

Comment on California Air Resources Board’s Proposed Clean Miles Standard

Matthew Bruchon*, Jeremy Michalek†
Department of Engineering and Public Policy
Carnegie Mellon University

May 2021

Dear Liane Randolph and California Air Resources Board members,

Thank you for the opportunity to provide comments on the proposed Clean Miles Standard. We are researchers at Carnegie Mellon University who study transportation, energy, environment, air quality, public policy, and related issues for shared mobility, ridesourcing, transportation network companies, and other transportation trends and technologies. We have reviewed the proposed rule and put together some comments informed by our research and the broader scientific literature that we hope will be helpful. We will first present some detailed comments on modeling assumptions used to inform the proposed regulation and then comment on the overall policy design.

1 Comments on CARB cost modeling assumptions

CARB has put substantial effort in to informing the proposed rule with a thoughtful and fairly comprehensive qualitative and quantitative analysis. In the spirit of supporting CARB’s efforts to inform policy with data and research, we have identified a set of CARB’s modeling assumptions and choices that differ from some of the established peer-reviewed scientific literature and may affect conclusions and policy implications, and we recommend that CARB consider these in the final version of a model that informs target setting.

Overall, the policy design is focused on setting ZEV targets to avoid imposing costs on drivers. As such, CARB intentionally makes conservative assumptions throughout the analysis in an effort to build confidence that the targets are achievable without imposing costs on drivers. We discuss, in turn, (1) factors that could lead CARB’s analysis to be too conservative, potentially imposing administrative compliance costs to firms, drivers, and government without changing outcomes or improving social welfare and, on the other side, (2) factors that could imply policy costs to drivers that are not modeled in CARB’s analysis.

*PhD Student; Department of Engineering and Public Policy; Carnegie Mellon University; mbruchon@andrew.cmu.edu

†Professor of Engineering and Public Policy; Professor of Mechanical Engineering; Director, Vehicle Electrification Group; Carnegie Mellon University; jmichalek@cmu.edu

1.1 CARB's conservative assumptions

If CARB's goal is to set the strictest GHG and eVMT targets that avoid inducing costs on drivers, the use of many conservative assumptions, when combined, may lead to targets that are so conservative that they may not affect behavior or outcomes, even as they impose administrative compliance costs to industry, drivers and government. We discuss a few of these conservative assumptions.

1.1.1 Purchase subsidy assumptions: As currently structured, the rule assumes that no drivers will receive federal tax benefits for ZEV switching. Federal tax benefits may substantially alter switching cost estimates, particularly if federal tax benefits are extended or expanded. A mechanism for adjusting estimates as federal policy evolves would make the targets for ZEV switching (and gCO₂/PMT by extension) more robust. A similar mechanism to adjust for the CVRP rebates as they evolve over time would similarly improve the targets.

For similar reasons, it may be warranted to estimate a fraction of TNC drivers who are currently eligible for income-based CVRP rebate increases, which currently more than double the credit for those with a household income up to four times the federal poverty level. A substantial share of TNC drivers may fall below that income threshold, and an additional \$2,500 rebate may substantively alter switching benefits.

1.1.2 Gasoline price forecast: Using statewide annual average gasoline prices as estimated by the Energy Information Administration, the ZEV switching cost model's assumption for gasoline price is similar to 2005 fuel prices, which is below the annual average of any year since. We recommend CARB or the California Energy Commission either provide methodology or justification for the forecasts used or adjust the forecast values, perhaps using an average of several independent estimates (or a conservative estimate out of these forecasts).

1.1.3 Charging behavior assumptions: In our understanding, the model assumes that all drivers will pay to have a their house rewired and a home Level 2 battery charger installed, but it also assumes that only half of drivers charge at home, with the remainder using DC fast chargers that are roughly twice as expensive. Given the potential for redundant infrastructure and the high range of BEVs indicated by CARB, this could be excessively conservative. However, we also discuss potential for related un-modeled costs in Section 1.2.

1.1.4 Used vehicle prices: When estimating used vehicle prices, CARB assumes equal depreciation for BEV and combustion vehicles. However, in practice, BEVs tend to depreciate faster, making used BEVs more cost competitive than assumed. See Section 1.3.

1.1.5 One year payback requirement: CARB requires a 1-year payback period before considering a BEV to be a cost-reducing investment. The vehicle asset retains substantial value after one year of TNC service that should be considered. Considering used vehicle asset value may improve the BEV value proposition and more closely represent economics to TNC drivers. See Section 1.3.

1.2 Potential costs to TNC drivers not modeled by CARB

CARB makes a concerted effort to set targets that avoid imposing costs on drivers. If CARB’s goal is to set the strictest targets that avoid inducing costs on drivers, we recommend that CARB consider several costs to drivers beyond those in the current analysis.

If drivers were fully informed, rational, and agreed with CARB’s assumptions, a policy designed to avoid imposing costs on drivers would not change behavior because drivers would adopt ZEVs at levels greater than prescribed even without the policy. Thus, for this policy to be effective (induce changes in driver behavior) without imposing costs on drivers, it must be true that either drivers are (1) not fully informed / rational or (2) do not fully agree with CARB’s assumptions.

For explanation (1), if drivers are not fully informed or not fully rational, it is possible that they are missing opportunities to invest in vehicles that would save them money, and the Clean Miles Standard could impose constraints on drivers that end up increasing driver welfare. The potential for this type of consumer myopia is a major justification for the federal CAFE standards; however, the scientific literature is mixed on findings about whether empirical data back up this assumption (for example, [2, 3]). We are not aware of literature about myopia among TNC drivers in particular, but TNC drivers have a greater incentive to compute costs to inform decisions, and we might expect less myopia under these conditions. Only if drivers are myopic and fail to make decisions that are in their own best interest could the Clean Miles Standard be both effective (change behavior) and reduce driver costs. We recommend that if CARB is assuming consumer myopia, CARB cite and discuss the scientific understanding of consumer myopia, and if CARB is not assuming consumer myopia that CARB explain how the policy can be effective (change behavior) without imposing costs on drivers if drivers are already successful at minimizing their own costs without policy constraints.

For explanation (2), we recommend that CARB explicitly consider factors that do not show up in CARB’s analysis but that may be of legitimate value to drivers in determining whether investment in a BEV is in the driver’s interest, including:

1. Heterogeneity in economics for specific gasoline vehicles owned by specific TNC drivers (e.g.: Prius), in contrast to the generic car/truck lookup table assumptions used by CARB; and
2. Lost income when pulling BEVs out of service to recharge. CARB does discuss that BEVs now have substantial range and access to DC fast charging; however, CARB may also want to consider (1) vehicles that are not able to consistently charge to 100% SOC before every TNC shift; (2) potential for cost of time spent navigating to and from stations and queuing at occupied fast charge stations; and (3) substantial drop in BEV range in cold and hot weather, aggressive or stop and go driving, and as vehicles age [6, 8, 7, 4]. CARB uses a generic barrier cost to account for un-modeled factors like these, but we recommend assessing (a) how the value to drivers of these factors might compare to barrier costs used and (b) whether these costs will ever go to zero, as CARB assumes barrier costs will.
3. BEVs depreciate more quickly than combustion vehicles, imposing additional ownership and operation costs that are larger for BEV drivers (see Section 1.3).

1.3 CARB’s consideration of depreciation

Our research suggests that considering TNC vehicle depreciation and/or TNC vehicle resale value in a used market at the end of the service period can have a substantial effect on the relative economics of BEVs and combustion vehicles for TNC drivers [1], and we recommend that CARB include depreciation or resale value in the analysis.

Our analysis finds that (although the used market for BEVs is evolving) BEVs tend to depreciate at a faster schedule than combustion vehicles. Further, because the State of California and the Federal Government offer purchase incentives for new BEVs, and because new BEVs are (imperfect) substitutes for used BEVs, new BEV incentives put downward pressure on used BEV prices, reducing the resale value of the driver’s BEV asset immediately upon purchase. Both of these factors suggest that BEV used market prices decline more rapidly than equivalent combustion vehicles. This has two implications for CARB’s analysis:

- Depreciation is a higher cost for BEV owners than combustion vehicle owners (or, equivalently, a combustion vehicle asset will retain a larger portion of its value at the end of a TNC service period than a BEV asset).
- For TNC drivers who purchase used vehicles, faster depreciation of BEVs lowers used BEV purchase prices more quickly than used combustion vehicle purchase prices, making BEVs more affordable to buy.

Given the above, we recommend the following:

Adopt powertrain-specific used vehicle prices: We recommend factoring powertrain into depreciated used vehicle market prices. We anticipate this change would encourage BEV switching for TNC drivers who purchase used vehicles.

Model depreciation: We recommend that CARB either include differentiated depreciation costs for BEVs and combustion vehicles as part of the cost of ownership for TNC drivers *or* model resale value at the end of the performance period. Taken alone, this change would capture an additional cost to drivers who buy BEVs because BEVs depreciate faster (however, used vehicle buyers also benefit from lower purchase costs, as above).

Model depreciation dependency on VMT: CARB models depreciation as a function of vehicle age alone. But estimates from Edmunds, Kelly Blue Book, and similar sources show that depreciation depends not only on vehicle age but also strongly depends vehicle miles traveled (VMT) (Figure 1). This matters particularly for high-mileage applications, like TNCs. We have found in our research that including the resale value (or salvage value) of the vehicle asset at the end of the TNC service period can have a substantial effect on the relative economics of BEVs and gasoline vehicles for TNC applications and that this effect differs for high-mileage (“full-time” or “baseload-serving”) TNC vehicles and low-mileage (“part-time” or “peak-period-serving”) TNC vehicles [1]. We recommend modeling depreciation as a function of VMT.

