian distance between the relocation's and the trip's starting latitude-longitude plus the Cartesian distance between the relocation's and trip's endpoint latitude-longitude.

The number of nearby neighbors to use for estimation, k, was determined separately for each of the 168 subsets using 10-fold cross-validation as follows:

1. Label passenger trip data as a randomly chosen 70% (training set) and 30% (validation set) of observations.
2. For each of the 168 weekday-hour subsets of data:
   (a) Divide its training observations into 10 equally sized folds of data.
   (b) Estimate the speed and distance of trips within each fold by training k-nearest neighbors on the other nine folds, sweeping k from 3 to 100.
   (c) Compute the root-mean-square error (RMSE) resulting from each k for each fold, where error is defined as the difference between actual and estimated speed (or distance).
   (d) Average each k's RMSE across the ten folds and use the RMSE-minimizing k for that weekday-hour subset of data.

Having determined an optimal k separately for each of the 168 subsets of data, we use the remaining 30% of data (the validation set) to evaluate the performance of k-nearest neighbors. Scatterplots of true versus estimated speed and distance are shown in Figures S18 and S17. Evaluated on the validation set, this method yields an RMSE of 5.71 miles per hour and a mean error of -0.53 miles per hour for speed. For distance, it yields an RMSE of 1.88 miles and a mean error of -0.22. As one comparison benchmark, using the Manhattan distance (relocation arc origin-destination's x-distance plus y-distance) as a distance estimator would have yielded a higher RMSE of 2.13 miles and a mean error of -0.76 (meaning that the bias towards shortening distance estimates would be more than tripled if we used Manhattan distance).

Each relocation's estimated travel time was computed as the product of estimated speed and estimated distance. For relocations with actual time available equal to the estimated travel time or less than 30 minutes longer, we assume the vehicle idles in place and that idling incurs no fuel consumption. This is consistent with our assumption that the fleet owner-operator has perfect information about demand, and it further assumes that there is curb space available at each trip's dropoff point (or another point along the route from dropoff to the next trip's pickup) for the vehicle to idle (without fuel consumption) as long as needed.

A.7. Graph reduction. As an initial step to reduce the size of the graph of arcs over which we optimize, estimated relocation travel time is used to exclude certain relocations from the model. Relocations meeting any of the following conditions are removed from consideration:

- An estimated relocation travel time greater than the actual time available indicates the relocation is impossible, so that relocation is removed. This applies to both trip-to-trip relocations and trip-to-base/base-to-trip relocations.
- An estimated relocation travel time more than 30 minutes less than the actual time between prior trip end and next trip start implies the vehicle would need to idle in place for 30 minutes or more. Those relocations are removed. This step applies only to trip-to-trip relocations, which may still occur in the optimization, but would require the vehicle to pass through the centrally-located home base where vehicles are able to park and charge.

This step removes trip-to-trip relocations that would require lengthy curbside idling and relocations of any kind that are likely to be infeasible.

As an additional step, a heuristic is used to remove trip-to-trip relocations that are relatively unlikely to be used in an optimal solution. That heuristic (Algorithm 1) is described in Detailed Optimization Methods below.

B. Vehicle technology parameters. Table S8 compares technological parameter assumptions across powertrains. These parameters were taken from the manufacturer website and from fueleconomy.gov.

C. Cost data.

C.1. Capital cost modeling. The method for discounting future cash flows from the resale of a vehicle were described in the main text. In order to determine depreciation factors, data was gathered from Kelley Blue Book's website across two ranges: (a) vehicles aged 12 years with mileage swept from 30,000 to 170,000 miles, and (b) vehicles at 170,000 miles with age swept from 0-12 years. The first set of values (a) corresponds to vehicles retired from TNC usage due to age and obsolescence; the second set of values (b) corresponds to vehicles retired from TNC usage due to high mileage and wear and tear. Separate regressions were run on these two sets of values. (There was no meaningful difference in resale value at 12 years for vehicles with fewer than 30,000 miles, so resale value was treated as usage-invariant within that range using the value taken from Kelley Blue Book). There were no BEVs with resale value data for a full set of 12 years, so the Nissan Leaf was chosen as a midmarket BEV with a relatively long history. For HEVs and BEVs, the Toyota Camry was chosen as a relatively midmarket reference vehicle. Table S9 shows the resale value factors each vehicle's purchase price (as given in Table S8) is multiplied by and the factors' dependence on annual vehicle-miles traveled.

Once each powertrain's regression function was determined, the nonlinear function was approximated via a set of several linear floor constraints on cost versus annualized VDT. Figure S20 illustrates how those linear floor constraints approximate the computed cost for HEVs.

C.2. Gasoline and electricity prices. Table S10 summarizes the gasoline and electricity prices used for each analysis region. These values were taken from EIA data for calendar year 2018 (2, 3).

C.3. Maintenance costs. Table S11 summarizes the per-mile maintenance costs that were used for each powertrain and the prior work they are derived from.

D. Damages data.

D.1. Damage models. To quantify a monetary value representing human health damages resulting from a unit of air emissions for a given pollutant, it is necessary to answer several questions: where the pollutant is emitted, what base levels of pollutants are in that area, where the pollutant may travel in the air, and what vulnerable populations are located in those areas. Estimation of emissions locations are described in the following section; the other questions are addressed by reduced complexity models that estimate health impacts resulting from air pollution. In contrast to the estimates generated by complex chemical transport-based air pollution models, reduced complexity models generate estimates at an acceptable level of accuracy while enabling estimates to be found for large numbers of scenarios quickly. For the base case, we use the AP3 model to compute and monetize estimated health damages associated with these emissions.(4) We consider additional sensitivity cases using EASIUR and InMAP.(5, 6)

D.2. Emissions from gasoline combustion and upstream processes. GREET was used to characterize tailpipe emissions from internal combustion, as well as "well-to-pump" emissions from gasoline production. Table S12 summarizes the values that were taken from GREET (identical emissions per gallon were used for conventional and hybrid electric vehicles). For each analysis region, tailpipe emissions are assumed to occur in the most highly populated county; Table S12 lists the county FIPS codes used for tailpipe emissions.

All upstream emissions from gasoline production processes are assumed to occur at the fuel refineries in the area surrounding each analysis region (the analysis region's state and, in the case of New York City, nearby states). EIA refining capacity data for individual refineries for 2018 was used to create a list of relevant refineries for each analysis region and their capacity (reported as atmospheric crude distillation capacity in barrels per calendar day). (**?**) Each refinery was then manually attributed to a specific county, and capacity was aggregated at the county label. To site upstream emissions, each of these counties was used and weighted by its operational refining capacity. Table S13 lists the county FIPS codes used upstream fuel refinery emissions for each analysis region.

D.3. Grid emissions. To model grid emissions, we use marginal grid emissions factors estimates generated by an existing Siler-Evans et al. (2012) methodology and downloaded from a Center for Climate and Energy Decisionmaking web portal. (7, 8) This methodology uses historical generation emissions and demand data to build a regression estimating typical marginal grid emissions and marginal grid health damages by hour of day and season of year. Table S14 lists the Emissions & Generation Resource Integrated Database (eGRID) subregion for which marginal emissions factors are calculated.

D.4. Manufacturing, recycling, and disposal damages. To model air emissions damages arising from vehicle manufacturing, recycling, and disposal, we use Argonne National Laboratory's GREET model. Table S15 shows the d

4. Optimization

A. Relevant optimization literature. Within the body of vehicle routing methodology studies, the term "Green Vehicle Routing Problem" (GVRP) denotes formulations involving range constraints. GVRP formulations are often motivated by either battery

capacity limits or alternative fuels' sparse refueling infrastructure. The GVRP with recharge or refueling stations was introduced by Erdogan et al. in 2012 (9). Extensions have added nuances including the option to partially recharge at stations of different battery charging speeds (10), complex energy consumption and charging models (11**?** , 12), time-variant congestion and emissions (13), and simultaneous charger siting and vehicle routing (14). GVRP studies almost always use a combination of mathematical optimization and heuristic solution generation, and some papers have focused on improving tractability of existing GVRP variants via reformulating the problem or improving heuristics (15–17). GVRP studies typically present exact solutions for no more than 10-150 trips or heuristic solutions for as many as 500 trips. They are generally more concerned with testing computational performance than with policy outcomes such as emissions.

Only four GVRP optimization studies were found that have considered routing a mixed-technology fleet including electric vehicles (18–21). Three of those four jointly optimized the fleet technology mix and vehicle routing for either passenger travel or goods delivery, and only one of those three considered both electric and combustion powertrains: Sassi et al. (2014) models a different problem than this study, the delivery of goods to a set of customers in any order over the course of a day (19). None of these four studies considered external costs. The four studies found heuristic solutions for up to 200 passenger trips or, in Sassi et al., 500 locations.

B. Detailed optimization methods.

B.1. Formulation. Our `FullMILP` formulation, summarized in Figure S4, finds the cost-minimizing fleet technology mix and assignment of vehicles to trip arcs where the set of decision variables $\mathcal{X}$, summarized in Table S5, includes the number of vehicles n_k of each type k purchased, assignments $a_{k,i,j}$ of vehicles k to arcs (i,j), charge level $q_{k,t}$ and energy charged from the grid $\Delta q^{\text{CHG}}_{k,t}$ for each vehicle k at each discrete time point t, and total annualized acquisition cost κ_k for each vehicle type (determined by vehicle utilization levels) for all vehicle types $k \in \mathcal{K}$, arcs $(i,j) \in \mathcal{A}$, and times $t \in \mathcal{T}$. Notation is summarized in Table S5. The set of vehicle types $k \in \mathcal{K}$ indexes individual vehicles for BEVs but groups vehicles into types for conventional and hybrid electric vehicles for efficiency because the model only needs to track charge level of BEVs.

The objective function, Eq.(5a), sums the vehicle purchase costs, dispatch costs, and time-dependent battery charging costs. In all test cases, the cost terms κ_k (vehicle purchase), $c_{k,i,j}$ (internal combustion and per-mile maintenance), and c_t (battery charging) include the costs described in Section **??**; in some cases, they also include the emissions externality costs described in Section **??**.

At the core of `FullMILP` are equations that are standard for many vehicle routing problems. Constraint 5b ensures preservation of flow for each vehicle through the network (forcing vehicles to return to the depot after serving trips), Constraint 5c requires that all passenger trips be satisfied, and Constraint 5d requires that a vehicle must be purchased to be dispatched. The remainder of the formulation is customized for our case.

Constraints **??**–5e model capital cost. Constraint **??** limits annualized distance per vehicle (default $d = 170,000$ miles); it is included to guarantee consistency with capital cost assumptions, but it is not binding in any scenarios. For all vehicles, Constraint 5e uses a set of linear constraints Ω to define a convex piecewise linear cost floor representing the sum of annualized vehicle costs, including salvage value, which is a function of vehicle assignment, and, in relevant cases, internalized externality costs. We discuss this aspect of our formulation in more detail in the SI.

Constraints 5g-5i manage BEV charge level. Constraint 5g applies to timesteps at which regular interval charging timesteps begin (default 15-minute intervals), defines charger usage, and tracks charge level changes. Constraint 5h applies to all other timesteps, at which there is no charging option, so that charge level is fully determined by traversed arcs' energy requirements. Constraint 5i enforces bounds of BEV charge levels. The implied amount of electricity purchased from the grid is quantified for the objective function in the "where" statement as the change in charge unexplained by travel.

For CVs and HEVs, `FullMILP` assumes that refueling can happen quickly such that tracking fuel level is unnecessary and `FullMILP` need not separately index each car. Refueling scheduling is not modeled, but CV and HEV dispatch is otherwise representative of discrete vehicles. Constraint 5j forces BEV assignment to be binary, but Constraint 5k only requires that CV and HEV assignments be a nonnegative integer.

B.2. Heuristics. Because `FullMIP` cannot be solved at meaningful scale using commercial branch-and-bound MIP solvers alone, this study uses two novel heuristics and one widely used heuristic to generate solutions for `FullMIP`. The tractability challenge comes from modeling the state of charge of BEVs. Because energy level of CVs and HEVs is not tracked, they need not be individually indexed (since an aggregate fleet dispatch can be easily decomposed into individual vehicles post-hoc) and they scale to reasonable problem sizes more easily.

The first heuristic step, `MCF_VaryingFleetSize`, reduces the set of arcs $\mathcal{A}$ to make the optimization tractable. It is inspired by a graph reduction heuristic developed by Bertsimas et al. (2017), but redesigns that heuristic to suit a different application: approximating behavior of a mixed-technology fleet in which certain arcs (e.g., city-style versus highway-style arcs) are better suited to certain powertrains (22). It relies on a variant of `MCF`, `MCF_CarLimit`, which has a limited fleet size; to dispatch this smaller fleet, it removes Constraint **??**'s mandate that every trip must be met and adds Constraint 2 to limit the number of vehicles dispatched to a fleet size F:

$$\sum_{i=r, j \in \mathcal{V} \setminus \{r\}} a_{i,j} \leq F \qquad [2]$$

`MCF_CarLimit` runs quickly but does not model individual vehicle routes.
`MCF_VaryingFleetSize` repeatedly runs `MCF_CarLimit`, sweeping the car limit from large to small car limits (i.e., fleet sizes).

For each of these fleet sizes, `MCF_CarLimit` is run repeatedly until all trips are served, randomizing arc costs at each iteration to represent a fleet of CVs, HEVs, or BEVs. The union of all arcs ever used across all of these fleet sizes and technology types is kept for subsequent heuristic steps.

Algorithm 1 `MinCostFlow_VaryingFleetSize` heuristic to reduce problem size using the `MCF_CarLimit` formulation.