Figure 2, based on analysis in [1]¹ summarizes an example using data-driven models of declining asset value as a function of age and VMT for TNC applications to determine whether BEVs or gasoline vehicles are more cost effective for TNC drivers.

¹See also SI Section C and Table S7 in [1]

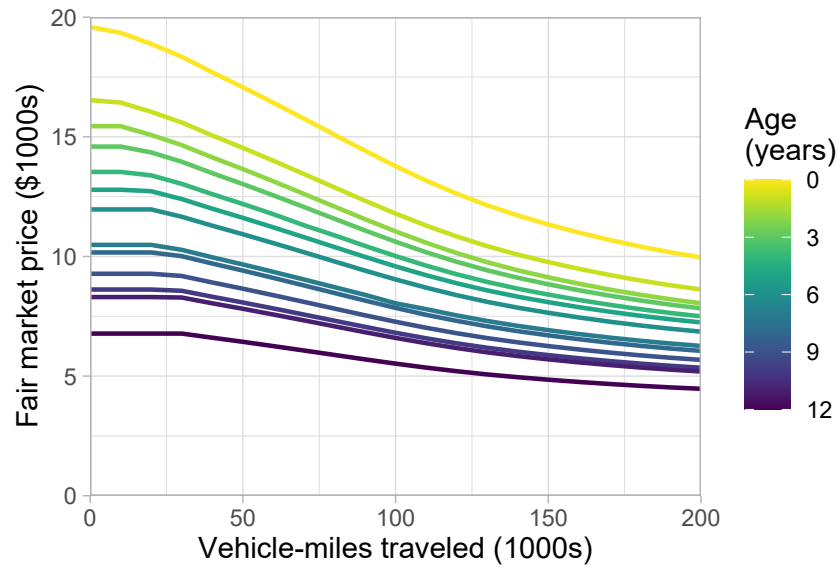


Figure 1: Kelly Blue Book fair market value in 2018 for a Toyota Camry LE four-door sedan of various ages and mileage.

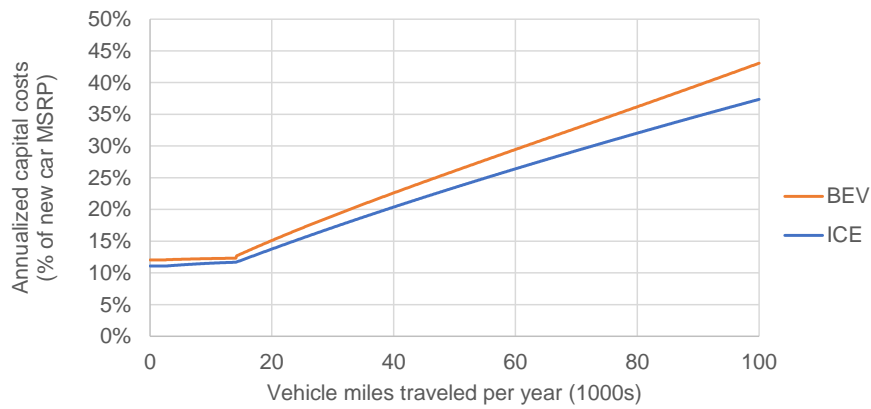


Figure 2: Annualized capital costs (% of new car MSRP) minus discounted future flow from vehicle resale, after depreciation is considered. This value is dependent on annual VMT. This example assumes the vehicle is a newly-purchased TNC vehicle, that decisions are made using a 7% real discount rate (with 2% annual inflation) and it compares a representative BEV (Nissan Leaf) and ICEV (Toyota Camry) assuming the TNC driver sells their car once it reaches 170,000 miles traveled or 12 years of age, whichever happens first (the kink point is where these two conditions intersect).

2 Comments on policy design and goals

We first summarize recommendations for setting targets in CARB’s current policy design and then discuss alternatives to the current approach that have some advantages.

2.1 gCO₂/PMT target design recommendations

2.1.1 Measure vehicle efficiency: The lookup tables provided for CO₂ emission factors (Tables 2 and 3) may be too coarsely defined to serve as an efficient incentive. Ideally, a lookup table derived per vehicle model from fueleconomy.gov or a similar source would more closely match each vehicle’s emissions to their compliance value. This may add to the rule’s complexity—and perhaps small TNCs (less than 5 million trips) should be exempted—but it would provide incentives to reduce GHG emissions through an additional pathway and improve forecasting uncertainty of future emission factors. Measuring actual vehicle on-road efficiency would further provide incentives for efficient driving practices as another pathway to reducing GHGs.

Short of this change, breaking vehicle classes into more granular categories (e.g., to distinguish smaller sedans and hatchbacks from medium-sized crossover SUVs) could improve outcomes.

2.1.2 Measure occupancy factors: We have a similar concern regarding compliance occupancy values (Table 4). When occupancy is assumed, rather than measured, it may create perverse incentives, such as TNCs intentionally pooling rides that involve long trips with large detours or little passenger overlap in order to reduce emissions per PMT on paper, but not in practice. If actual occupancy (which the proposed rule requires TNCs to report annually) were used rather than the assumed values, the gCO₂/PMT metric would offer an efficient incentive to pool when the benefits are greatest (i.e., by pooling rides that require fewer detours and overlap for a substantial portion of the ride durations). It also could offer an incentive to pool more heavily (more than two rides per vehicle), perhaps even by using higher-occupancy shuttle-style vehicles that could substantially reduce actual gCO₂/PMT in urban cores with greater demand density. Instead of serving as the primary means of computing occupancy, assumed lookup-table values could serve as an easier option that small TNCs (less than 5 million trips) could choose to use.

2.1.3 Consider a centralized fund for infrastructure investment: If pedestrian and bicycle infrastructure investment is to be used as an option for CO₂ emissions offsets, CARB may consider whether more equitable outcomes could be obtained by a board charged with considering efficiency, effectiveness and equity in making investment decisions rather than allowing TNCs to selectively invest in neighborhoods of their choosing (for example, TNCs may prioritize projects where bicycle infrastructure would not compete with ridesourcing).

2.1.4 Revise CO₂ credit for transit connections: The current calculation for integrated fare mass transit trips (Equation 5) credits the TNC with offsetting distance equal to the TNC leg of the trip. The current approach may create perverse incentives, encouraging trips that use TNCs to take passengers to the last stop on a transit line before their destination, rather than the nearest stop. A more efficient version of this offset would take

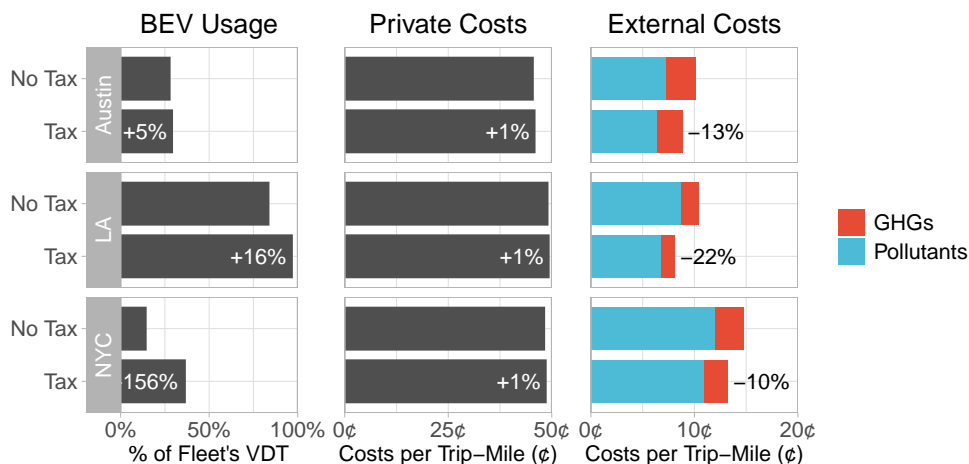


Figure 3: Summary of cost minimizing TNC fleet, its private costs, and the public costs of its emissions when the TNC pays (1) only private costs versus (2) private costs plus the public health and environmental costs of air pollutants (“tax”), from [1]

into account the how many miles were served by transit, not how many were served by the TNC. If TNC’s app handles integrated payment such that the user specifies what their final destination will be (rather than only the “handover point” between TNC and transit), that transit distance estimate could determine zero-emissions miles.

2.2 Recommendations on policy approach

2.2.1 Public costs: Avoiding costs to TNC drivers is an understandable goal. However, it is also true that vehicle travel by TNC drivers causes costs to others, through air emissions, congestion, and other externalities. Our research has found that small increases in private TNC cost can produce larger reductions in public costs, including health and environmental costs of air pollution [1]. In particular, we estimate that if TNCs were to pay for the external health and environmental costs of the air emissions associated with producing and operating their vehicles, they would have incentives to electrify more of their vehicle fleet and reduce air pollution-related health and environmental costs by about \$30 million per year in Los Angeles alone (Figure 3; [1]). These improved outcomes only require an estimated 1% increase in private costs of vehicle acquisition and operation. Constraining policy to avoid imposing any costs on drivers may be missing opportunities to reduce overall costs to society.

It is important to note, though, that most transportation modes produce air emissions, and incomplete regulation requiring only TNCs to pay the cost of their emissions (but not other modes, like personal vehicle travel) could have undesirable implications, like emissions leakage to other modes.

2.2.2 Perverse incentives: The current policy structure has potential for some perverse incentives, including:

1. **Deadheading hot potato:** The proposed standard creates incentives for TNCs to avoid taking credit for driver travel while waiting for ride requests. For example, one possible TNC work-around is for TNCs to ask drivers to log off until they receive a text letting them know that a ride request is ready so that travel while waiting for a ride request is not counted. Additionally, as described on p50 of the proposed standard, many drivers multi-app, logging on to multiple TNC company apps at the same time and switching between them as ride requests come in. TNCs may have incentive to figure out how to assign VMT between ride requests to competitors.
2. **Inefficient ride pooling:** CARB assumes that matched pooled rides always have 2.5 passengers. The use of assumed occupancy for compliance, rather than measured occupancy, may incentivize TNCs to strategically pool rides that involve long distances between passengers (long detours with a single passenger to pick up another passenger) because such a trip will increase the denominator in CARB's compliance calculation of CO2/PMT even though pooling for such a combination of trips may increase VMT overall and contribute to congestion, crash risk, and air emissions. Pooling rides that have origins and destinations that are not generally along the same trajectory can increase VMT, relative to separate vehicles serving these rides, but under the proposed regulation TNCs may have incentive to do this kind of pooling anyway because it helps their compliance calculations.
3. **Inefficient BEV use:** The eVMT targets may incentivize TNCs to increase BEV deadheading because the more BEVs drive (whether with or without a passenger), the more it helps the TNC comply.
4. **PMT:** Regulating emissions rates (CO2/PMT) without direct incentives to reduce PMT may encourage increased PMT through mechanisms like the rebound effect and compliance work-arounds discussed above.

2.2.3 Improving policy efficiency and avoiding perverse incentives: Environmental economics has long established that market-based policies that target the underlying problem's cost directly via externality pricing, rather than setting targets and standards using proxies, can produce more benefits to society at lower cost while avoiding perverse incentives [5].

CARB and the CPUC do not have the authority to impose prices on externalities. SB 1014 gives CARB authority to develop, and CPUC to implement, annual emission-related targets to reduce GHG emissions for TNCs. However, CARB and CPUC could use flexibility in the policy structure to partially approximate the benefits of market-based approaches.

For the Clean Miles Standard, CPUC will need to establish penalties for non-compliance and begin implementation by 2023 (p16). If said penalties for non-compliance are calibrated so that they are comparable to the external costs to society caused by non-compliance, this could potentially offer a pathway to approximate an efficient market-based policy under current regulatory authority. Under this approach, standards could be set more tightly and penalties calibrated to impose external costs of air emissions on the emitters so that non-compliance is an acceptable response to the policy that simply triggers TNC payment for the societal costs of externalities it generates. Such an approach could potentially improve efficiency of the policy, improve flexibility of TNCs and drivers to respond in efficient ways, and reduce potential for perverse incentives.

References

- [1] Matthew B. Bruchon and Jeremy J. Michalek. Effects of air emissions externalities on optimal ride-hailing fleet electrification and operations. *Environmental Science & Technology*, 55(5):3188–3200, mar 2020.
- [2] Meghan R Busse, Christopher R Knittel, and Florian Zettelmeyer. Are consumers myopic? evidence from new and used car purchases. *American Economic Review*, 103(1):220–56, 2013.
- [3] Kenneth Gillingham, Sébastien Houde, and Arthur Van Benthem. Consumer myopia in vehicle purchases: evidence from a natural experiment. Technical report, National Bureau of Economic Research, 2019.
- [4] Orkun Karabasoglu and Jeremy Michalek. Influence of driving patterns on life cycle cost and emissions of hybrid and plug-in electric vehicle powertrains. *Energy policy*, 60:445–461, 2013.
- [5] Nathaniel Keohane and Sheila Olmstead. *Markets and the Environment*. Island Press, 2016.
- [6] Tugce Yuksel, Shawn Litster, Venkatasubramanian Viswanathan, and Jeremy J Michalek. Plug-in hybrid electric vehicle lifepo4 battery life implications of thermal management, driving conditions, and regional climate. *Journal of Power Sources*, 338:49–64, 2017.
- [7] Tugce Yuksel and Jeremy J Michalek. Effects of regional temperature on electric vehicle efficiency, range, and emissions in the united states. *Environmental science & technology*, 49(6):3974–3980, 2015.
- [8] Tugce Yuksel, Mili-Ann M Tamayao, Chris Hendrickson, Inês ML Azevedo, and Jeremy J Michalek. 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, 2016.

Attachment: Copy of Publication [1]

Effects of Air Emission Externalities on Optimal Ridesourcing Fleet Electrification and Operations

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



Cite This: <https://dx.doi.org/10.1021/acs.est.0c05141>



Read Online

ACCESS |



Metrics & More



Article Recommendations



Supporting Information


ABSTRACT: Ridesourcing services from transportation network companies, like Uber and Lyft, serve the fastest growing share of U.S. passenger travel demand.¹ Ridesourcing vehicles' high use intensity is economically attractive for electric vehicles, which typically have lower operating costs and higher capital costs than conventional vehicles. We optimize fleet composition (mix of conventional vehicles (CVs), hybrid electric vehicles (HEVs), and battery electric vehicles (BEVs)) and operations to satisfy demand at minimum cost and compare findings across a wide range of present-day and future scenarios for three cities. In nearly all cases, the optimal fleet includes a mix of technologies, HEVs and BEVs make up the majority of distance traveled, and CVs are used primarily for periods of peak demand (if at all). When life cycle air pollution and greenhouse gas emission externalities are internalized via a Pigovian tax, fleet electrification increases and externalities decrease, suggesting a role for policy. Externality reductions vary from 10% in New York (where externality costs for both gasoline and electricity consumption are relatively high and a Pigovian tax induces a partial shift to BEVs), to 22% in Los Angeles (where high gasoline and low electric grid externalities lead a Pigovian tax to induce a near-complete shift to BEVs).

We optimize fleet purchases and operations for two objectives:

1. Private Costs

2. Private Costs + Emission Costs

- Fleet uses mix of powertrains
- Regional energy costs matter
- Electrification increases
- Emission costs cut by 10–22%



1. BACKGROUND

Passenger cars produce the largest share of greenhouse gas (GHG) emissions from U.S. transportation, which recently surpassed electric power as the country's highest-emitting economic sector.² Passenger cars also emit substantial conventional air pollution, and premature mortality from U.S. air pollution (28% of which results from transportation) is comparable to automobile accident fatalities, with an annual social cost of \$886 billion.³

Ridesourcing 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 ridesourcing services,¹ and by 2016, 15% of intraurban trips in San Francisco were served by Uber and Lyft.⁴

Vehicle electrification has the potential to drastically reduce ridesourcing 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 manufacturing 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 the potential of transportation network company (TNC) fleet electrification 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-electric vehicle (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 ridesourcing 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 nontrivial;^{9,10} and, depending on region, timing, and vehicle design, electric vehicles do not always reduce air pollutant emissions or greenhouse gas emission externalities compared to CVs (with lower-income census block groups more

Received: July 31, 2020

Revised: January 22, 2021

Accepted: January 25, 2021

likely to face increased emission externalities from BEVs).^{11–14} 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 (e.g., peak versus off-peak) being served by different technologies.

We investigate the optimal technology mix and operations of a ridesourcing fleet whose operator has perfect foresight of exogenous (and inflexible) passenger trip requests and total control over fleet acquisition and routing. Centralized control of fleet vehicle choices may represent ridesourcing companies that have owned or leased vehicles in some locales,^{13,16} a future with autonomous vehicle fleets,^{17,18} or vehicles that are purpose-built for ridesourcing fleets.^{19,20} Centralized vehicle routing may become widespread as autonomous vehicle technology advances, whereas today's ridesourcing services only approximate centralized routing via human drivers responding to ride requests and price signals. Also, regulations and incentives that operate at the level of the fleet, rather than the individual driver (e.g., California's under-development Clean Miles Standard, which will regulate fleet-wide annual CO₂ emissions per passenger-mile²¹), increase the role of centralized fleet-wide planning, coordination, and control.

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 emission externality costs in fleet optimization, we assess the degree to which unpriced emission externalities bias fleet outcomes away from socially optimal solutions and consider whether policy intervention may be therefore justified on economic efficiency grounds.

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

Some studies use agent-based modeling (ABM) to explore the operational impacts of homogeneous all-electric fleets. Bauer et al. 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, and Bauer et al. find that fleet-wide coordination of charging would allow BEVs serving demand New York City or San Francisco to meet the same level of service as CVs at a lower cost even if charge networks are relatively sparse.^{22,23} Our fleet differs in its optimization of the fleet mix under different objectives and its consideration of multiple cities. Chen et al. find that electrification can meet ridesourcing demand while barely increasing empty VDT, but only if the fleet size is increased.²⁴

Other studies use ABMs in combination with a second model. The scenarios simulated in Chen et al.²⁴ were used as a case study and defined the inputs for a life cycle assessment framework in Gawron et al., which found that a fleet of electric autonomous taxis could reduce cumulative greenhouse gas emissions by 60% in the period from 2020 to 2050 in the base

case and up to 87% in additional scenarios.²⁵ Sheppard et al. use an ABM to generate simplified operational parameters for a national-scale optimal sizing of vehicles and infrastructure for an all-electric fleet, estimating that 12.5 million vehicles could replace the fleet of 276 million personally owned vehicles.²⁶ Chen and Kockelman incorporate a logit choice model into an ABM to estimate that a shared, autonomous, all-electric vehicle fleet could capture 14–39% of all passenger trips within the Austin, Texas region, depending on pricing.²⁷

Studies employing ABMs use simplifying assumptions or heuristics to model agents' behavior. These 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. find that for vehicle routing problems, optimization coupled with well-designed heuristics increases fleet revenue results by as much as 9% relative to a heuristic alone and that heuristics perform unevenly across problem instances;²⁸ 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 with and without internalized air emission externality costs. 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), they do not consider external costs, and they rarely jointly optimize purchases and routing even in cases when a fixed mixture of powertrains is assumed.^{29,30} The **Supporting Information (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 our study contributes to its problem (applied to an instance of 5000 trips).