Data: Graph G; maximum allowed number of vehicles, max.fleet
Result: G_{small}, a version of G with certain relocation arcs removed
Initialize $G_{small} \leftarrow$ all arcs from G except for trip-to-trip relocations
for *i from floor(log₂(max.fleet)) to floor(log₂(min.fleet))* **do**
  Initialize $G' \leftarrow G$
  while *unassigned trips remain in* G' **do**
    Randomly assign *technology.type* ← CV, HEV, or BEV
    Assign arc costs to G using per-arc energy consumption of *technology.type*
    Run `MCF_CarLimit` on G with a fleet size F=2^i
    Remove trips from G that were traversed by `MCF_CarLimit`
    Add trip-to-trip relocation arcs traversed by `MCF_CarLimit` to G_{small}
  end
end
return G_{small}

After the graph is reduced for tractability, a second novel heuristic, `ShrinkingBattery`, generates an initial solution. Rather than optimizing all individually indexed BEVs at once in `FullMIP`, it runs `FullMIP` with only one specific BEV at a time alongside the CV and HEV portions of the fleet. To prevent avoid "greedy" heuristic behavior in which each BEV captures as much demand as possible (rather than acting as one member of a larger fleet), an aggregated representation of a BEV fleet is also included, represented as one vehicle *big.ev* with integer assignments ($a_{big.ev,i,j} \in \mathbb{Z}_+$). To ensure this aggregated BEV itself does not capture an unrealistic amount of demand—since it is less impeded by range constraints than a realistic BEV—Constraint 3 is added to `FullMIP`. It ensures the individually indexed BEV traverses a "longer" tour of trips than the per-car average tour of the aggregated BEV, measured via total energy lost on trip arcs. big.ev.size denotes the number of cars represented by the aggregated BEV.

$$big.ev.size \times \sum_{(i,j)\in trips} a_{small.ev,i,j} - \sum_{(i,j)\in trips} a_{big.ev,i,j} \geq 0 \quad [3]$$

To ensure that the battery capacity of the aggregated BEV is determined by the quantity (in aggregate) of BEVs purchased, Constraint 4 is added to `FullMIP`:

$$l_{big.ev,t+1} \leq p_{large.ev} b_{small.ev} \quad \forall t \in \mathcal{T}_+ \quad [4]$$

At each step, the single BEV's assignments are added to the current solution; the aggregation shrinks and its operations iteratively come closer to approximating an individual vehicle.

Algorithm 2 `ShrinkingBattery` heuristic to generate BEV routes using the `FullMIP` formulation.

Data: Graph G; maximum allowed vehicles of each type, $max.CV$, $max.HEV$, & $max.BEV$
Result: *bev.arcs*, the assignment of specific BEVs and arcs in G
Initialize *remaining.bev* ← *max.ev*, $G' \leftarrow G$, and vehicle.arcs ← NULL **while** *remaining.bev > 0 and remaining.graph is not empty* **do**
  Retrieve a solution for `FullMIP` on *remaining.graph* with 4 vehicle indices: (1) $max.CV$ CVs, (2) $max.HEV$ HEVs, (3) a single BEV, and (4) an index representing an aggregated BEV fleet of size *remaining.bev* Add the single BEV's assigned arcs to *bev.arcs* Remove the single BEV's fulfilled trips from G';
  remaining.bev ← *remaining.bev* − 1
end
Add arcs traversed by CVs in the most current run of `FullMIP` to *bev.arcs* **return** *bev.arcs*

After an initial solution is found, the `RuinAndRecreate` heuristic is used to randomly re-optimize smaller pieces of the full set of BEV vehicle routes, incrementally improving the quality of the solution. This general concept is widely used in vehicle routing literature (23–25).

Algorithm 3 `RuinAndRecreate` heuristic to improve BEV routes of a solution to `FullMIP`.

Data: An instance of `FullMIP`, M; a candidate solution, S; desired iterations *limit*
Result: S', an improved solution to M
Initialize $S' \leftarrow S$
for *i from 1 to limit* **do**
  Starting point solution of $M \leftarrow S'$;
  Randomly select a subset of two BEVs, *subset*, to optimize;
  Upper & lower bound of M's vehicle-to-arc-assignments not in *subset* $\leftarrow S'$;
  Upper & lower bound of M's vehicle-to-arc-assignments in *subset* $\leftarrow (0, 1)$
  $S' \leftarrow$ solution of `FullMIP` on M with updated bounds, optimizing only *subset*
end
return S'

After `RuinAndRecreate` improves the candidate solution, `FullMIP` is run a final time to establish a lower bound on cost and (in some cases) to find minor improvements to the candidate solution.

Fig. S1. Summary of optimized private cost components for three analysis regions, with and without a Pigovian tax on air emissions. Except where otherwise stated, scenarios reflect a 7% real private firm discount rate, no labor costs, $33,950 BEV price (2019 Kia Soul), $50/ton CO_2 externality price (for Pigovian tax cases), the AP3 damage model, $9.41 million (2018) value of statistical life, and the Pope et al. (2019) concentration-response function. RideAustin trip data is used in all three cases. Region-specific values include: gasoline prices, electricity prices, marginal grid emissions estimates, and criteria pollutant externality cost estimates.

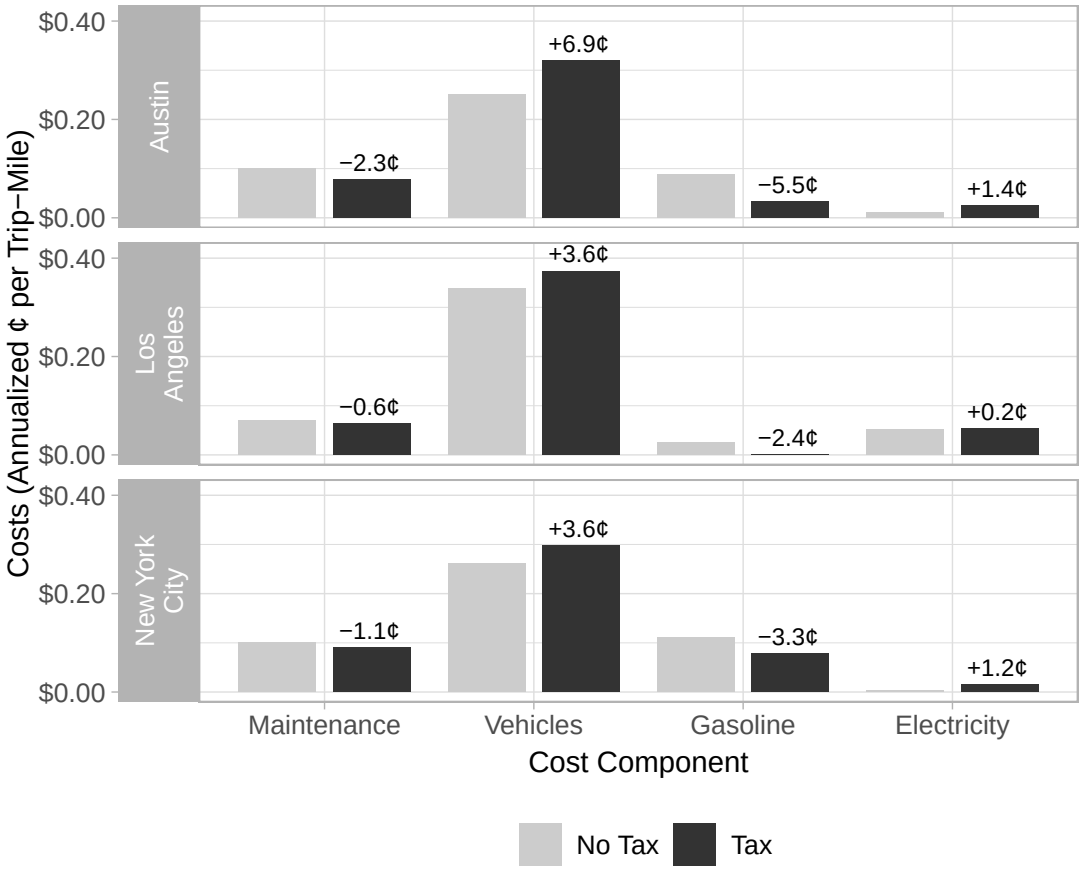


Fig. S2. Summary of optimized external cost components for three analysis regions, with and without a Pigovian tax on air emissions. Except where otherwise stated, scenarios reflect a 7% real private firm discount rate, no labor costs, $33,950 BEV price (2019 Kia Soul), $50/ton CO_2 externality price (for Pigovian tax cases), the AP3 damage model, $9.41 million (2018) value of statistical life, and the Pope et al. (2019) concentration-response function. RideAustin trip data is used in all three cases. Region-specific values include: gasoline prices, electricity prices, marginal grid emissions estimates, and criteria pollutant externality cost estimates.

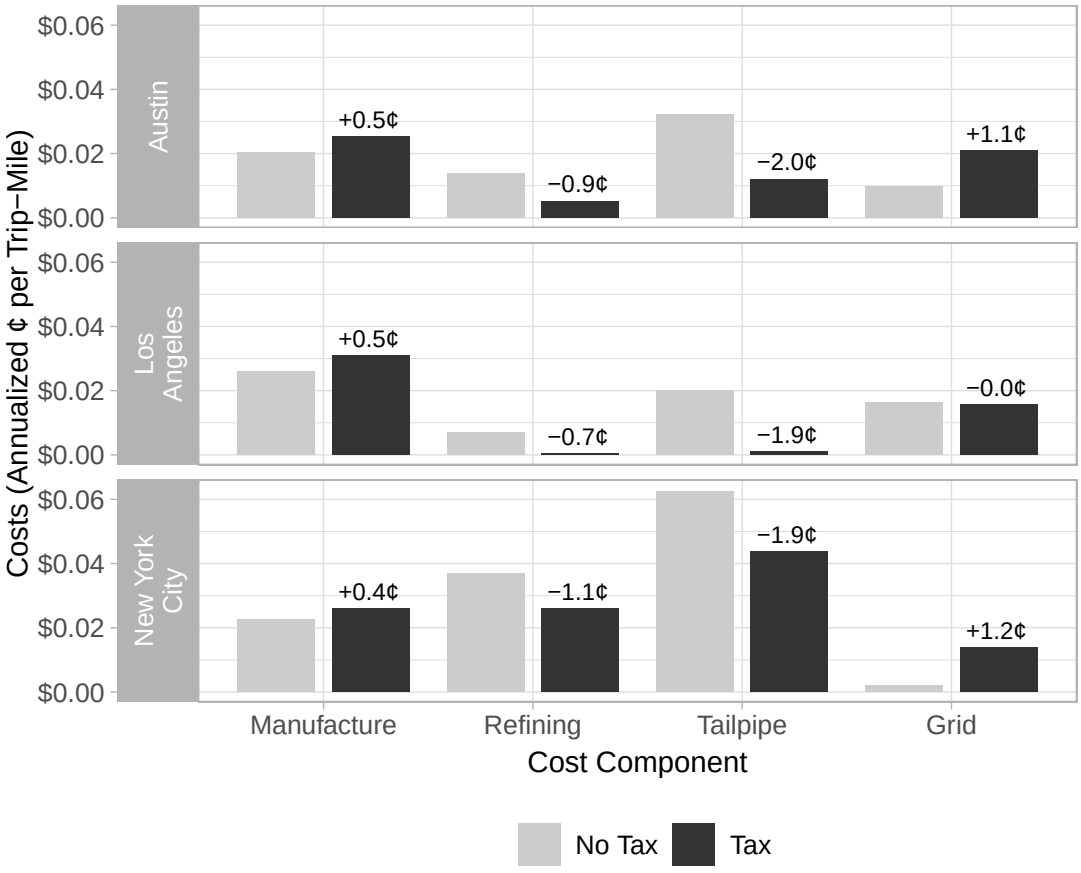


Fig. S3. Summary of results for different fleet configurations across three analysis regions. Except where otherwise stated, scenarios reflect a 7% real private firm discount rate, no labor costs, $33,950 BEV price (2019 Kia Soul), $50/ton CO_2 externality price (for Pigovian tax cases), the AP3 damage model, $9.41 million (2018) value of statistical life, and the Pope et al. (2019) concentration-response function. RideAustin trip data is used in all three cases. Region-specific values include: gasoline prices, electricity prices, marginal grid emissions estimates, and criteria pollutant externality cost estimates.

	Powertrains Available	Externalities Internalized?	Vehicles Purchased	Vehicle-Distance Traveled (1K miles/year)	Private Costs (¢/trip-mile)		External Costs (¢/trip-mile)		Social Costs (¢/trip-mile)	
					Value	Change	Value	Change	Value	Change
Austin	All		24 7 5	340 601 346	45.0	1.3%	7.6	-16.7%	52.6	-1.3%
		X	1 16 19	14 395 881	45.6		6.3		51.9	
	BEV Only		42	1306	48.5	0.7%	6.5	-2.9%	55.0	0.3%
		X	43	1307	48.8		6.3		55.1	
	ICE Only		36	1289	47.4	0.0%	10.0	0.0%	57.4	0.0%
		X	36	1289	47.4		10.0		57.4	
	HEV Only		36	1289	45.4	0.0%	7.4	0.0%	52.9	0.0%
		X	36	1289	45.4		7.4		52.9	
Los Angeles	All		14 22	198 1092	48.5	1.7%	7.0	-30.8%	55.4	-2.3%
		X	2 2 35	2 8 1286	49.3		4.8		54.1	
	BEV Only		42	1308	50.3	0.1%	5.0	-1.9%	55.3	-0.1%
		X	42	1309	50.3		4.9		55.2	
	ICE Only		36	1289	56.7	0.0%	28.5	0.0%	85.2	0.0%
		X	36	1289	56.7		28.5		85.2	
	HEV Only		36	1289	51.5	0.0%	19.7	0.0%	71.2	0.0%
		X	36	1289	51.5		19.7		71.2	
New York City	All		5 22 9	71 1012 204	47.7	0.9%	12.4	-11.5%	60.1	-1.7%
		X	14 22	782 506	48.1		11.0		59.1	
	BEV Only		42	1310	52.9	0.5%	11.4	-3.8%	64.3	-0.3%
		X	43	1308	53.2		10.9		64.1	
	ICE Only		36	1289	51.4	0.0%	19.2	0.0%	70.7	0.0%
		X	36	1289	51.4		19.2		70.7	
	HEV Only		36	1289	48.1	0.0%	13.5	0.0%	61.6	0.0%
		X	36	1289	48.1		13.5		61.6	

Conventional Vehicle (CV) | Hybrid Electric Vehicle (HEV) | Battery Electric Vehicle (BEV)

Fig. S4. Summary of results for three different private firm discount rates across three analysis regions. Except where otherwise stated, scenarios reflect no labor costs, $33,950 BEV price (2019 Kia Soul), $50/ton CO_2 externality price (for Pigovian tax cases), the AP3 damage model, $9.41 million (2018) value of statistical life, and the Pope et al. (2019) concentration-response function. RideAustin trip data is used in all three cases. Region-specific values include: gasoline prices, electricity prices, marginal grid emissions estimates, and criteria pollutant externality cost estimates.