In the somewhat-related context of round-trip car-sharing fleets, in which the user pays for short-term rental of a car and drives it themselves, Zoepf finds that BEVs do have a niche to fill, reducing private costs when 20–40% of a gasoline fleet is electrified but increasing private costs beyond that threshold.³¹ For TNC fleets, in the grey literature, a 2019 International Council on Clean Transportation report examined powertrain choice from the perspective of TNC driver costs of vehicle ownership.³² 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. A later analysis by the same group found that a per-trip fee, indexed to tailpipe emissions, between \$0.58 and 1.12 would suffice to make BEVs economically superior to HEVs.³³

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 and (2) applying the model to characterize how the optimal technology mix,




Figure 1. Illustration of the time–space graph showing passenger trips (red), charging arcs (yellow), dispatch arcs (gray), and relocation arcs (blue). Some arcs are omitted for simplicity.

Table 1. Formulation of the FullMILP Optimization Problem

$\begin{aligned} & \text{minimize } \sum_{k \in \mathcal{K}} \kappa_k + \sum_{(i,j) \in \mathcal{A}} c_{k,i,j} a_{k,i,j} + \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}_+} c_t \Delta q_{k,t}^{\text{CHG}} \\ & + \tau \left[\sum_{k \in \mathcal{K}} \delta_k + \sum_{(i,j) \in \mathcal{A}} d_{k,i,j} a_{k,i,j} + \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}_+} d_t \Delta q_{k,t}^{\text{CHG}} \right] \end{aligned}$	<p>(1a) minimize private and external car, fuel, and electricity costs</p>
<p>subject to</p> $\sum_{i \in \mathcal{V}} a_{k,i,j} = \sum_{i \in \mathcal{V}} a_{k,i,i}, \quad \forall k \in \mathcal{K}, j \in \mathcal{V} \setminus \{r, s\}$	<p>(1b) flow is preserved across nodes except source and sink</p>
$\sum_{k \in \mathcal{K}} a_{k,i,j} = n_{i,j}, \quad \forall (i, j) \in \{\mathcal{A}: n_{i,j} > 0\}$	<p>(1c) demand is satisfied</p>
$\sum_{j \in \mathcal{V} \setminus r} a_{k,r,j} = n_k, \quad \forall k \in \mathcal{K}$	<p>(1d) dispatched cars are purchased</p>
$\kappa_k \geq \alpha_{\omega,k}^{\text{COSTS}} n_k + \beta_{\omega,k}^{\text{COSTS}} \sum_{(i,j) \in \mathcal{A}} m_{i,j} a_{k,i,j}, \quad \forall k \in \mathcal{K}, \omega \in \Omega_k$	<p>(1e) per-vehicle private capital costs vary with usage</p>
$\delta_k \geq \alpha_{\omega,k}^{\text{X-COSTS}} n_k + \beta_{\omega,k}^{\text{X-COSTS}} \sum_{(i,j) \in \mathcal{A}} m_{i,j} a_{k,i,j}, \quad \forall k \in \mathcal{K}, \omega \in \Omega_k$	<p>(1f) per-vehicle manufacturing external costs vary with usage</p>
$q_{k,t+1} \leq q_{k,t} + \sum_{(i,j) \in \{\mathcal{A}: t_i=t\}} a_{k,i,j} \Delta q_{k,i,j}^{\text{MAX}}, \quad \forall k \in \mathcal{K}_{\text{B}}, t \in \mathcal{T}_{\text{Q}}$	<p>(1g) BEV charge level is tracked (times with a charger starting timeslot)</p>
$q_{k,t+1} = q_{k,t} + \sum_{(i,j) \in \{\mathcal{A}: t_i=t\}} a_{k,i,j} \Delta q_{k,i,j}^{\text{MAX}}, \quad \forall k \in \mathcal{K}_{\text{B}}, t \in \mathcal{T} \setminus \mathcal{T}_{\text{Q}}$	<p>(1h) BEV charge level is tracked (times with no charger starting timeslot)</p>
$q_{k,t}^{\text{END}} = q_{k,t}^{\text{START}}, \quad \forall k \in \mathcal{K}_{\text{B}}$	<p>(1i) final charge level equals initial charge level</p>
$0 \leq q_{k,t} \leq q_k^{\text{MAX}}, \quad \forall k \in \mathcal{K}_{\text{B}}, t \in \mathcal{T}$	<p>(1j) charge level does not exceed battery capacity</p>
$a_{k,i,j} \in \{0, 1\}, n_k \in \{0, 1\}, \quad \forall k \in \mathcal{K}_{\text{B}}, (i, j) \in \mathcal{A}$	<p>(1k) BEV routing and purchase decisions are binary</p>
$a_{k,i,j} \in \mathbb{Z}_+, n_k \in \mathbb{Z}_+, \quad \forall k \in \mathcal{K} \setminus \mathcal{K}_{\text{B}}, (i, j) \in \mathcal{A}$	<p>(1l) CV HEV routing and purchase decisions are integral</p>
$q_{k,t} \in \mathbb{R}_+, \quad \forall k \in \mathcal{K}_{\text{B}}, t \in \mathcal{T}$	<p>(1m) BEV charge level is always nonnegative</p>
$\kappa_k \in \mathbb{R}_+, \quad \forall k \in \mathcal{K}$	<p>(1n) per-vehicle private capital costs are nonnegative</p>
$\delta_k \in \mathbb{R}_+, \quad \forall k \in \mathcal{K}$	<p>(1o) per-vehicle manufacturing external costs are nonnegative</p>
<p>where</p> $\Delta q_{k,t}^{\text{CHG}} = q_{k,t+1} - q_{k,t} + \sum_{(i,j) \in \{\mathcal{A}: t_i=t, \Delta q_{k,i,j}^{\text{MAX}} < 0\}} a_{k,i,j} \Delta q_{k,i,j}^{\text{MAX}},$	<p>(1p) charging of each BEV at each timestep equals charge change minus routing losses</p>
$\forall k \in \mathcal{K}_{\text{B}}, t \in \mathcal{T}_{\text{Q}}$	

operations, and life cycle air emission externalities of a TNC fleet change across scenarios representing geographic and temporal variation, uncertainty, and the internalization of air emission externalities (as a Pigovian tax passed through to the fleet operator). Our model is also unique in its treatment of vehicle costs, incorporating into the optimization the effect of vehicle

usage on the 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 com-

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}_B	set	battery electric vehicles (subset of \mathcal{K} , indexed individually)
\mathcal{T}	set	all unique arc start and end times
\mathcal{T}_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
$\Delta q_{k,t}^{\text{CHG}}$	variable	energy charged to vehicle k from the grid at time t
κ_k	variable	private acquisition cost for vehicle k
δ_k	variable	externality costs of manufacturing, disposal, and recycling emissions k
τ	parameter	flag controlling whether air emission 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
t^{START}	parameter	earliest time in \mathcal{T}
t^{END}	parameter	latest time in \mathcal{T}
$n_{i,j}$	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
$\Delta q_{k,i,j}^{\text{MAX}}$	parameter	maximum energy change for car k induced by travel on arc (i, j) (positive for charging arcs, negative for all others)
$\alpha_{\omega,k}^{\text{COSTS}}$	parameter	intercept term for line Ω , representing a portion of the convex piecewise linear VDT-dependent capital costs
$\beta_{\omega,k}^{\text{COSTS}}$	parameter	slope term for line Ω , representing a portion of the convex piecewise linear VDT-dependent capital costs
$\alpha_{\omega,k}^{\text{X-COSTS}}$	parameter	intercept term for line Ω , representing a portion of the convex piecewise linear VDT-dependent manufacturing external costs
$\beta_{\omega,k}^{\text{X-COSTS}}$	parameter	slope term for line Ω , representing a portion of the convex piecewise linear VDT-dependent manufacturing external costs

pounds) using reduced complexity models that estimate health costs 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 costs 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. 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) 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.

2.1. Optimization Model. Figure 1 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 (divided into 15 min charging increments) at a charging location while recharging a battery or waiting for the next trip.
- Dispatch arcs: Deadheading from a vehicle’s home base to the first passenger trip request.
- Return arcs: Deadheading from a vehicle’s final passenger trip back to its home base.
- Relocation arcs: Deadheading from the end of one passenger trip to the beginning of a next passenger trip or between passenger trips and recharge locations.

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.

2.2. 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 5, sums the relevant vehicle purchase costs κ_b , gasoline and per mile maintenance costs $c_{k,i,j}$ and time-varying battery charging costs c_t . In cases where air emission externalities are internalized, $\tau = 1$, so the fleet also considers a Pigovian tax on externalities from manufacturing, disposal, and recycling emissions δ_b , tailpipe and fuel refining emissions $d_{k,i,j}$ and grid emissions d_g .

At the core of FullMILP are equations that are standard for many vehicle routing problems. Constraint 6 ensures preservation of flow for each vehicle through the network (forcing vehicles to return to the depot after serving trips), Constraint 7 requires that all passenger trips be satisfied, and Constraint 8 requires that a vehicle must be purchased to be dispatched. The remainder of the formulation is customized for our case.