	Real Discount Rate	Externalities Internalized?	Vehicles Purchased	Vehicle-Distance Traveled (1K miles/year)	Private Costs (¢/trip-mile)		External Costs (¢/trip-mile)		Social Costs (¢/trip-mile)	
					Value	Change	Value	Change	Value	Change
Austin	1%		1 14 21	14 747 527	40.5	0.2%	5.9	-4.8%	46.4	-0.4%
		X	14 22	404 886	40.6		5.6		46.3	
	7%		24 7 5	340 601 346	45.0	1.3%	7.6	-16.7%	52.6	-1.3%
		X	1 16 19	14 395 881	45.6		6.3		51.9	
	13%		27 8 1	382 784 121	48.9	0.6%	8.4	-8.6%	57.3	-0.7%
		X	21 6 9	298 327 663	49.2		7.7		56.9	
Los Angeles	1%		6 32	85 1207	43.2	0.8%	5.0	-17.8%	48.2	-1.1%
		X	2 2 35	1 11 1284	43.5		4.1		47.7	
	7%		14 22	198 1092	48.5	1.7%	7.0	-30.8%	55.4	-2.3%
		X	2 2 35	2 8 1286	49.3		4.8		54.1	
	13%		9 8 19	127 114 1049	53.4	2.6%	8.8	-36.4%	62.3	-2.9%
		X	2 3 33	2 15 1280	54.9		5.6		60.5	
New York City	1%		20 16	938 349	42.9	1.0%	10.9	-6.6%	53.8	-0.6%
		X	13 25	721 568	43.3		10.2		53.5	
	7%		5 22 9	71 1012 204	47.7	0.9%	12.4	-11.5%	60.1	-1.7%
		X	14 22	782 506	48.1		11.0		59.1	
	13%		26 10	368 919	51.9	1.9%	15.4	-23.1%	67.4	-3.9%
		X	18 18	889 398	52.9		11.9		64.8	

Conventional Vehicle (CV) | Hybrid Electric Vehicle (HEV) | Battery Electric Vehicle (BEV)

Fig. S5. Summary of results across three analysis regions, with and without considering resale cash flows in the model. Except where otherwise stated, scenarios reflect no labor costs, $33,950 BEV price (2019 Kia Soul), $50/ton CO_2 externality price (for Pigovian tax cases), the AP3 damage model, $9.41 million (2018) value of statistical life, and the Pope et al. (2019) concentration-response function. RideAustin trip data is used in all three cases. Region-specific values include: gasoline prices, electricity prices, marginal grid emissions estimates, and criteria pollutant externality cost estimates.

	BEV Purchase Price	Externalities Internalized?	Vehicles Purchased	Vehicle-Distance Traveled (1K miles/year)	Private Costs (¢/trip-mile)		External Costs (¢/trip-mile)		Social Costs (¢/trip-mile)	
					Value	Change	Value	Change	Value	Change
Austin	Resale		24 7 5	340 601 346	45.0	1.3%	7.3	-24.1%	52.3	-2.2%
		X	1 17 18	14 257 1019	45.6		5.6		51.2	
	No Resale		21 8 7	298 736 253	53.5	0.8%	7.5	-14.0%	61.0	-1.0%
		X	1 15 20	14 600 674	54.0		6.4		60.4	
Los Angeles	Resale		14 22	198 1092	48.5	1.3%	6.3	-34.3%	54.8	-2.8%
		X	2 3 33	2 13 1280	49.1		4.1		53.2	
	No Resale		12 24	222 1068	57.1	0.4%	7.2	-21.1%	64.3	-2.0%
		X	6 33	35 1210	57.3		5.7		63.0	
New York City	Resale		5 22 9	71 1012 204	47.7	0.8%	12.4	-11.5%	60.1	-1.8%
		X	15 21	809 479	48.0		11.0		59.0	
	No Resale		2 23 11	28 1007 252	56.1	1.0%	12.1	-7.8%	68.2	-0.6%
		X	14 23	785 503	56.7		11.1		67.8	

Conventional Vehicle (CV) | Hybrid Electric Vehicle (HEV) | Battery Electric Vehicle (BEV)

Fig. S6. Summary of results for no labor costs (base case) and a case of $12/hour labor costs across three analysis regions. Except where otherwise stated, scenarios reflect a 7% real private firm discount rate, $33,950 BEV price (2019 Kia Soul), $50/ton CO_2 externality price (for Pigovian tax cases), the AP3 damage model, $9.41 million (2018) value of statistical life, and the Pope et al. (2019) concentration-response function. RideAustin trip data is used in all three cases. Region-specific values include: gasoline prices, electricity prices, marginal grid emissions estimates, and criteria pollutant externality cost estimates.

	Labor Costs ($/hour)	Externalities Internalized?	Vehicles Purchased	Vehicle-Distance Traveled (1K miles/year)	Private Costs (¢/trip-mile)		External Costs (¢/trip-mile)		Social Costs (¢/trip-mile)	
					Value	Change	Value	Change	Value	Change
Austin	$0		24 7 5	340 601 346	45.0	1.3%	7.6	-16.7%	52.6	-1.3%
		X	1 16 19	14 395 881	45.6		6.3		51.9	
	$12		25 11	354 946	158.7	0.2%	8.0	-12.4%	166.7	-0.4%
		X	2 25 9	28 1074 198	159.0		7.0		166.0	
Los Angeles	$0		14 22	198 1092	48.5	1.7%	7.0	-30.8%	55.4	-2.3%
		X	2 2 35	2 8 1286	49.3		4.8		54.1	
	$12		20 16	898 403	164.0	0.1%	15.1	-48.6%	179.0	-4.0%
		X	6 33	244 1061	164.2		7.8		171.9	
New York City	$0		5 22 9	71 1012 204	47.7	0.9%	12.4	-11.5%	60.1	-1.7%
		X	14 22	782 506	48.1		11.0		59.1	
	$12		8 21 7	113 1035 152	161.4	0.1%	13.0	-9.7%	174.4	-0.6%
		X	21 15	988 313	161.6		11.7		173.3	

Conventional Vehicle (CV) | Hybrid Electric Vehicle (HEV) | Battery Electric Vehicle (BEV)

Fig. S7. Summary of results for the base case BEV (2019 Kia Soul with a 30 kWh battery priced at $33,950) and the 2020 Chevrolet Bolt (with a higher-capacity and more energy-dense 66 kWh battery priced at $36,620) across three analysis regions. Except where otherwise stated, scenarios reflect a 7% real private firm discount rate, no labor costs, $50/ton CO_2 externality price (for Pigovian tax cases), the AP3 damage model,$9.41 million (2018) value of statistical life, and the Pope et al. (2019) concentration-response function. RideAustin trip data is used in all three cases. Region-specific values include: gasoline prices, electricity prices, marginal grid emissions estimates, and criteria pollutant externality cost estimates.

	Powertrains Available	BEV Model	Externalities Internalized?	Vehicles Purchased	Vehicle-Distance Traveled (1K miles/year)	Private Costs (¢/trip-mile)		External Costs (¢/trip-mile)		Social Costs (¢/trip-mile)	
						Value	Change	Value	Change	Value	Change
Austin	All	2019 Kia Soul		24 7 5	340 601 346	45.0	1.3%	7.6	-16.7%	52.6	-1.3%
			X	1 16 19	14 395 881	45.6		6.3		51.9	
		2020 Chevy Bolt		23 10 3	326 808 154	45.0	1.3%	7.5	-23.7%	52.5	-2.3%
			X	1 19 16	14 588 686	45.6		5.7		51.3	
	BEV Only	2019 Kia Soul		42	1306	48.5	0.7%	6.5	-2.9%	55.0	0.3%
			X	43	1307	48.8		6.3		55.1	
		2020 Chevy Bolt		37	1300	48.1	-0.6%	5.6	-3.2%	53.8	-0.9%
			X	36	1300	47.9		5.5		53.3	
Los Angeles	All	2019 Kia Soul		14 22	198 1092	48.5	1.7%	7.0	-30.8%	55.4	-2.3%
			X	2 2 35	2 8 1286	49.3		4.8		54.1	
		2020 Chevy Bolt		12 24	170 1117	48.2	0.7%	6.2	-29.9%	54.5	-2.8%
			X	8 28	8 1282	48.6		4.4		53.0	
New York City	All	2019 Kia Soul		5 22 9	71 1012 204	47.7	0.9%	12.4	-11.5%	60.1	-1.7%
			X	14 22	782 506	48.1		11.0		59.1	
		2020 Chevy Bolt		3 20 13	42 712 533	47.4	1.3%	10.1	-18.2%	57.5	-2.1%
			X	12 24	409 879	48.1		8.3		56.3	

Conventional Vehicle (CV)	Hybrid Electric Vehicle (HEV)	Battery Electric Vehicle (BEV)

Fig. S8. Summary of results for the base case BEV price ($33,950) and a price reduction of $5,000 across three analysis regions. Except where otherwise stated, scenarios reflect a 7% real private firm discount rate, no labor costs, $50/ton CO_2 externality price (for Pigovian tax cases), the AP3 damage model, $9.41 million (2018) value of statistical life, and the Pope et al. (2019) concentration-response function. RideAustin trip data is used in all three cases. Region-specific values include: gasoline prices, electricity prices, marginal grid emissions estimates, and criteria pollutant externality cost estimates.

	BEV Purchase Price	Externalities Internalized?	Vehicles Purchased	Vehicle-Distance Traveled (1K miles/year)	Private Costs (¢/trip-mile)		External Costs (¢/trip-mile)		Social Costs (¢/trip-mile)	
					Value	Change	Value	Change	Value	Change
Austin	$33950		24 7 5	340 601 346	45.0	1.3%	7.3	-24.1%	52.3	-2.2%
		X	1 17 18	14 257 1019	45.6		5.6		51.2	
	$28950		6 33	85 1208	41.9	-0.4%	5.7	-7.1%	47.6	-1.2%
		X	6 32	85 1208	41.7		5.3		47.0	
Los Angeles	$33950		14 22	198 1092	48.5	1.3%	6.3	-34.3%	54.8	-2.8%
		X	2 3 33	2 13 1280	49.1		4.1		53.2	
	$28950		5 34	71 1221	44.0	0.2%	4.9	-17.5%	48.9	-1.6%
		X	2 2 35	52 1290	44.1		4.1		48.1	
New York City	$33950		5 22 9	71 1012 204	47.7	0.8%	12.4	-11.5%	60.1	-1.8%
		X	15 21	809 479	48.0		11.0		59.0	
	$28950		8 30	501 789	45.7	0.1%	10.6	-2.1%	56.3	-0.3%
		X	7 31	411 879	45.8		10.4		56.1	

Conventional Vehicle (CV) | Hybrid Electric Vehicle (HEV) | Battery Electric Vehicle (BEV)

Fig. S9. Summary of results estimating criteria pollutant externalities using the AP3 reduced complexity model (base case) and two other reduced complexity models across three analysis regions. Except where otherwise stated, scenarios reflect a 7% real private firm discount rate, no labor costs, $33,950 BEV price (2019 Kia Soul), $50/ton CO_2 externality price (for Pigovian tax cases), $9.41 million (2018) value of statistical life, and the Pope et al. (2019) concentration-response function. RideAustin trip data is used in all three cases. Region-specific values include: gasoline prices, electricity prices, marginal grid emissions estimates, and criteria pollutant externality cost estimates.

	Damage Model	Externalities Internalized?	Vehicles Purchased	Vehicle-Distance Traveled (1K miles/year)	Private Costs (¢/trip-mile) Value	Private Costs (¢/trip-mile) Change	External Costs (¢/trip-mile) Value	External Costs (¢/trip-mile) Change	Social Costs (¢/trip-mile) Value	Social Costs (¢/trip-mile) Change
Austin	AP3		24 7 5	340 601 346	45.0	1.3%	7.6	-16.7%	52.6	-1.3%
		X	1 16 19	14 895 881	45.6		6.3		51.9	
	INMAP		24 7 5	340 601 346	45.0	1.6%	6.9	-38.3%	51.9	-3.7%
		X	14 22	198 1092	45.7		4.3		50.0	
	EASIUR		24 7 5	340 601 346	45.0	1.0%	5.2	-18.4%	50.2	-1.0%
		X	5 14 17	71 341 876	45.4		4.2		49.6	
Los Angeles	AP3		14 22	198 1092	48.5	1.7%	7.0	-30.8%	55.4	-2.3%
		X	22 35	2 8 1286	49.3		4.8		54.1	
	INMAP		14 22	198 1092	48.5	0.7%	4.0	-29.2%	52.5	-1.6%
		X	13 34	2 42 1251	48.8		2.8		51.7	
	EASIUR		14 22	198 1092	48.5	0.8%	3.3	-10.4%	51.8	0.0%
		X	5 33	71 1223	48.8		3.0		51.8	
New York City	AP3		5 22 9	71 1012 204	47.7	0.9%	12.4	-11.5%	60.1	-1.7%
		X	14 22	782 506	48.1		11.0		59.1	
	INMAP		5 22 9	71 1012 204	47.7	7.2%	19.5	-80.8%	67.2	-18.3%
		X	5 32	70 1225	51.1		3.8		54.9	
	EASIUR		5 22 9	71 1012 204	47.7	1.3%	9.5	-21.2%	57.2	-2.4%
		X	13 23	700 587	48.3		7.5		55.8	

Conventional Vehicle (CV) | Hybrid Electric Vehicle (HEV) | Battery Electric Vehicle (BEV)

Fig. S10. Summary of results estimating criteria pollutant externalities using a social cost of carbon for CO2-equivalent GHG emissions of $50/tonne (base case) and $300/tonne. Except where otherwise stated, scenarios reflect a 7% real private firm discount rate, no labor costs, $33,950 BEV price (2019 Kia Soul), $50/ton CO_2 externality price (for Pigovian tax cases), $9.41 million (2018) value of statistical life, and the Pope et al. (2019) concentration-response function. RideAustin trip data is used in all three cases. Region-specific values include: gasoline prices, electricity prices, marginal grid emissions estimates, and criteria pollutant externality cost estimates.

	Social Cost of Carbon	Externalities Internalized?	Vehicles Purchased	Vehicle-Distance Traveled (1K miles/year)	Private Costs (¢/trip-mile)		External Costs (¢/trip-mile)		Social Costs (¢/trip-mile)	
					Value	Change	Value	Change	Value	Change
Austin	$50		24 7 5	340 601 346	45.0	1.3%	7.6	-16.7%	52.6	-1.3%
		X	1 16 19	14 395 881	45.6		6.3		51.9	
	$300		24 7 5	340 601 346	45.0	3.9%	21.5	-43.3%	66.5	-11.4%
		X	4 34	57 1239	46.7		12.2		58.9	
Los Angeles	$50		14 22	198 1092	48.5	1.7%	7.0	-30.8%	55.4	-2.3%
		X	2 2 35	2 8 1286	49.3		4.8		54.1	
	$300		14 22	198 1092	48.5	1.3%	13.6	-28.4%	62.1	-5.2%
		X	2 2 34	3 6 1290	49.1		9.7		58.9	
New York City	$50		5 22 9	71 1012 204	47.7	0.9%	12.4	-11.5%	60.1	-1.7%
		X	14 22	782 506	48.1		11.0		59.1	
	$300		5 22 9	71 1012 204	47.7	4.9%	25.1	-31.3%	72.8	-7.5%
		X	6 31	285 1006	50.0		17.3		67.3	

Conventional Vehicle (CV) | Hybrid Electric Vehicle (HEV) | Battery Electric Vehicle (BEV)

Fig. S11. Trips are sampled from eight unique categories of weekday/weekend and season (top).

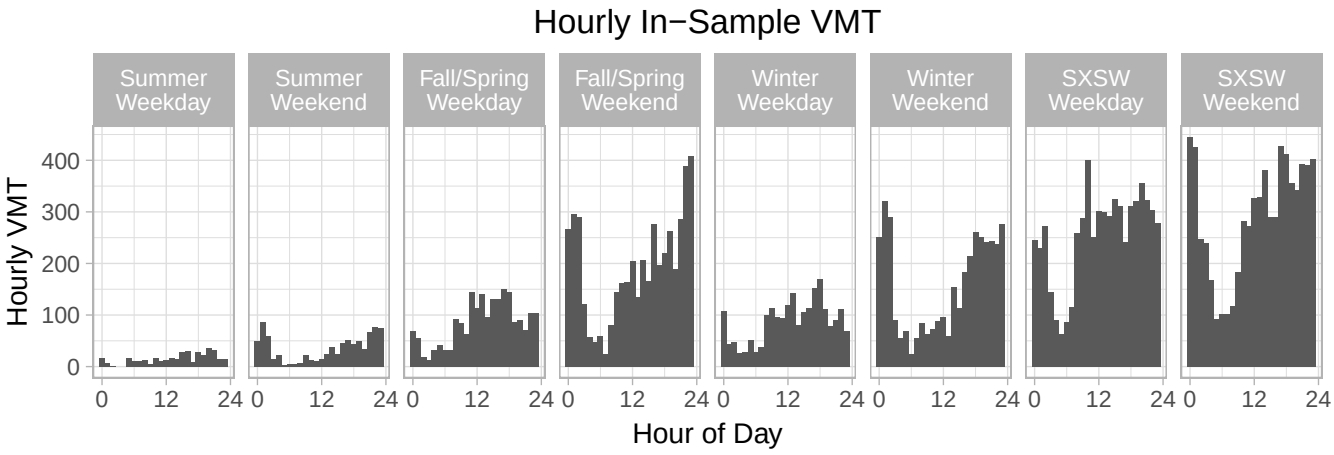


Fig. S12. Density map of trips that begin or end in the general vicinity of Austin's city center. Demand is densely concentrated in the core of downtown (center of plot), with a secondary peak of demand at Austin-Bergstrom International Airport (southeast corner of plot) and some demand concentrated near the University of Texas-Austin (just north of downtown).

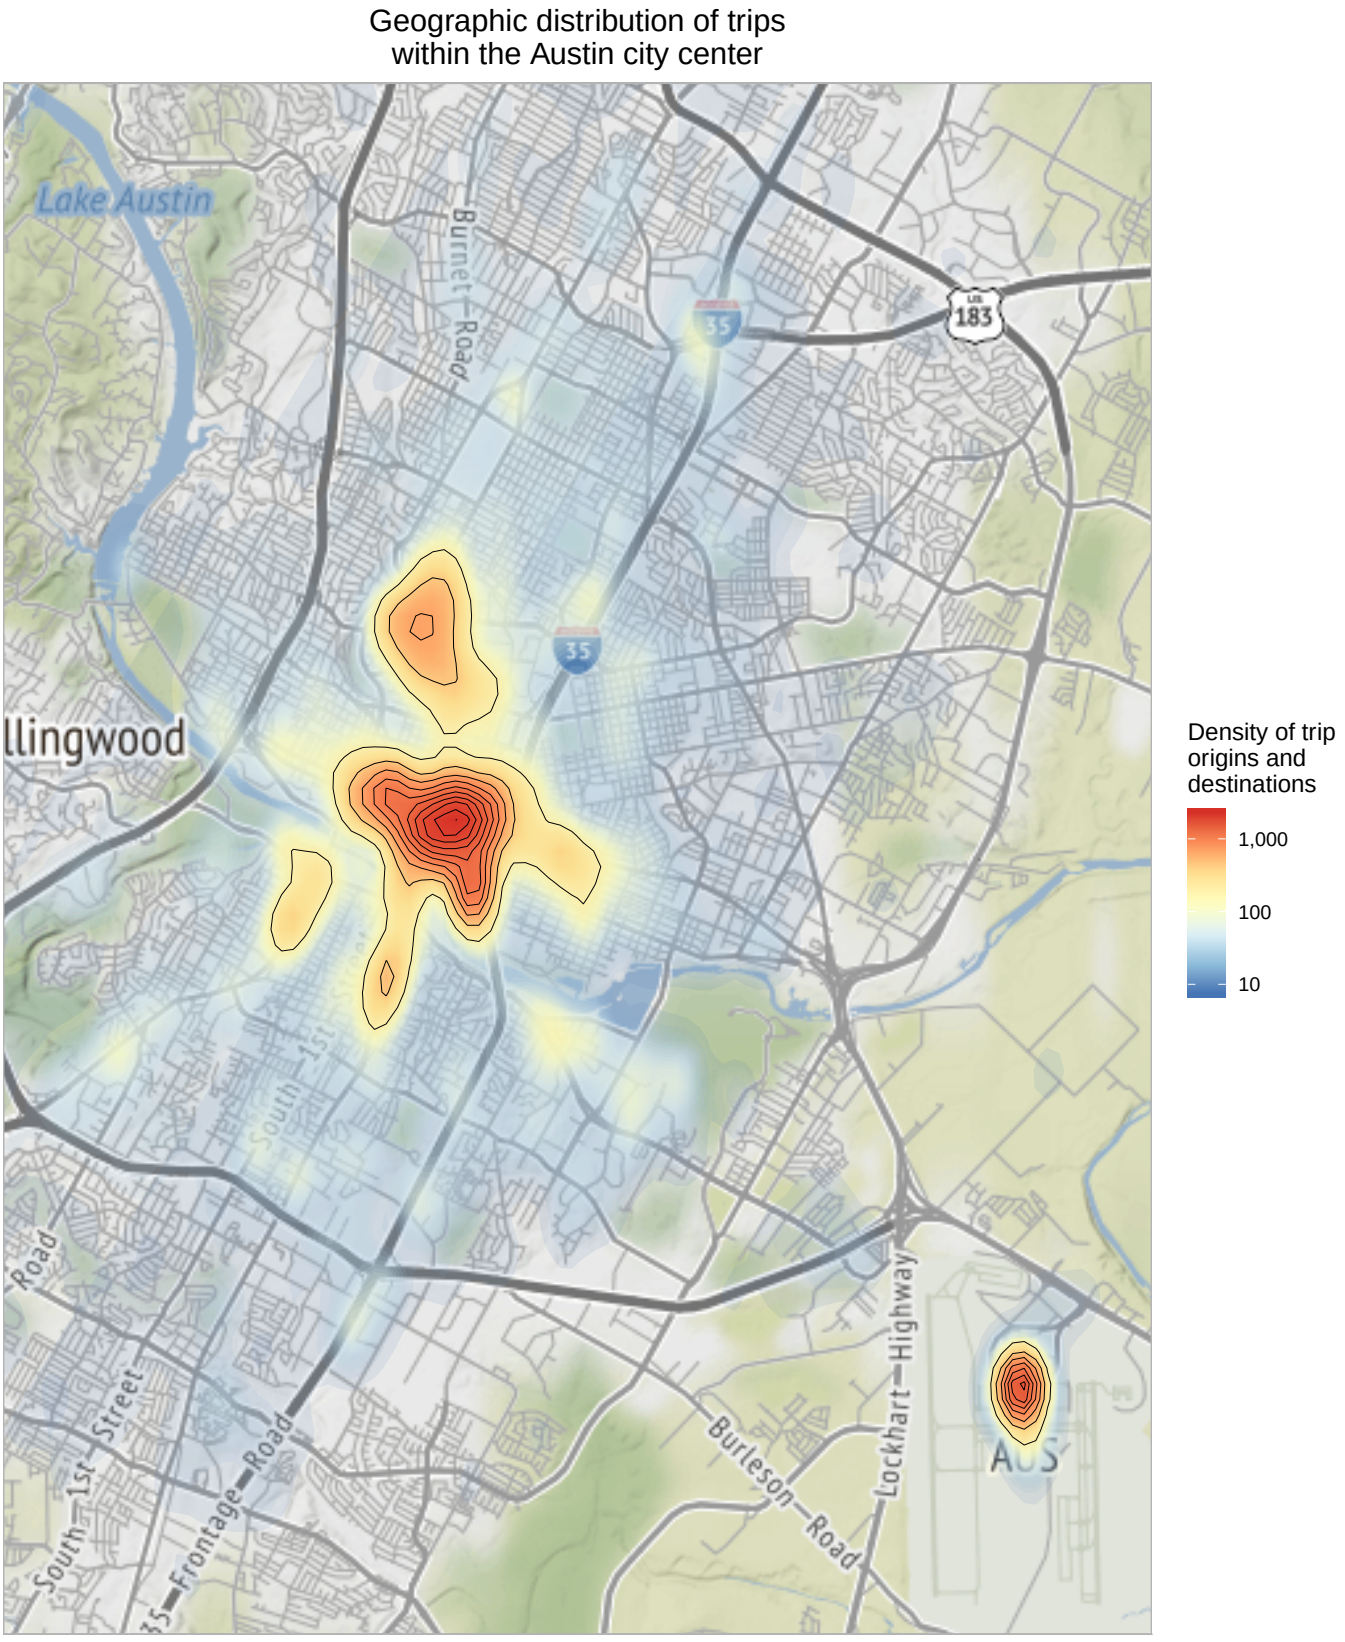


Fig. S13. Voronoi diagram showing the centroid and boundary of each cluster of trip origin-destination pairs. Grouping trips via k-means clustering has the effect of marginally reducing the size of the optimization network, particularly for high-demand periods in the downtown core. Clusters are colored according to their trip volumes. Over 1 million trips (out of approximately 1.5 million in total) originate and/or end at the highest-volume cluster, in downtown Austin.