In our model, annualized mileage determines in what future year each vehicle is sold (either due to age or high mileage), its resale value, and the resulting discounted resale cash flow. Constraints 9–10 model capital costs and manufacturing external. For all vehicles, Constraint 9 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 12–14 manage the BEV charge level. Constraint 11 applies to timesteps at which regular (15 min) interval charging timesteps begin, defines charger usage, and tracks charge level changes. Constraint 12 applies to all other timesteps, at which there is no charging option, so that the charge level is fully determined by traversed arcs' energy requirements. Constraint 14 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 is 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.

2.3. 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 5000 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 the 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 work²⁸ and uses MCF_CarLimit, a customization of the widely known minimum-cost network flow problem.³⁴

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 reoptimize, improving the ShrinkingBattery solution.
4. The 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 a 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 a 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.

2.4. 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 ridesourcing service in Austin, Texas. We use the same set of trips from Austin to also model Los Angeles and New York City (varying private and external costs by region but not travel demand). We extract passenger trip origin and destination, starting and ending timestamps, and distance traveled to define trip arcs $n_{i,j}$. All demand must be satisfied, and passenger pickup times are inflexible. We sample down to 5000 trips using the weekday-season categories shown in Figure S17, plus a separate category for the high-demand days of the South By Southwest Festival. This sample size equates to a fleet size ranging from 37 to 39 vehicles in the base case (37–44 vehicles across all sensitivity cases), depending on the optimal technology mix. 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 5 min. 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/h, respectively).³⁵

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

regression³⁶ on the RideAustin trips. This method and its implications are described further in the SI.

Relocations from the prior trip to the next trip were disallowed if the actual time gap (from the first trip's end to the second trip's start) was shorter than the estimated relocation duration or longer than 30 min. For tractability, vehicles may chain trips more than 30 min apart but must park at the central charge station depot between those trips. When the estimated duration is shorter than the time gap between trips, we assume that the vehicle travels at the estimated speed for the estimated trip duration, then idles for the remainder of the 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 that the vehicle parks immediately at the station and departs the station as late as possible.

2.5. Vehicle and Charger Technology. We model a typical present-day ridesourcing 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. Figure S25 shows that its efficiency and range are representative of model year 2018 BEVs excluding Teslas (likely 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. In practical contexts, optimal sizing and siting of charge capacity is a challenging problem that requires planning and investment, and it would add a great deal of complexity to our optimization. However, because our results do not show a substantial increase in VDT from BEVs routing to and from charging stations—perhaps due to perfect demand information—we do not consider sensitivity cases with more charging stations or a different charge station location. 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 min (linearized to a rate of 35.2 kWh/h 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/gallon (mpg), and highway efficiency of 30.9 mpg. The hypothetical hybrid version of the Soul's parameters is 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;³⁷ 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/gal).^{38,39} These time-invariant energy prices are shown in Table S8.

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 , as shown in eq 2

$$f_{\text{CR}}(r, N) = \frac{r(1+r)^{N-1}}{(1+r)^N - 1} \quad (2)$$

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 a noninteger) 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

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

where $N = \min\left(N_{\text{MAX}}, \frac{d_{\text{MAX}}}{d}\right)$ 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 $N_{\text{MAX}} = 12$ years, and a maximum VDT of $d_{\text{MAX}} = 170\,000$ miles, based on Argonne National Laboratory's Greenhouse gases, Regulated Emissions, and Energy use in Transportation Model (GREET).⁴⁰ 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.

2.6. Air Emission Externality Costs. In scenarios where external costs of emissions are considered, air emission externalities from the manufacturing stage are added as a Pigovian tax on vehicle investments

$$T_{\text{MFG}} = \sum_{i \in \mathcal{P}} \gamma_i^{\text{MFG}} c_i^{\text{P}} \quad (4)$$

where \mathcal{P} is the set of pollutants considered, γ_i^{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 CO₂, methane, and N₂O; we consider health costs from PM_{2.5}, SO_x, 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 CO₂ equivalent estimated by the Interagency Working Group on the Social Cost of Carbon.⁴¹ For conventional air pollutants, external costs depend on emission location, and we use the AP3 model⁴² to compute and monetize estimated health costs associated with these emissions. AP3 is one of several reduced complexity models that estimate the 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,⁴⁰ 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.

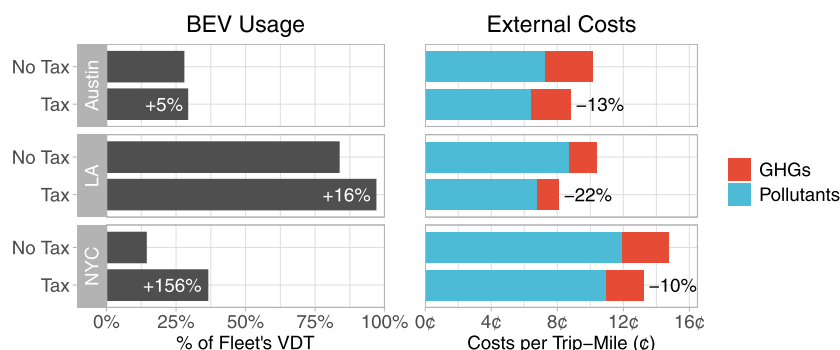


Figure 2. Summary of changes to the optimal ridesourcing fleet when air emission externalities are internalized, including share of fleet-wide vehicle-distance traveled (VDT) from BEVs (left) and total air emission externality costs per trip-mile (right) in three cities for the optimal fleet technology mix and routing to serve exogenous travel demand. Each measure's relative change induced by a Pigovian tax (expressed as a percentage of the "no-tax" case) is annotated. Assumed private and external costs of energy inputs vary by city, as described in Sections 2.5 and 2.6. All cases use a 7% real private firm discount rate, no labor costs, the vehicles described in Section 2.5 including the 2019 Kia Soul BEV, \$50/tonne CO₂ externality price, the AP3 external cost model, \$9.41 million (2018) value of statistical life, and the Pope et al.⁴⁷ concentration–response function. Results using alternative assumptions are summarized in Section 3.3.

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.

$$T_{\text{OP}} = \sum_{i \in \mathcal{P}} \gamma_i^{\text{OP}} c_i^{\text{P}} \quad (5)$$

As with manufacturing emissions, we use the social cost of carbon and AP3 to estimate external costs per unit of pollutant emitted from vehicle operations. We adopt GREET tailpipe and upstream estimates of emissions per gallon of gasoline consumed and compute emissions 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, there is a body of literature estimating the time-varying marginal grid emissions from end uses, including studies by Graff Zivin, Kotchen, and Mansur and Siler-Evans et al.^{43,44} that use regression approaches. We use the Siler-Evans et al. methodology, recalculated on recent data and averaged by season and hour of day in 2017 for each analysis city's eGrid subregion (ERCOT in the case of Austin).⁴⁵ 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 costs factors of refinery counties in the city's region. This approach is described further in the SI, and values used as inputs are given in Tables S2 (time-invariant gasoline externalities) and S3 (average values of time-variant grid externalities).

For emission external costs from all sources, we use the AP3 external cost model, a value of statistical life of \$9.41 million (2018), a carbon price of \$50/tonne, and the Pope et al. concentration–response relationship.⁴⁷

3. RESULTS AND DISCUSSION

Across a wide 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 emission externality

costs, and we compare resulting outcomes of policy interest. The second case assumes 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. Each test case has the same total trip miles, since demand is 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 life cycle air emission 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/tonne CO₂ externality valuation, the AP3 model of conventional air pollution emission mortality effects, \$9.41 million value of statistical life, and the Pope et al.⁴⁷ 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 Materials and Methods.

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

3.1. Impact of a Pigovian Tax. Figure 2 summarizes key cost outcomes in each city when optimized with and without a Pigovian tax on air emission externalities. Across cities in our base case, with no Pigovian tax, private costs range from 45.5 to 49.0¢ and external costs range from 10.1 to 14.8¢ per trip-mile (that is, total annualized life cycle costs divided by the number of annualized miles of passenger trips served). Depending on the variability of regional costs, a tax leads the fleet to increase its usage of BEVs by 5–156% and dispatch these vehicles in a manner that reduces emission externalities per trip-mile by 10–22%. In absolute terms (x -axis of Figure 2), these reductions range from 1.3 to 2.3¢ per trip-mile. These values are broken down in greater detail in Figures S1–S4, SI.

External cost reductions are greatest in percentage and absolute terms in Los Angeles, where fuel emission externalities are high and electricity generation externalities are low relative to the other cities modeled (a larger relative difference for

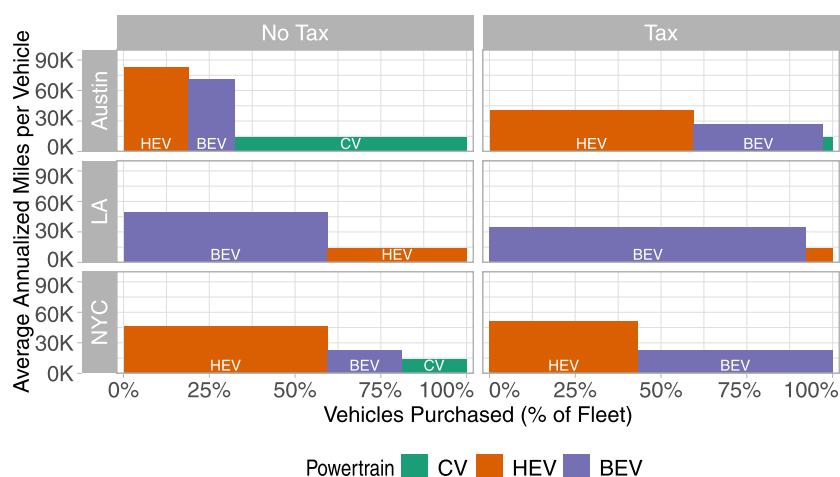


Figure 3. Vehicle purchases (*x*-axis) and average utilization (*y*-axis) by powertrain type for cost-minimizing fleets when excluding (left) and including (right) a Pigovian tax on air emissions. Private and external costs of energy inputs vary across cities, as described in Sections 2.5 and 2.6. All cases use a 7% real private firm discount rate, no labor costs, the vehicles described in Section 2.5 including the 2019 Kia Soul BEV, \$50/tonne CO₂ externality price, the AP3 air emission external cost model, \$9.41 million (2018) value of statistical life, and the Pope et al.⁴⁷ concentration–response function.

criteria pollutant external costs than for GHGs). In percentage terms, they are smallest in New York City, where the external costs of electricity generation are highest of the three cities. Austin sees the largest increase in BEV usage, partially because of lower gas prices that lead a private cost-minimizing fleet to use many CVs and few BEVs. However, due to lower health external costs per unit of tailpipe emissions and a less “clean” grid than Los Angeles, Austin’s external cost reductions fall between the other two cities’ in percentage terms and are smallest in absolute terms.