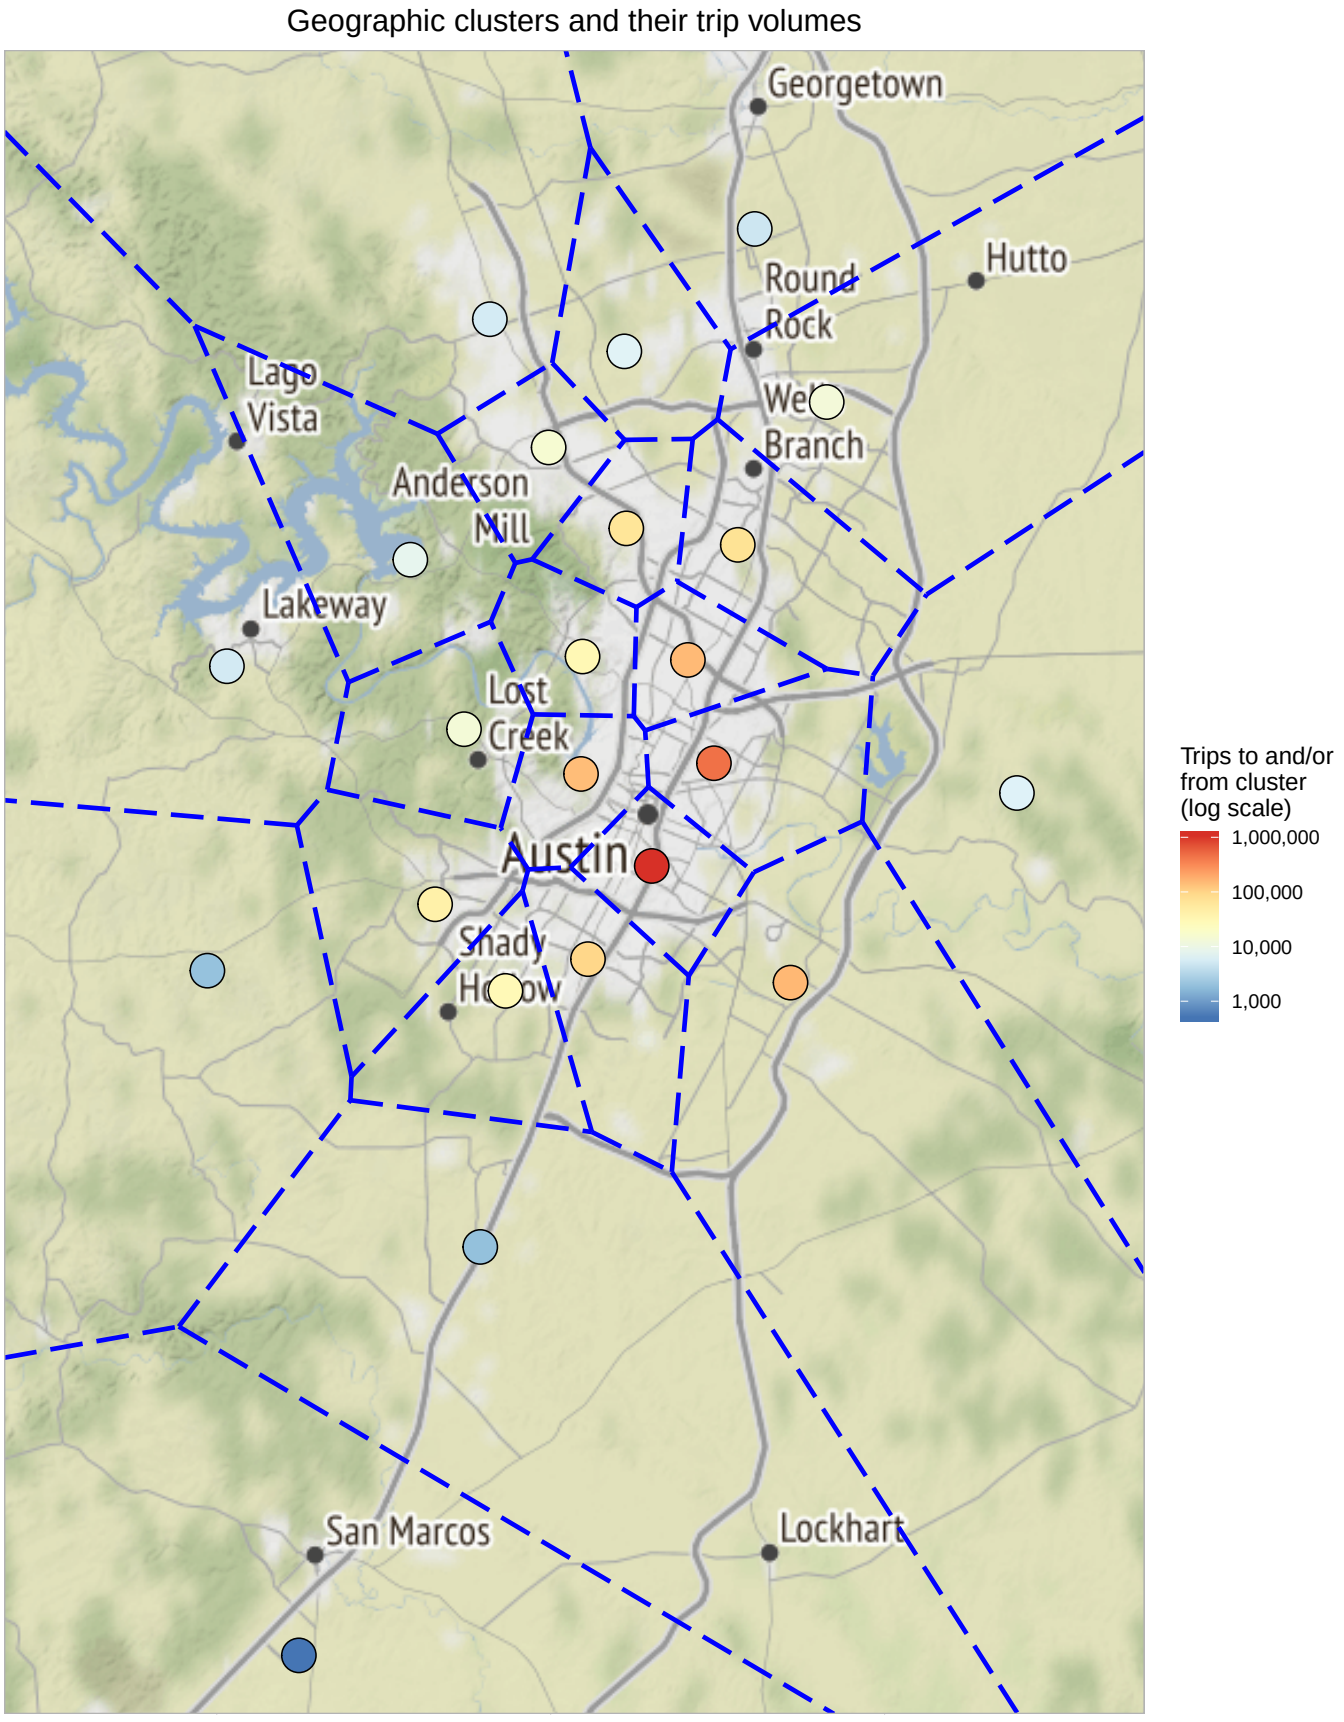


Fig. S14. Clustering validation plot showing within-cluster sum of squares (log scale), a measure of within-cluster similarity, as a function of the total number of clusters. By this measure, increasing number of clusters with hierarchical clustering reaches diminishing returns prior to the chosen number, 23.

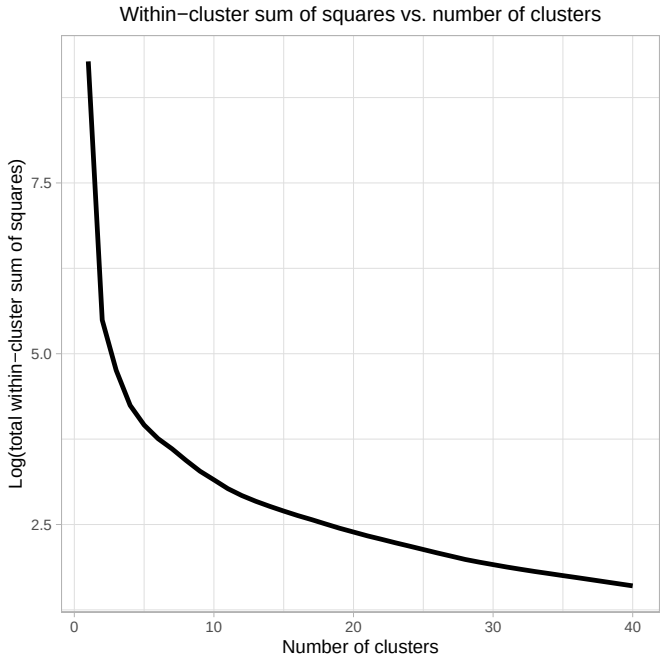


Fig. S15. Clustering validation plot showing average cluster silhouette width, which measures how well-clustered the data is via the average distance between clusters, as a function of the total number of clusters. By this measure, hierarchical clustering performs similarly well for any number of clusters above 10.

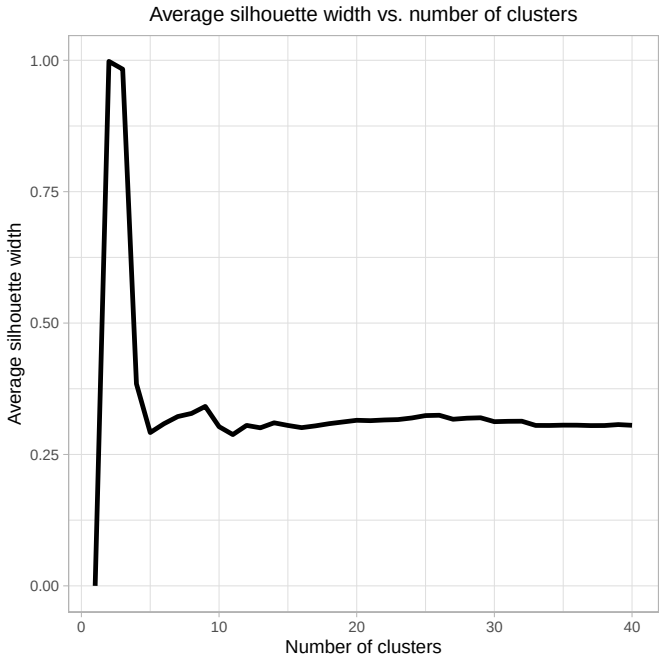


Fig. S16. Clustering validation plot showing gap statistics, which measure how different the clustered data's structure is than randomly distributed data. By this measure, hierarchical clustering performs similarly well for any number of clusters above 20.

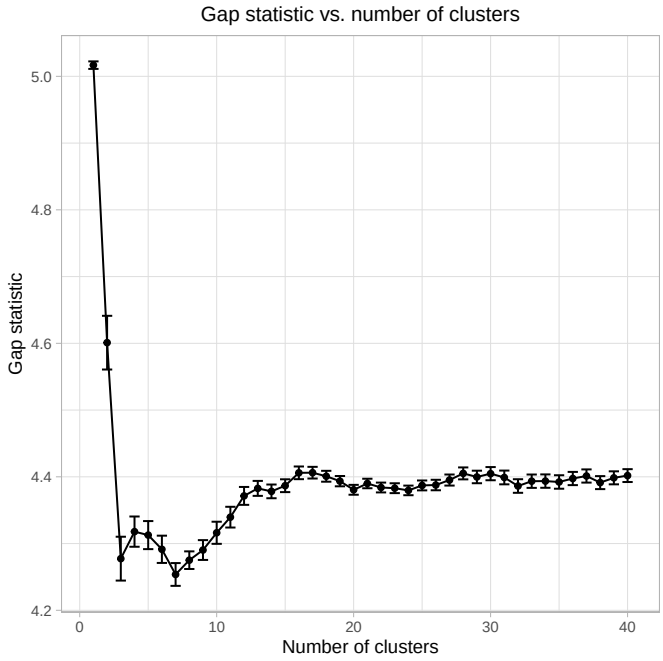


Fig. S17. Validation set prediction intervals for relocation arc distances. The x-axis represents actual values of trip distances in the validation set and the y-axis represents the distances predicted using k-nearest neighbors. Coloring, and black lines, show the empirical prediction intervals on the validation set for each x-value. The red line indicates where a perfect prediction would fall (i.e., predicted value equals actual value). Histograms along each axis show the marginal distributions of actual (top of plot) and predicted (right-hand side of plot) distance.

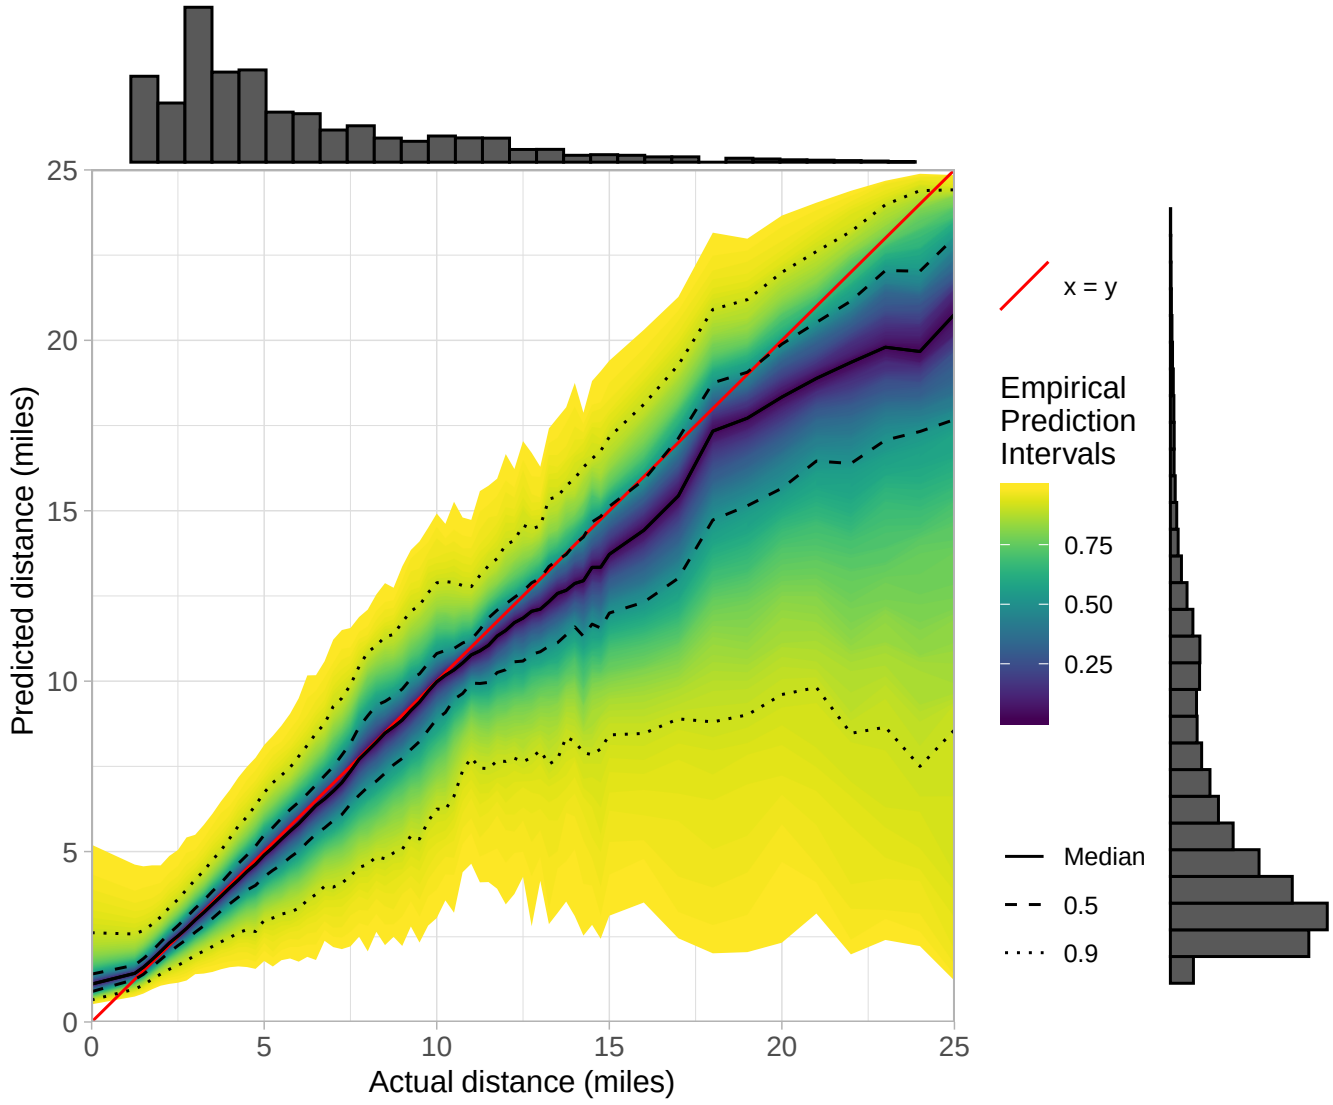
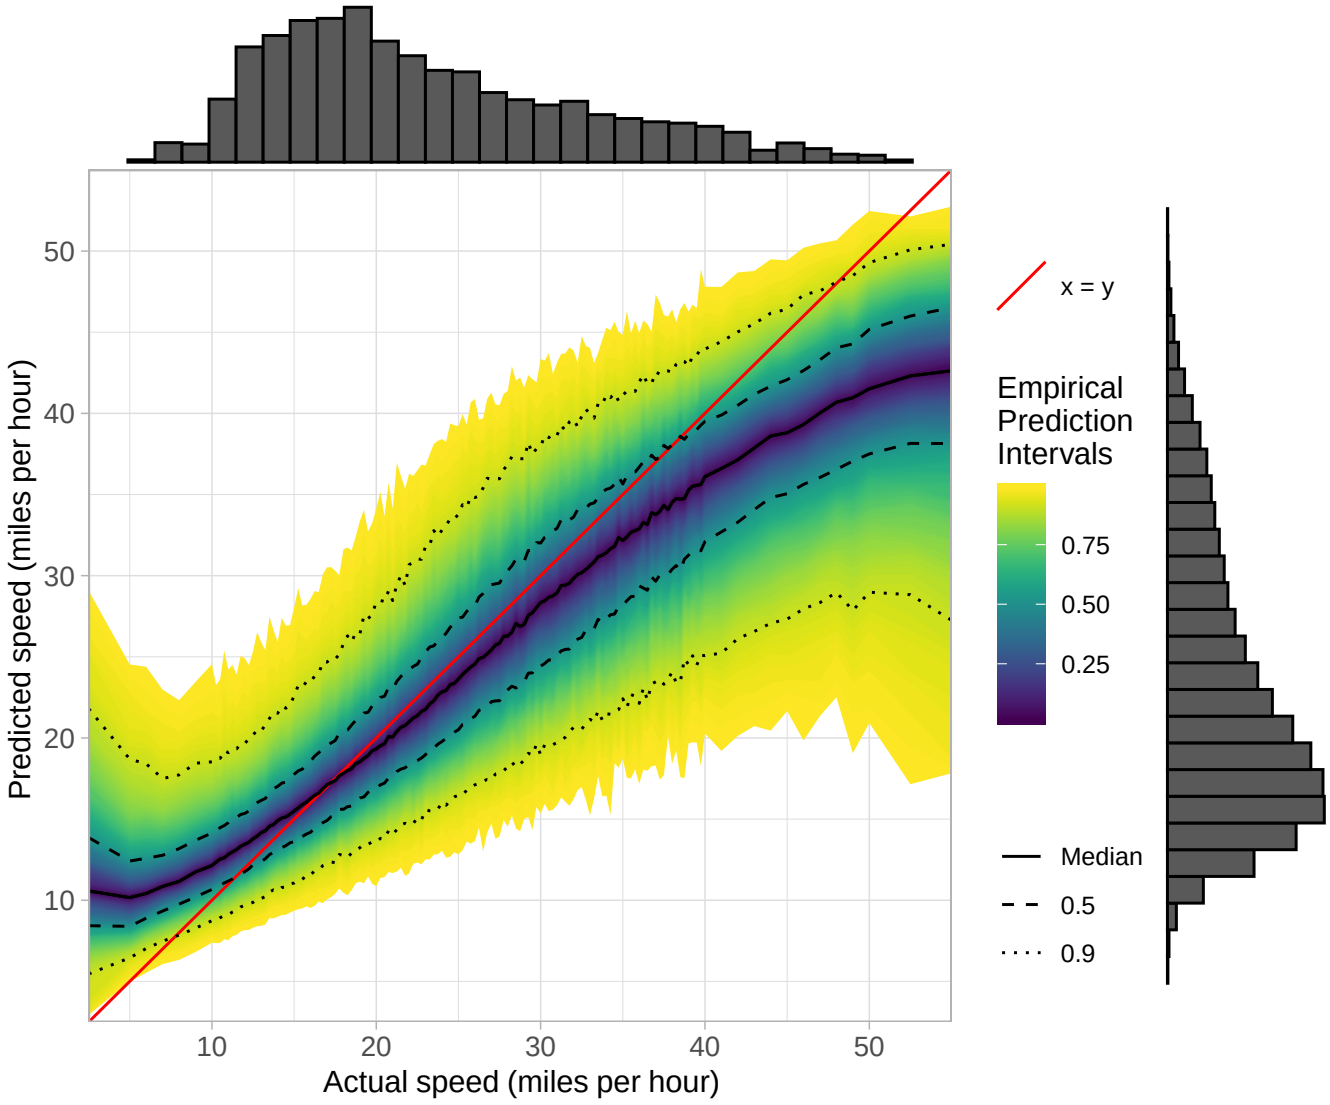


Fig. S18. Validation set prediction intervals for relocation arc speeds. The x-axis represents actual values of trip speeds in the validation set and the y-axis represents the speeds predicted using k-nearest neighbors. Coloring, and black lines, show the empirical prediction intervals on the validation set for each x-value. The red line indicates where a perfect prediction would fall (i.e., predicted value equals actual value). Histograms along each axis show the marginal distributions of actual (top of plot) and predicted (right-hand side of plot) speed.



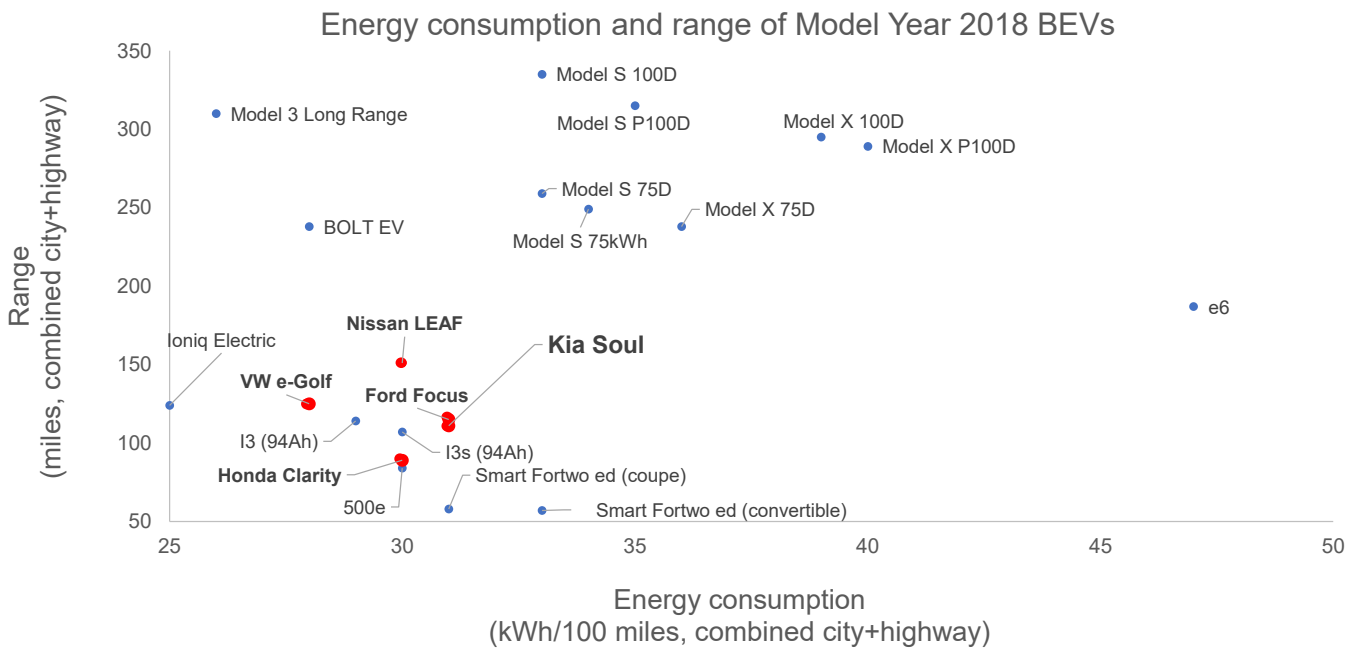


Fig. S19. The Kia Soul's range and energy consumption are typical of model year 2018 BEVs, excluding the longer ranges of the Chevrolet Bolt and the Tesla Model S, X, and 3 and the higher energy consumption of the BYD e6. Data collected from fueleconomy.gov (26), with red dots for BEVs that have similar CV counterparts.

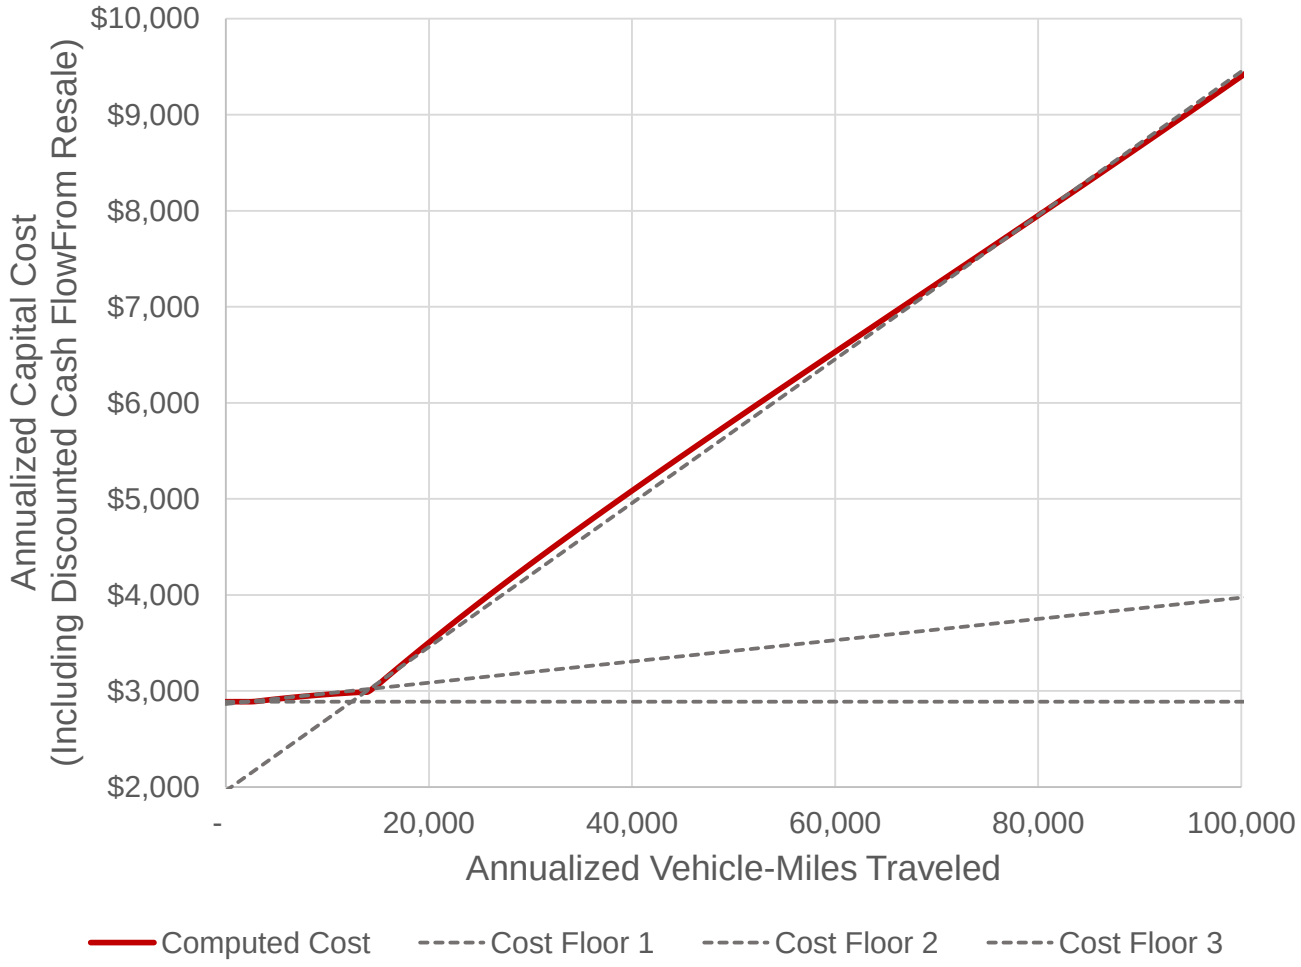


Fig. S20. Illustration (using data for a HEV) of how three linear floor constraints are used to approximate the nonlinear regressions that relate each vehicle's annualized capital costs to the vehicle's annualized vehicle-distance traveled.