To put these per trip-mile results in context, a recent Fehr and Peers consulting report estimated that Uber and Lyft drive 104 million monthly trip miles in Los Angeles.⁴⁸ Multiplying those trip miles by the 2.3¢ per trip-mile decrease in externalities, we can roughly estimate external cost reductions of \$29 million per year in Los Angeles (\$24 million in reduced criteria pollutant emissions and the remainder in reduced GHG emissions).

As shown in Figure S2, these external cost reductions occur alongside private cost increases up to 1% in our base case (increases no higher than 0.4¢ per trip-mile). The net effect of these cost changes is a reduction in overall social costs (private costs plus external costs) ranging from 2 to 3% (0.9–2.0¢ per trip-mile). While this net effect is small in relative terms, 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 effects are not uniform across life cycle stages. Figure S5 shows that across analysis regions, as more BEVs are used, per trip-mile manufacturing external costs increase by 6–11% (0.3–0.5¢), almost entirely due to criteria pollutants. Tailpipe and refining external costs drop by 17–80% (0.7–2.9¢) as internal combustion engines are used less. In New York City, where BEV usage is low without a Pigovian tax and the increase is largest in relative terms, grid external costs increase by a factor of four (0.9¢); in Austin and Los Angeles, where the shift is less drastic in relative terms, changes in charge scheduling offset emissions from increased grid energy usage.

These effects also vary by type of emissions. The share of per trip-mile external cost reductions attributed to reduced criteria pollutant emissions ranges from 63 to 85% (0.8–2.0¢), with the

remaining 15 to 37% coming from reduced GHG emissions (0.3–0.5¢). Criteria pollutant external cost reductions range from 8 to 11%, while greenhouse gas external cost reductions range from 16 to 19%.

These external cost reductions are accomplished in each city not only by shifting VDT away from gasoline usage (in CVs and HEVs) and toward electricity usage (in BEVs), but also by a corresponding change in vehicle purchases. Figure 3 illustrates for each city, with and without the Pigovian tax, the share of vehicle purchases for each powertrain (out of an optimal fleet size ranging from 37 to 39 vehicles in the base case) and the annual miles driven per car 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 used infrequently, primarily during periods of high demand, while HEVs serve as “baseload” supply and are responsible for a plurality of total miles driven. When air emission externalities are internalized, the fleet uses HEVs to serve baseload and BEVs for nearly all remaining trips, almost 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 relatively high externalities of electricity generation, HEVs are still used as the baseload.

Across the three cities, the number of BEVs in the optimal fleet increases by 63–180% when a Pigovian tax is imposed on the fleet, and BEVs’ total vehicle-distance traveled increases by 5–




Figure 4. 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 emission 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, the vehicles described in Section 2.5 including the 2019 Kia Soul BEV, \$50/tonne CO₂ externality price, the AP3 external cost model, \$9.41 million (2018) value of statistical life, and the Pope et al.⁴⁷ concentration–response function.

156%. 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.

3.2. 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 emission externalities it produces. Figure 4 illustrates each cost component for four fleet configurations: (1) an all-CV fleet optimized for private costs with no Pigovian tax on emission externalities (the single-technology fleet option that is 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 emission 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 net capital costs).

In all three cities, an all-CV fleet is suboptimal enough that many of the external cost reductions seen from a Pigovian tax are also achieved simply by having the fleet optimally choose powertrains (while still providing it with the full foresight to know what that optimal mix is). Relative to an all-CV fleet, a fleet optimized for private costs reduces private costs by 5–14% and, in doing so, also reduces emission externalities by 14–66%. 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 1–4% and the mixed fleet with a Pigovian tax reduces social costs by 1–4%.

Unlike all-CV and all-HEV fleets, internalizing an all-BEV fleet’s emission externalities can shift charging to lower-polluting times of the day to reduce externalities (assuming perfect day-ahead information regarding external costs of the grid’s marginal

generator). The magnitude of this reduction ranges from 4 to 6% depending on the scenario. All-BEV fleets are also slightly larger than other fleets, due to the need for some portion of the fleet to recharge during high-demand hours. These results are provided in the SI.

3.3. Sensitivity to Model Inputs. Here, we briefly summarize findings across additional test cases. In Figures S7–S16, we provide more exhaustive results from all sensitivity cases.

3.3.1. External Cost Model and Social Cost of Carbon. Our base case assumes a \$50/tonne CO₂-equivalent externality valuation and uses the AP3 reduced complexity model to estimate health costs from criteria air pollutants. Using a very high CO₂ 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 under a Pigovian tax, reducing externalities as much as 39% (in Austin).

We also consider EASIUR and InMAP, two alternative reduced complexity models downloaded from a Center for Air, Climate, and Energy Solutions Database.^{46,49,50} Using either model rather than AP3 in Austin (where those models’ marginal grid external cost estimates are each roughly half of AP3’s) leads to substantially increased BEV uptake in Pigovian tax scenarios. In tax scenarios for Los Angeles, BEV usage is nearly maximized with all three models. In New York City, EASIUR results in similar outcomes as AP3, but InMAP leads to estimated externality reductions of 75% due to near-total electrification with a tax. This change is primarily driven by a large difference in county-level tailpipe external cost estimates for Manhattan (where InMAP’s are around 5 times greater than AP3’s and 3 times greater than EASIUR’s). These models’ structural differences, and regional variations in differences between their external cost estimates, are examined systematically in Gilmore et al.⁵¹

3.3.2. Discount Rate, Resale, and Labor. Our base case assumes that the fleet pays drivers no hourly wages (instead effectively assuming either a flat percentage of fare or driverless cars) and uses 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 is described in Section 2.5, eq 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 no more than 22% of VDT and BEVs still serve 23–97% 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 the resale value is removed, the effective purchase costs of CVs 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 27% in the base case (7% discount rate) to 26% 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 26 (in New York City) to 52 (in Los Angeles) percentage points in no-tax scenarios. This demonstrates that charging may not only increase planning complexity but also impose labor costs that change a fleet's optimal strategy.

3.3.3. Battery Capacity and Cost. Our base case uses the 2019 Kia Soul with a retail price of \$33 950 and a 30 kWh battery. Rather than model an explicit cost per kWh, we consider a sensitivity case in which the BEV's sticker price is reduced to \$28 950 (a cost reduction of \$167/kWh if all price reductions are attributed to lower battery costs). In that case, the private cost-minimizing fleet would be the majority BEV in each city (in terms of purchases and VDT) and a Pigovian tax would lead BEVs to serve 60–96% 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.

3.3.4. Electricity Rates. Our base case uses average retail electricity rates from 2017. In the future, higher rates could result from large amounts of demand induced from economy-wide BEV charging, which may alter grid dispatch and require additional generating capacity. For sensitivity cases, we focused on Austin, which has the most evenly mixed fleet in the base case, and varied its electricity rates from 5.5¢ (half of the base case value) to 21.8¢ (double the base case value) per kilowatt hour. Across all price ranges, usage of CVs does not change much: they make up 68–73% of the fleet with no-tax (but primarily serve peak demand, around 30% of total VDT) and are nearly eliminated with a tax. However, the portion of the remaining demand served by BEVs varies by electricity price. When rates are doubled, they are used only in the tax scenario (for 17% of VDT versus 29% in the base case); when rates are halved, at least 68% of VDT are served by BEVs with or without a tax. Since this affects scenarios with and without a tax in similar ways, the range of reductions in external costs induced by a tax is relatively consistent, from 10% (with rates doubled) to 15% (with rates halved). Because Austin is the least favorable for BEVs in our base case, it loosely corresponds to a pessimistic lower bound on BEV usage under increased electricity rates.

3.3.5. Marginal External Costs from Electricity Generation. Our base case (described in Section 2.6) uses marginal generation estimates from 2017 for each eGrid subregion. It is unclear how external cost values from marginal generation may change in the future: it is possible that in some regions, demand induced from economywide BEV charging could shift less efficient coal plants to the margin, but it is also possible that long-term increases in renewable energy may result in some hours of the day having zero marginal emissions. Rather than explicitly model these possibilities, in each region, we consider three additional sensitivity cases with three time-invariant values for marginal external costs: the region's highest-external cost marginal emissions from 2017 (the most damaging plant is always on the margin), the region's lowest-external-cost marginal emissions from 2017 (the least damaging plant is always on the margin), and 50% of the region's lowest-external-cost marginal emissions from 2017 (on average, renewables and the least damaging plant are each on the margin half the time). For Austin's lowest-external-cost case, BEVs serve 79% of VDT with a tax; in the highest-external-cost case, they serve only 18%. External cost estimates shift with generation assumptions across tax and no-tax scenarios, and the external cost reductions induced by a tax vary in the range from 9 to 19% (0.8–1.8¢ per trip-mile). The trend is less pronounced in New York City, where BEV VDT in tax cases ranges from 31 to 42% and external cost reductions range from 7 to 15% (1.6–2.3¢). In Los Angeles, where BEVs serve a majority of VDT across all cases, external cost reductions shift even less, ranging from 17 to 22% (1.9–2.3¢).