Table S1. Comparison of air emissions externalities estimated to result from vehicle manufacture for each powertrain type.

	Powertrain		
External Cost Category	CV	HEV	BEV
Greenhouse gas	$258	$281	$351
Health	$2,436	$2,752	$4,025
Total	$2,694	$3,033	$4,376

Table S2. Comparison of air emissions externalities (2018 ¢/gallon) estimated to result from consumption of gasoline in each analysis region.

		City		
External Cost Category	Stage	Austin	Los Angeles	New York City
Greenhouse gas	Tailpipe	46.1¢	46.1¢	46.1¢
	Refining	11.2¢	11.2¢	11.2¢
Health	Tailpipe	34.5¢	232.9¢	108.6¢
	Refining	23.5¢	87.8¢	80.3¢
Total		115.2¢	378.9¢	246.2¢

Table S3. Comparison of air emissions externalities (2018 ¢/kWh) estimated to result from consumption of electricity in each analysis region. Within the optimization model, these externalities vary by region, season, and time of day. The values shown in this table are averages across 24 hours of day and 8 seasonal sampling segments, weighted by each segment's multiplication factor (see Table S6).

	City		
External Cost Category	Austin	Los Angeles	New York City
Greenhouse gas	0.35¢	0.28¢	0.28¢
Health	0.87¢	0.40¢	3.04¢
Total	1.22¢	0.68¢	3.32¢

Table S4. Formulation of the `FullMILP` optimization problem.

$$\underset{\mathcal{X}}{\text{minimize}} \sum_{k\in\mathcal{K}} \kappa_k + \sum_{(i,j)\in\mathcal{A}} c_{k,i,j} a_{k,i,j} + \sum_{k\in\mathcal{K}} \sum_{t\in\mathcal{T}_+} c_t \Delta q_{k,t}^{\text{CHG}} + \tau \left[\sum_{k\in\mathcal{K}} \delta_k + \sum_{(i,j)\in\mathcal{A}} d_{k,i,j} a_{k,i,j} + \sum_{k\in\mathcal{K}} \sum_{t\in\mathcal{T}_+} d_t \Delta q_{k,t}^{\text{CHG}} \right]$$

Minimize private and external car, fuel, and electricity costs [5a]

subject to

$\sum_{i\in\mathcal{V}} a_{k,i,j} = \sum_{i\in\mathcal{V}} a_{k,j,i}$	$\forall k \in \mathcal{K}, j \in \mathcal{V}\backslash\{r,s\}$	Flow is preserved across nodes except source and sink	[5b]
$\sum_{k\in\mathcal{K}} a_{k,i,j} = n_{i,j}$	$\forall (i,j) \in \{\mathcal{A} : n_{i,j} > 0\}$	Demand is satisfied	[5c]
$\sum_{j\in\mathcal{V}\backslash r} a_{k,r,j} = n_k$	$\forall k \in \mathcal{K}$	Dispatched cars are purchased	[5d]
$\kappa_k \geq \alpha_{\omega,k}^{\text{COSTS}} n_k + \beta_{\omega,k}^{\text{COSTS}} \sum_{(i,j)\in\mathcal{A}} m_{i,j} a_{k,i,j}$	$\forall k \in \mathcal{K}, \omega \in \Omega_k$	Per-vehicle private capital costs vary with usage	[5e]
$\delta_k \geq \alpha_{\omega,k}^{\text{DAMAGES}} n_k + \beta_{\omega,k}^{\text{DAMAGES}} \sum_{(i,j)\in\mathcal{A}} m_{i,j} a_{k,i,j}$	$\forall k \in \mathcal{K}, \omega \in \Omega_k$	Per-vehicle manufacturing external costs vary with usage	[5f]
$q_{k,t+1} \leq q_{k,t} + \sum_{(i,j)\in\{\mathcal{A}:t_i=t\}} a_{k,i,j} \Delta q_{k,i,j}^{\text{MAX}}$	$\forall k \in \mathcal{K}_{\text{B}}, t \in \mathcal{T}_{\text{Q}}$	BEV charge level is tracked (times with a charger starting timeslot)	[5g]
$q_{k,t+1} = q_{k,t} + \sum_{(i,j)\in\{\mathcal{A}:t_i=t\}} a_{k,i,j} \Delta q_{k,i,j}^{\text{MAX}}$	$\forall k \in \mathcal{K}_{\text{B}}, t \in \mathcal{T}\backslash\mathcal{T}_{\text{Q}}$	BEV charge level is tracked (times with no charger starting timeslot)	[5h]
$0 \leq q_{k,t} \leq q_k^{\text{MAX}}$	$\forall k \in \mathcal{K}_{\text{B}}, t \in \mathcal{T}$	Charge level does not exceed battery capacity	[5i]
$a_{k,i,j} \in \{0,1\},\ n_k \in \{0,1\}$	$\forall k \in \mathcal{K}_{\text{B}}, (i,j) \in \mathcal{A}$	BEV routing and purchase decisions are binary	[5j]
$a_{k,i,j} \in \mathbb{Z}_+,\ n_k \in \mathbb{Z}_+$	$\forall k \in \mathcal{K}\backslash\mathcal{K}_{\text{B}}, (i,j) \in \mathcal{A}$	CV & HEV routing and purchase decisions are integral	[5k]
$q_{k,t} \in \mathbb{R}_+$	$\forall k \in \mathcal{K}_{\text{B}}, t \in \mathcal{T}$	BEV charge level is always nonnegative	[5l]
$\kappa_k \in \mathbb{R}_+$	$\forall k \in \mathcal{K}$	Per-vehicle private capital costs are nonnegative	[5m]
$\delta_k \in \mathbb{R}_+$	$\forall k \in \mathcal{K}$	Per-vehicle manufacturing external costs are nonnegative	[5n]

where

$\Delta q_{k,t}^{\text{CHG}} = q_{k,t+1} - q_{k,t} + \sum_{(i,j)\in\{\mathcal{A}:t_i=t, \Delta q_{k,i,j}^{\text{MAX}}<0\}} a_{k,i,j} \Delta q_{k,i,j}^{\text{MAX}}$	$\forall k \in \mathcal{K}_{\text{B}}, t \in \mathcal{T}_{\text{Q}}$	Charging of each BEV at each timestep equals charge change minus routing losses	[5o]

Table S5. Sets, decision variables, and input parameters

Label	Type	Description
$\mathcal{X}$	Set	All decision variables
$\mathcal{V}$	Set	Vertices representing points in space-time
$\mathcal{A}$	Set	Arcs connecting feasible pairs of vertices in $\mathcal{V}$
$\mathcal{K}$	Set	Vehicles or vehicle types (BEVs are represented individually, whereas CVs and HEVs are each tracked as a group)
$\mathcal{K}_{\mathrm{B}}$	Set	Battery electric vehicles (subset of $\mathcal{K}$, indexed individually)
$\mathcal{T}$	Set	All unique arc start and end times
$\mathcal{T}_{\mathrm{Q}}$	Set	All unique charging arc start times (subset of $\mathcal{T}$)
Ω_k	Set	Linear constraints that make up the piecewise linear convex cost floor for capital cost κ_k for vehicle type k
n_k	Variable	Number of vehicle k purchased (BEVs are tracked individually, whereas CVs and HEVs are tracked as a group)
$a_{k,i,j}$	Variable	Assignment of vehicle k to arc (i,j)
$q_{k,t}$	Variable	Charge level of vehicle k at time t
$\Delta q_{k,t}^{\mathrm{CHG}}$	Variable	Energy charged to vehicle k from the grid at time t
κ_k	Variable	Private acquisition cost for vehicle k
δ_k	Variable	Manufacturing damages for vehicle k
τ	Parameter	Flag controlling whether air emissions externalities are included as a tax
r	Parameter	Source vertex from which all routes originate
s	Parameter	Sink vertex at which all routes terminate
t_i	Parameter	Time of vertex i
$n_{i,j}$	Parameter	Number of trips requested along arc (i,j)
$m_{i,j}$	Parameter	Travel distance along arc (i,j) (annualized)
m_{MAX}	Parameter	Maximum lifetime travel distance of a vehicle
q_k^{MAX}	Parameter	Energy capacity of vehicle k (∞ for CVs and HEVs)
$c_{k,i,j}$	Parameter	Private cost for vehicle k to traverse arc (i,j)
$d_{k,i,j}$	Parameter	External cost from vehicle k traversing arc (i,j)
c_t	Parameter	Private cost per kWh of electricity from the grid at time t
d_t	Parameter	External cost per kWh of electricity from the grid at time t
$\Delta q_{k,i,j}^{\mathrm{MAX}}$	Parameter	Maximum energy change for car k induced by travel on arc (i,j) (positive for charging arcs, negative for all others)
$\alpha_{\omega,k}^{\mathrm{COSTS}}$	Parameter	Intercept term for line ω, representing a portion of the convex piecewise linear VDT-dependent capital costs
$\beta_{\omega,k}^{\mathrm{COSTS}}$	Parameter	Slope term for line ω, representing a portion of the convex piecewise linear VDT-dependent capital costs
$\alpha_{\omega,k}^{\mathrm{DAMAGES}}$	Parameter	Intercept term for line ω, representing a portion of the convex piecewise linear VDT-dependent manufacturing external costs
$\beta_{\omega,k}^{\mathrm{DAMAGES}}$	Parameter	Slope term for line ω, representing a portion of the convex piecewise linear VDT-dependent manufacturing external costs

Table S6. Multiplication factors used within the optimization to compute annualized costs and distances. These factors convert eight representative days (24 hours per segment) to one average year of 366.25 days. For sampling purposes, each day begins at 6 AM. For example, "Weekends", labeled here for simplicity as Friday-Saturday, are sampled from trips occurring from Friday at 6:00:00 AM to Sunday at 5:59:59 AM.

Segment	Definition	Multiplier for annualized costs and distances
Summer Weekday	Sunday-Thursday, May-September	109.29
Summer Weekend	Friday-Saturday, May-September	43.71
Fall/Spring Weekday	Sunday-Thursday, April & October	43.57
Fall/Spring Weekend	Friday-Saturday, April & October	17.43
Winter Weekday	Sunday-Thursday, November-March*	103.04
Winter Weekend	Friday-Saturday,November-March**	39.21
South by Southwest Weekday	March 12-16, 2017	5.00
South by Southwest Weekend	March 10,11,17, & 18, 2017	4.00

*Minus 5 South by Southwest days
** Minus 4 South by Southwest days

Table S7. Frequency table of passenger trip requests per trip arc. Clustering start/end locations by geography and rounding start/end times to the nearest three minutes preserves full geographic and temporal fidelity for 4,370 trips, and the remainder of trips are assigned the average start/end location and time of two or more trips.

Demand (number of trip requests per trip arc)	Number of trips
1	4,370
2	271
3	45
4	8
5	2

Table S8. Assumed vehicle parameters (prices in 2018 USD).

Powertrain	Reference Vehicle	Purchase Price	Energy Capacity	Efficiency		
				City	Highway	Units
CV	2019 Kia Soul CV (1.6L, naturally aspirated, automatic transmission)	$20,490	Unlimited	26.1	36.9	miles/gallon
HEV	Reference CV, with price and fuel consumption rate (gallons/mile) multiplied by the ratio of relevant values from the 2019 Kia Optima sedan (HEV value : LX 2.4 CV value)	$25,040	Unlimited	40.7	39.7	miles/gallon
BEV	2019 Kia Soul EV	$33,950	30 kWh	27.3	36.1	kWh/100mi

Table S9. Resale values estimated for each powertrain type (2018 USD) as a function of m, annualized vehicle-miles traveled per car. When $m \leq 2,610$, the vehicle is resold after 12 years at a low enough mileage that resale value is independent of mileage; when $2,610 < m \leq 14,167$, the vehicle is resold after 12 years with a total mileage up to 170,000; when $m > 14,167$, the vehicle is resold prior to 12 years with a mileage of 170,000.

Powertrain	Vehicle Used for Regression	Resale Value Factor (Proportion of Purchase Price) as a function of m, annualized vehicle-miles traveled		
		$m \leq 2,610$	$2,610 \leq m \leq 14,167$	$m > 14,167$
CV	Toyota Camry CV	0.38	$0.425 - 1.85 \times 10^{-5} m + 4.24 \times 10^{-10} m^2$	$-0.724 + 0.101 \ln(m)$
HEV	Toyota Camry HEV	0.28	$0.313 - 1.39 \times 10^{-5} m + 2.98 \times 10^{-10} m^2$	$-1.69 + 0.199 \ln(m) - 1.60 \times 10^{-6} m$
BEV	Nissan Leaf	0.17	$0.180 - 7.50 \times 10^{-6} m + 1.72 \times 10^{-10} m^2$	$-2.23 + 0.239 \ln(m) - 1.93 \times 10^{-6} m$

Table S10. Base case assumptions for gasoline and electricity prices (2018 USD).

Analysis Region	Gasoline Prices	Utility	Electricity Rate
Austin	$2.20/gallon	Austin Energy	10.9¢/kWh
Los Angeles	$3.51/gallon	Southern California Edison	25.8¢/kWh
New York City	$2.77/gallon	Consolidated Edison	16.9¢/kWh

Table S11. Sampling of recent maintenance cost assumptions made in the literature comparing different powertrains (2017 USD).