Detailed results from sensitivity analyses, along with an expanded discussion, are available in the SI.

3.4. Discussion. Across a wide range of scenarios, our results consistently suggest that internalizing air emission 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 emission externality costs (by 10–22% in the base case and 4–75% across sensitivity cases, depending on the city) and lower social costs (by 2–3% in the base case and 0–18% across sensitivity cases, depending on the city). This suggests a potential role for policy because when emission externalities are unpriced, firms have incentives to lower private cost in ways that increase air emissions, implement a lower degree of electrification, and charge BEVs when the grid is less clean than socially optimal. While the change in social cost is fairly small across most of the scenarios examined, the change in who bears the cost (private versus external costs) can be significant—as estimated above for Los Angeles, as high as \$29 million of annual environmental and health outcomes.

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.

It is worth noting that unlike private vehicles, ridesourcing fleet vehicles spend a substantial portion of time deadheading with no passengers while they wait for their next ride request and

travel to its pickup location. Across test cases, even with our model's assumption of perfect information, around 47% of total distance traveled result from these empty miles (due in part to relatively low demand density in the RideAustin dataset); this implies that a similar share of external costs is due to deadheading—a largely unavoidable aspect of ridesourcing—regardless of powertrain decisions. Because demand is exogenous in our model, the Pigovian tax (and the resulting electrification) does little to reduce the degree of deadheading.

These results should be interpreted in context. Our model is relatively detailed in its treatment of supply-side investment and operation costs and constraints, but it considers a single ridesourcing firm with perfect information and full control of fleet acquisition and operation that must satisfy all demand with inflexible pickup times. In practice, current ridesourcing fleets in the U.S. are staffed by workers who choose their own vehicles, which often serve dual uses as personal vehicles and choose when to work in response to incentives. Where vehicles are purchased by each worker, it is unclear what options (e.g., a powertrain externality based driver incentive program) may be most viable for fleet owners to induce these shifts. 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 fleet centrally owned and routed by the ridesourcing firm (e.g., a fleet of autonomous vehicles; a fleet owned and leased to drivers for TNC work) or perhaps one responding to policies requiring greater fleet-wide coordination and optimization.²¹

Our exclusion of vehicles' outside value for dual uses overestimates the extent to which CVs (which have lower capital costs) reduce costs for peak demand hours relative to BEVs and HEVs. However, this overestimation is partially mitigated by our formulation of endogenous capital costs, which lowers them for less heavily used cars. Accordingly, excluding dual uses affects the composition of vehicle purchases, particularly for low-use vehicles, but its effect on the vehicle distance traveled by each vehicle type is smaller by comparison.

Assuming perfect information and control may overestimate the fleet's ability to opportunistically schedule battery charging around gaps in demand and fluctuations in marginal grid emissions, overestimating the number of trips each BEV may be able to serve. Ridesourcing services also need not meet all demand at the exact start and end time they were served in the RideAustin data. If we allowed flexible time windows for passenger pickup and dropoff time (perhaps with a cost for additional waiting time), the fleet could improve operational efficiency and the optimal fleet composition could potentially change.⁵²

We provide results for three cities with varying private and external assumptions but use RideAustin data from 2016 to 2017 in all scenarios. Because this study does not have accurate and current inputs for each city's TNC travel demand, it misses differences resulting from the urban form that may alter the trip and relocation distances and speeds across regions (e.g., compact versus sprawled development, gridded versus irregular design, congested versus free-flowing travel, and low-speed neighborhood streets versus urban highways). This may change optimal fleets. For example, where trips are more stop-and-go, BEVs may be more optimal; where trips are longer, BEVs may face more difficult charging constraints.

We do not consider distributional impacts of ridesourcing fleet externalities, but rather optimize total social costs across disparate regions. Approaches considering equitable outcomes could, for example, require that no region may see external costs increase by more than some margin. By shifting emissions from the tailpipe to the grid, fleet electrification could increase health impacts of air emissions in areas outside the city even while reducing total air emission health impacts.¹³ Distributional impacts within a status-quo city also warrant further attention; ridesourcing fleets may be used by relatively affluent people but may impose disproportionately large health costs on populations less likely to use the services.

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 the implications of unpriced externalities in technology choice and operations.

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

■ ASSOCIATED CONTENT

SI Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.est.0c05141>.

Detailed results for test cases, summary results for additional sensitivity cases, discussion of model limitations, detailed methods for dataset preparation and estimation of uncertain inputs, and discussion of optimization methods and related literature (PDF)

■ AUTHOR INFORMATION

Corresponding Author

Jeremy J. Michalek – *Department of Engineering & Public Policy and Department of Mechanical Engineering, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, United States*; orcid.org/0000-0001-7678-8197;
Email: jmichalek@andrew.cmu.edu

Authors

Matthew B. Bruchon – *Department of Engineering & Public Policy, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, United States*; orcid.org/0000-0003-4321-1345

Inês L. Azevedo – *Department of Energy Resources Engineering, Woods Institute for the Environment, and Precourt Institute for Energy, Stanford University, Stanford, California 94305, United States*; orcid.org/0000-0002-4755-8656

Complete contact information is available at:
<https://pubs.acs.org/10.1021/acs.est.0c05141>

Notes

The authors declare no competing financial interest. Data and code that support the findings of this study are available at <https://github.com/mbruchon/FleetElectrification>.

■ ACKNOWLEDGMENTS

The authors thank Julian Marshall for providing guidance on appropriate usage and interpretation of reduced complexity air quality models. This work was supported by the Carnegie Mellon Center for Climate and Energy Decision Making through a cooperative agreement (Awards SES-0949710 and SES-1463492) between the National Science Foundation and Carnegie Mellon University. This work was also developed as

part of the Center for Clean Air Climate Solution (CACES), which was supported under Assistance Agreement no. R835873 awarded by the U.S. Environmental Protection Agency. It has not been formally reviewed by EPA. The views expressed in this document are solely those of authors and do not necessarily reflect those of the Agency. EPA does not endorse any products or commercial services mentioned in this publication. This work was also supported by the Department of Engineering & Public Policy at Carnegie Mellon University. This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number ACI-1548562. Specifically, it used the Bridges system, which is supported by NSF award number ACI-1445606, at the Pittsburgh Supercomputing Center (PSC).^{53,54}