Study	Year	Data or method (if documented)	Per-mile maintenance cost assumption			
			CV	HEV	PHEV-30	BEV
Oak Ridge Nat'l Lab (27)	2010	Partially itemized maintenance schedules from EPRI (2002) report (28)	5.1	4.6	3.8	**2.6
Propfe et al. (29)	2012	Mean time between failure, estimated costs to replace	10.9	10.2	10.1	5.5
Mitropoulous et al. (30)	2017	AAA costs for tire wear; 2001 study on ICEV & BEV life cycle costs (31); EPRI (2002) (28)	6.9	5.6	-	5.2
Hummel et al. (32)	2017	Teardown of Chevy Bolt BEV & VW Gold CV; itemized maintenance schedules	6.1	-	-	2.6
Pavlenko et al. (33)	2019	CV & BEV estimates from (32)	*6.1	3.7	-	*2.6
This study	2019	CV: median of original estimates from (27, 29, 30, 32, 33) HEV & BEV: median estimated percent of CV costs from (27, 29, 30, 33)	6.5	5.6		3.3

*Not an original estimate (Borrowed directly from prior work)
**Computed as PHEV-30 cost minus air filter and oil replacement as in Michalek et al. (2011) (34)

Table S12. Assumed emissions from gasoline production and combustion, estimated using GREET (grams per gallon of gasoline).

Stage	VOC	NOx	PM10	PM2.5	SOx	GHG
Well-to-Pump	3.556	4.116	0.392	0.280	1.960	2,240.000
Tailpipe	6.132	3.080	0.140	0.140	0.056	9,268.000

Table S13. County FIPs codes used for each analysis region's upstream fuel refinery emissions. Percent of upstream emissions are assumed to be proportional to percent of upstream refining capacity.

Analysis region	Upstream refinery FIPS code	Percent of region's refining capacity
Austin	48245 (Jefferson)	27.1%
	48201 (Harris)	24.6%
	48355 (Nueces)	14.8%
	48071 (Chambers)	7.3%
	48039 (Brazoria)	6.9%
	48167 (Galveston)	6.3%
	48341 (Moore)	3.3%
	48233 (Hutchinson)	2.5%
	48141 (El Paso)	2.3%
	48297 (Live Oak)	1.5%
	48423 (Smith)	1.3%
	48227 (Howard)	1.2%
	48029 (Bexar)	0.3%
	48177 (Gonzales)	0.1%
	48493 (Wilson)	0.1%
Los Angeles	06037 (Los Angeles)	53.3%
	06013 (Contra Costa)	36.3%
	06095 (Solano)	7.7%
	06029 (Kern)	2.2%
	06083 (Santa Barbara)	0.5%
New York City	42101 (Philadelphia)	27.9%
	34039 (Union)	21.5%
	42045 (Delaware)	15.8%
	10003 (New Castle)	15.2%
	34015 (Gloucester)	13.3%
	42123 (Warren)	5.4%
	42083 (McKean)	0.9%

Table S14. County FIPs codes used for each city's tailpipe emissions and eGRID region used for each region's marginal grid emissions.

Analysis Region	Tailpipe Emissions FIPS code	eGRID subregion
Austin	48453 (Travis County, TX)	ERCT
Los Angeles	06037 (Los Angeles County, CA)	CAMX
New York City	36061 (New York County, NY)	NYCW

Table S15. Estimated emissions from vehicle production, computed by GREET (grams per car).

Powertrain	Component	Stage	VOC	NOx	PM10	PM2.5	SOx	GHG
CV	Car	Copper mining	1.30	8.73	2.31	1.13	7.63	7,331.21
CV	Car	Copper production	7.95	140.08	9.07	5.64	1,237.12	66,886.88
CV	Car	Remainder	6,824.82	5,136.33	2,238.70	1,042.79	16,400.33	5,114,470.99
HEV	Car	Copper mining	2.91	19.60	5.18	2.54	17.13	16,456.61
HEV	Car	Copper production	17.85	314.44	20.35	12.67	2,777.01	150,143.13
HEV	Car	Remainder	6,953.34	5,241.65	2,349.75	1,098.03	17,176.82	5,292,367.65
BEV	Car	Copper mining	3.62	24.35	6.44	3.16	21.29	20,450.74
BEV	Car	Copper production	22.18	390.76	25.29	15.74	3,451.01	186,583.83
BEV	Car	Remainder	6,306.51	4,607.20	1,959.53	912.78	14,636.78	4,694,669.37
HEV	Battery	Copper mining	0.21	1.43	0.38	0.19	1.25	1,202.05
HEV	Battery	Nickel mining	0.33	7.27	1.24	0.63	3.02	1,825.89
HEV	Battery	Cobalt mining	0.69	7.13	37.62	4.22	31.46	2,339.81
HEV	Battery	Manganese mining	0.01	0.07	0.29	0.15	0.12	85.41
HEV	Battery	Copper production	1.30	22.97	1.49	0.93	202.84	10,967.03
HEV	Battery	Nickel production	0.46	15.39	7.21	3.61	702.19	3,440.74
HEV	Battery	Cobalt production	1.64	9.95	1.08	0.66	43.13	8,788.99
HEV	Battery	Lithium production	1.58	15.13	1.97	1.53	6.97	4,384.26
HEV	Battery	Graphite production	2.33	25.14	10.34	5.05	149.35	9,075.20
HEV	Battery	Remainder	34.01	118.35	28.46	15.02	289.55	141,576.58
BEV	Battery	Copper mining	1.24	8.37	2.21	1.09	7.32	7,032.86
BEV	Battery	Nickel mining	5.86	130.22	22.24	11.23	54.03	32,702.99
BEV	Battery	Cobalt mining	12.31	127.71	673.75	75.63	563.46	41,907.66
BEV	Battery	Manganese mining	0.01	0.07	0.29	0.15	0.12	85.41
BEV	Battery	Copper production	7.63	134.38	8.70	5.41	1,186.78	64,164.84
BEV	Battery	Nickel production	8.20	275.57	129.18	64.75	12,576.75	61,626.12
BEV	Battery	Cobalt production	29.31	178.17	19.26	11.87	772.54	157,417.08
BEV	Battery	Lithium production	28.38	271.07	35.23	27.32	124.88	78,525.21
BEV	Battery	Graphite production	41.41	446.87	183.72	89.81	2,654.27	161,287.11
BEV	Battery	Remainder	333.40	1,311.81	362.30	191.10	3,648.99	1,556,346.80

Table S16. Summary of NAICS (2012 revision) industry codes used to characterize upstream processes related to elements serving as key inputs of battery manufacturing processes. To create a rough estimate of emissions damages for processes with a negligible level of activity in the U.S., their emissions are assumed to be co-located with processes sharing the same NAICS code. For example, there is negligible nickel mining activity in the U.S., so nickel mining emissions are assumed to occur in U.S. counties where copper mining activity are located.

	Mining	Refining or Manufacture
Nickel	212234: Copper and Nickel Ore Mining	331410: Nonferrous Metal (except Aluminum) Smelting and Refining
Copper		
Cobalt	212299: All Other Metal Ore Mining	
Manganese		325180: Other Basic Inorganic Chemical Manufacturing
Lithium		
Graphite		335991: Carbon and Graphite Product Manufacturing

= Part of the modeled supply chain and US activity is non-negligible
= Part of the modeled supply chain but US activity is negligible
= Not part of the supply chain as modeled

Table S17. County FIPS codes.

Process	County FIPS code	Description
Copper mining & production	04*	Arizona
	16*	Idaho
	26*	Michigan
	29*	Missouri
	30*	Montana
	32*	Nevada
	35*	New Mexico
	48*	Texas
	49*	Utah
Cobalt mining	16059	Lemhi County, Idaho
	26103	Marquette County, Michigan
	27137	St. Louis County, Minnesota
	30095	Stillwater County, Montana
	30097	Sweet Grass County, Montana
Cobalt production	16059	Lemhi County, Idaho
	30095	Stillwater County, Montana
Manganese production	24*	Maryland
	32*	Nevada
	39*	Ohio
	47*	Tennessee
	54*	West Virginia
Lithium production	32009	Esmerelda County, Nevada

References

1. Yuksel T, Tamayao MAM, Hendrickson C, Azevedo IML, Michalek JJ (2016) Effect of regional grid mix, driving patterns and climate on the comparative carbon footprint of gasoline and plug-in electric vehicles in the United States. *Environmental Research Letters* 11(4):044007.
2. US Energy Information Administration (2018) Retail Gasoline and Diesel Prices: Regular - Conventional Areas, Annual Period.
3. US Energy Information Administration (2018) Average retail price of electricity, Texas, annual.
4. Muller NZ, Mendelsohn R (2006) The Air Pollution Emission Experiments and Policy analysis model (APEEP) Technical Appendix. *Yale University2* 1.
5. Heo J, Adams PJ, Gao HO (2016) Public Health Costs of Primary PM2.5 and Inorganic PM2.5 Precursor Emissions in the United States. *Environmental Science and Technology* 50(11):6061–6070.
6. Tessum CW, Hill JD, Marshall JD (2017) InMAP: A model for air pollution interventions. *PLoS ONE* 12(4):e0176131.
7. Siler-Evans K, Azevedo IL, Morgan MG (2012) Marginal Emissions Factors for the U.S. Electricity System. *Environmental Science & Technology* 46(9):4742–4748.
8. Azevedo IL, Horner N, Siler-evans K, Vaishnav P (2017) Electricity Marginal Factors Estimates. *Center for Climate and Energy Decision Making*.
9. Erdoğan S, Miller-Hooks E (2012) A Green Vehicle Routing Problem. *Transportation Research Part E: Logistics and Transportation Review* 48(1):100–114.
10. Felipe Ortuño MT, Righini G, Tirado G (2014) A heuristic approach for the green vehicle routing problem with multiple technologies and partial recharges. *Transportation Research Part E: Logistics and Transportation Review* 71:111–128.
11. Lin J, Zhou W, Wolfson O (2016) Electric Vehicle Routing Problem in *Transportation Research Procedia*. Vol. 12, pp. 508–521.
12. Montoya A, Guéret C, Mendoza JE, Villegas JG (2017) The electric vehicle routing problem with nonlinear charging function. *Transportation Research Part B: Methodological* 103.
13. Xiao Y, Konak A (2016) The heterogeneous green vehicle routing and scheduling problem with time-varying traffic congestion. *Transportation Research Part E: Logistics and Transportation Review* 88:146–166.
14. Schiffer M, Walther G (2017) The electric location routing problem with time windows and partial recharging. *European Journal of Operational Research* 260(3):995–1013.
15. Koç Karaoglan I (2016) The green vehicle routing problem: A heuristic based exact solution approach. *Applied Soft Computing* 39:154–164.
16. Andelmin J, Bartolini E (2017) An Exact Algorithm for the Green Vehicle Routing Problem. *Transportation Science* 51(August):trsc.2016.0734.
17. Bruglieri M, Mancini S, Pezzella F, Pisacane O (2016) A new Mathematical Programming Model for the Green Vehicle Routing Problem. *Electronic Notes in Discrete Mathematics* 55:89–92.
18. Hiermann G, Puchinger J, Ropke S, Hartl RF (2016) The Electric Fleet Size and Mix Vehicle Routing Problem with Time Windows and Recharging Stations. *European Journal of Operational Research* 252(3):995–1018.
19. Sassi O, Cherif WR, Oulamara A (2015) Vehicle Routing Problem with Mixed fleet of conventional and heterogenous electric vehicles and time dependent charging costs. *International Journal of Mathematical, Computational, Physical, Electrical and Computer Engineering* 9(3):167–177.
20. Vaz Penna PH, Afsar HM, Prins C, Prodhon C (2016) A Hybrid Iterative Local Search Algorithm for The Electric Fleet Size and Mix Vehicle Routing Problem with Time Windows and Recharging Stations. *IFAC-PapersOnLine* 49(12):955–960.
21. Goeke D, Schneider M (2015) Routing a mixed fleet of electric and conventional vehicles. *European Journal of Operational Research* 245(1):81–99.
22. Bertsimas D, Jaillet P, Martin S (2017) Online Vehicle Routing : The Edge of Optimization in Large-Scale Applications. *Operations Research* pp. 1–48.
23. Koç Bektaş T, Jabali O, Laporte G (2016) Thirty years of heterogeneous vehicle routing.
24. Repoussis PP, Tarantilis CD (2010) Solving the fleet size and mix vehicle routing problem with time windows via adaptive memory programming. *Transportation Research Part C: Emerging Technologies* 18(5):695–712.
25. Wang Z, Li Y, Hu X (2015) A heuristic approach and a tabu search for the heterogeneous multi-type fleet vehicle routing problem with time windows and an incompatible loading constraint. *Computers and Industrial Engineering* 89:162–176.
26. US Environmental Protection Agency (2018) Download Fuel Economy Data: 2018 Datafile.
27. Laboratory ORN (2010) Plug-in hybrid electric vehicle value proposition study. *U.S. DEPARTMENT OF ENERGY* p. 218.
28. Graham R, Little AD (2001) Comparing the Benefits and Impacts of Hybrid Electric Vehicle Options, (Electric Power Research Institute), Technical Report July.
29. Propfe B, Redelbach M, Santini DJ, Friedrich H (2012) Cost analysis of plug-in hybrid electric vehicles including maintenance & repair costs and resale values. *World Electric Vehicle Journal* 5(4):886–895.
30. Mitropoulos LK, Prevedouros PD, Kopelias P (2017) Total cost of ownership and externalities of conventional, hybrid and electric vehicle in *Transportation Research Procedia*. (Elsevier B.V.), Vol. 24, pp. 267–274.
31. Delucchi MA, Lipman TE (2001) An analysis of the retail and lifecycle cost of battery-powered electric vehicles. *Transportation Research Part D: Transport and Environment* 6(6):371–404.

32. Hummel P, et al. (2017) Electric Car Teardown-Disruption Ahead?, (UBS), Technical Report May.
33. Pavlenko N, Slowik P, Lutsey N (2019) When does electrifying shared mobility make economic sense?, (ICCT), Technical Report January.
34. Michalek JJ, et al. (2011) Valuation of plug-in vehicle life-cycle air emissions and oil displacement benefits. *Proceedings of the National Academy of Sciences of the United States of America* 108(40):16554–8.