REFERENCES

- (1) Conway, M.; Salon, D.; King, D. Trends in Taxi Use and the Advent of Ridehailing, 1995-2017: Evidence from the US National Household Travel Survey. *Urban Sci.* **2018**, *2*, No. 79.
- (2) United States Environmental Protection Agency. *Inventory of U.S. Greenhouse Gas Emissions and Sinks*; United States Environmental Protection Agency, Office of Air and Radiation, Climate Change Division, 2018; pp 1–505.
- (3) Goodkind, A. L.; Tessum, C. W.; Coggins, J. S.; Hill, J. D.; Marshall, J. D. Fine-scale damage estimates of particulate matter air pollution reveal opportunities for location-specific mitigation of emissions. *Proc. Natl. Acad. Sci. U.S.A.* **2019**, *116*, 8775–8780.
- (4) San Francisco County Transportation Authority. TNCs Today: A Profile of San Francisco Transportation Network Company Activity, 2017. <https://www.sfcta.org/emerging-mobility/tncs-today>.
- (5) *Strengthening and Implementing the Global Response*; IPCC, 2018; Chapter 4.
- (6) George, S. R.; Zafar, M. Electrifying the Ride-Sourcing Sector in California, 2018. <http://www.cpuc.ca.gov/General.aspx?id/6442457050>.
- (7) Uber. The Clean Air Plan: Helping London Go Electric, 2018. <https://www.uber.com/en-GB/blog/a-cleaner-future-for-our-cities/>.
- (8) Fulton, L.; Mason, J.; Meroux, D. *Three Revolutions in Urban Transportation*; ITDP, 2017; p 38.
- (9) Tessum, C. W.; Hill, J. D.; Marshall, J. D. Life cycle air quality impacts of conventional and alternative light-duty transportation in the United States. *Proc. Natl. Acad. Sci. U.S.A.* **2014**, *111*, 18490–18495.
- (10) Michalek, J. J.; Chester, M.; Jaramillo, P.; Samaras, C.; Shiao, C.-S. N.; Lave, L. B. Valuation of plug-in vehicle life-cycle air emissions and oil displacement benefits. *Proc. Natl. Acad. Sci. U.S.A.* **2011**, *108*, 16554–16558.
- (11) Weis, A.; Jaramillo, P.; Michalek, J. Consequential life cycle air emissions externalities for plug-in electric vehicles in the PJM interconnection. *Environ. Res. Lett.* **2016**, *11*, No. 024009.
- (12) Yuksel, T.; Tamayao, M.-A. M.; Hendrickson, C.; Azevedo, I. M. L.; Michalek, J. J. 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. *Environ. Res. Lett.* **2016**, *11*, No. 044007.
- (13) Holland, S. P.; Mansur, E. T.; Muller, N. Z.; Yates, A. J. Distributional Effects of Air Pollution from Electric Vehicle Adoption. *J. Assoc. Environ. Resour. Econ.* **2019**, *6*, S65–S94.
- (14) Holland, S. P.; Mansur, E. T.; Muller, N. Z.; Yates, A. J. Are there environmental benefits from driving electric vehicles? the importance of local factors. *Am. Econ. Rev.* **2016**, *106*, 3700–3729.
- (15) Corporation, T. M. Toyota develops “Total-care Service” Designed for Ridesourcing Companies, 2018. <https://global.toyota/en/newsroom/corporate/25737570.html>.
- (16) Lyft, Inc. Lyft Express Drive, 2020. <https://www.lyft.com/expressdrive>.
- (17) Lyft, Inc. Lyft Level 5 Self-Driving Division, 2020. <https://level5.lyft.com/>.
- (18) Uber Technologies, Inc., Uber Advanced Technologies Group, 2020. <http://www.uber.com/us/en/atg/>.
- (19) Pram, D. GM’s Cruise Shows Off Its Self-Driving Origin Shuttle with No Room for a Driver, 2020. <https://www.bloomberg.com/news/articles/2020-01-21/gm-s-cruise-to-show-first-car-built-to-operate-without-a-driver>.
- (20) Sun, Y.; Shirouzu, N. China’s Didi Sets Up Joint Ventures with BAIC Unit, Volkswagen, 2019. <https://www.reuters.com/article/us-china-didi-chuxing/chinas-didi-sets-up-joint-ventures-with-baic-unit-volkswagen-idUSKCN1PM079>.
- (21) California State Senate, Senate Bill #1014, 2018. https://leginfo.ca.gov/faces/billTextClient.xhtml?bill_id=201720180SB1014.
- (22) Bauer, G. S.; Greenblatt, J. B.; Gerke, B. F. Cost, Energy, and Environmental Impact of Automated Electric Taxi Fleets in Manhattan. *Environ. Sci. Technol.* **2018**, *52*, 4920–4928.
- (23) Bauer, G. S.; Phadke, A.; Greenblatt, J. B.; Rajagopal, D. Electrifying urban ridesourcing fleets at no added cost through efficient use of charging infrastructure. *Transp. Res. Part C: Emerging Technol.* **2019**, *105*, 385–404.
- (24) Chen, T. D.; Kockelman, K. M.; Hanna, J. P. Operations of a shared, autonomous, electric vehicle fleet: Implications of vehicle & charging infrastructure decisions. *Transp. Res. Part A: Policy Pract.* **2016**, *94*, 243–254.
- (25) Gawron, J. H.; Keoleian, G. A.; De Kleine, R. D.; Wallington, T. J.; Kim, H. C. Deep decarbonization from electrified autonomous taxi fleets: Life cycle assessment and case study in Austin, TX. *Transp. Res. Part D: Transp. Environ.* **2019**, *73*, 130–141.
- (26) Sheppard, C. J. R.; Bauer, G. S.; Gerke, B. F.; Greenblatt, J. B.; Jenn, A. T.; Gopal, A. R. Joint Optimization Scheme for the Planning and Operations of Shared Autonomous Electric Vehicle Fleets Serving Mobility on Demand. *Transp. Res. Rec.* **2019**, *2673*, 579–597.
- (27) Chen, T. D.; Kockelman, K. M. Management of a Shared Autonomous Electric Vehicle Fleet. *Transp. Res. Rec.* **2016**, *2572*, 37–46.
- (28) Bertsimas, D.; Jaillet, P.; Martin, S. Online Vehicle Routing: The Edge of Optimization in Large-Scale Applications. *Oper. Res.* **2017**, *1–48*.
- (29) Pourazarm, S.; Cassandras, C. G.; Wang, T. Optimal routing and charging of energy-limited vehicles in traffic networks. *Int. J. Robust Nonlinear Control* **2016**, *26*, 1325–1350.
- (30) Sassi, O.; Cherif, W. R.; Oulamara, A. Vehicle Routing Problem with Mixed fleet of conventional and heterogenous electric vehicles and time dependent charging costs. *Int. J. Math., Comput., Phys., Electr. Comput. Eng.* **2015**, *9*, 167–177.
- (31) Zoepf, S. M. Plug-in Vehicles and Carsharing: User Preferences, Energy Consumption and Potential for Growth. Ph.D. Thesis, Massachusetts Institute of Technology, 2015.
- (32) Pavlenko, N.; Slowik, P.; Lutsey, N. *When Does Electrifying Shared Mobility Make Economic Sense*; International Council on Clean Transportation, 2019; pp 1–15.
- (33) Slowik, P.; Wappelhorst, S.; Lutsey, N. *How Can Taxes and Fees on Ride-Hailing Fleets Steer Them to Electrify*; International Council on Clean Transportation, 2019; pp 1–15.
- (34) Klein, M. A. A Primal Method for Minimal Cost Flows with Applications to the Assignment and Transportation Problems. *Manage. Sci.* **1967**, *14*, 205–220.
- (35) U.S. Department of Energy Office of Energy Efficiency and U.S. Environmental Protection Agency. Detailed Test Information, 2018. https://www.fueleconomy.gov/feg/fe_test_schedules.shtml.
- (36) Wang, H.; Kuo, Y.-H.; Kifer, D.; Li, Z. In *A Simple Baseline for Travel Time Estimation Using Large-Scale Trip Data*, Proceedings of the 24th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems—GIS ’16. New York, New York, 2016; pp 1–4.
- (37) U.S. Environmental Protection Agency. Download Fuel Economy Data: 2018 Datafile, 2018. <https://www.fueleconomy.gov/feg/download.shtml>.

- (38) U.S. Energy Information Administration. Retail Gasoline and Diesel Prices: Regular—Conventional Areas, Annual Period, 2018. https://www.eia.gov/dnav/pet/pet_pri_gnd_a_epmru_pte_dpgal_a.htm.
- (39) U.S. Energy Information Administration. Average Retail Price of Electricity, Texas, Annual, 2018. <https://www.eia.gov/electricity/data/browser/>.
- (40) Argonne National Laboratory. Greenhouse Gases, Regulated Emissions, and Energy Use in Transportation Model, 2019. <https://greet.es.anl.gov/greet.models>.
- (41) *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; 2016, pp 1–21.
- (42) Muller, N. Z.; Mendelsohn, R. *The Air Pollution Emission Experiments and Policy Analysis Model (APEEP): Technical Appendix*; Yale University, 2006; Vol. 1.
- (43) Graff Zivin, J. S.; Kotchen, M. J.; Mansur, E. T. Spatial and temporal heterogeneity of marginal emissions: Implications for electric cars and other electricity-shifting policies. *J. Econ. Behav. Org.* **2014**, *107*, 248–268.
- (44) Siler-Evans, K.; Azevedo, I. L.; Morgan, M. G. Marginal Emissions Factors for the U.S. Electricity System. *Environ. Sci. Technol.* **2012**, *46*, 4742–4748.
- (45) Azevedo, I. L.; Donti, P. L.; Horner, N.; Siler-Evans, K.; Vaishnav, P. T. *Electricity Marginal Factors Estimates*; Center For Climate and Energy Decision Making, Carnegie Mellon University: Pittsburgh, 2021. <https://cedm.shinyapps.io/MarginalFactors/>.
- (46) Center for Air, Climate, and Energy Solutions, CACES Reduced Complexity Models Data Download. 2018, <http://airclimateenergy.org/data/>.
- (47) Pope, C. A.; Lefler, J. S.; Ezzati, M.; Higbee, J. D.; Marshall, J. D.; Kim, S. Y.; Bechle, M.; Gilliat, K. S.; Vernon, S. E.; Robinson, A. L.; Burnett, R. T. Mortality Risk and Fine Particulate Air Pollution in a Large, Representative Cohort of U.S. Adults. *Environ. Health Perspect.* **2019**, *127*, No. 077007.
- (48) Balding, M.; Whinery, T.; Leshner, E.; Womeldorff, E. *Estimated TNC Share of VMT in Six US Metropolitan Regions*, 2019.
- (49) Heo, J.; Adams, P. J.; Gao, H. O. Reduced-form modeling of public health impacts of inorganic PM_{2.5} and precursor emissions. *Atmos. Environ.* **2016**, *137*, 80–89.
- (50) Tessum, C. W.; Hill, J. D.; Marshall, J. D. InMAP: A model for air pollution interventions. *PLoS One* **2017**, *12* (4), e0176131.
- (51) Gilmore, E. A.; Heo, J.; Muller, N. Z.; Tessum, C. W.; Hill, J. D.; Marshall, J. D.; Adams, P. J. An inter-comparison of the social costs of air quality from reduced-complexity models. *Environ. Res. Lett.* **2019**, *14*, No. 074016.
- (52) Alonso-Mora, J.; Samaranayake, S.; Wallar, A.; Frazzoli, E.; Rus, D. On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment. *Proc. Natl. Acad. Sci. U.S.A.* **2017**, *114*, 462–467.
- (53) Towns, J.; Cockerill, T.; Dahan, M.; Foster, I.; Gauthier, K.; Grimshaw, A.; Hazelwood, V.; Lathrop, S.; Lifka, D.; Peterson, G. D.; Roskies, R.; Scott, J. R.; Wilkens-Diehr, N. XSEDE: Accelerating scientific discovery. *Comput. Sci. Eng.* **2014**, *16*, 62.
- (54) Nystrom, N. A.; Levine, M. J.; Roskies, R. Z.; Scott, J. In *Bridges: A Uniquely Flexible HPC Resource for New Communities and Data Analytics*, Bridges: A Uniquely Flexible HPC Resource for New Communities and Data Analytics, 2015.