

Research Statement

Andrey Makhanov

Ph.D. Candidate, Business Administration — Operations Management and Data Analytics

UNC Kenan-Flagler Business School

Andrey_Makhanov@kenan-flagler.unc.edu | andreymakhanov.com

Overview

Governments do not decarbonize transportation. Households buy cars, one at a time, through channels and over infrastructure that someone else designed. My research studies that intervening layer. I use archival data, econometric methods, and machine learning to ask how the structure of a distribution channel and the design of a charging network determine whether a clean technology actually reaches the consumers a policy was written to serve.

The organizing observation is that the operational layer between a policy and an adopter is not neutral. A state franchise statute written in the era of gasoline vehicles decides who is permitted to sell an electric vehicle, and therefore who can capture the demand a purchase subsidy creates. A public charging network built one site at a time decides which households find an electric vehicle usable at all, and the technology, siting, and geographic spread of that network turn out to matter more than its raw size. In both cases the policy instrument is only an input. The outcome is produced by an operational structure standing between the intention and the household, and treating that structure as the object of study—rather than as a friction to be assumed away—is what unites my dissertation and the agenda I intend to build from it.

Both essays are set in Washington State, which is deliberate rather than incidental. Washington combines an unusually sharp legal discontinuity, a granular vehicle registration record maintained by the state, high adoption rates, and a geography that runs from dense Seattle neighborhoods to sparsely populated eastern counties. That combination lets me hold the population fixed and vary the mechanism under study: first the legal structure of the sales channel, then the physical structure of the charging network. Methodologically, both papers share a commitment I hold strongly—that when a setting does not deliver clean identification, the right response is to triangulate across estimators with non-overlapping identifying assumptions and to state plainly what the data can and cannot support, rather than to select the specification that reads best.

That orientation reflects how I came into the field. Before and alongside my doctoral training I worked as a software engineer, machine learning practitioner, and security analyst, building systems that had to survive real users rather than idealized ones. It left me with a durable skepticism toward models whose conclusions depend on implementing parties behaving as assumed, and with the technical capacity to assemble and work with the large administrative and transactional data required to study those parties empirically.

Franchise Laws and the Distribution of Electric Vehicles

My job market paper, “Automobile Franchise Laws and EV Adoption: Evidence from Washington,” with Nur Sunar and Jayashankar M. Swaminathan, asks how the legal structure of automotive distribution shapes the diffusion of electric vehicles.

The setting is well suited to the question. State franchise laws generally require manufacturers to sell

through independent franchised dealers rather than directly to consumers. These statutes predate electric vehicles entirely; they were written to protect dealers from opportunism by the manufacturers whose brands they carried. But they now govern the channel through which most electric vehicles must reach buyers, and they do so asymmetrically. Washington permitted direct-to-consumer sales by all manufacturers until 2014, when RCW 46.96.185 restricted direct sales to manufacturers that already held a vehicle license before January 1, 2014—in the electric vehicle market, effectively Tesla alone. A February 2014 provision then allowed the exempted manufacturer to open an unlimited number of stores. A statute designed to solve a vertical contracting problem in one era became, without anyone intending it, an instrument of environmental policy in another, and it did so in a form that granted one incumbent a channel advantage its rivals were denied.

The theoretical prediction is not obvious, which is what makes the question interesting. We develop a quantity-setting model with two manufacturers—the exempted incumbent selling direct, and a rival that is either vertically integrated or forced through an independent dealer—and characterize equilibrium under both channel structures. Forcing the rival through a dealer introduces double marginalization, raising the rival’s effective retail price. We prove that this unambiguously increases the exempted manufacturer’s equilibrium quantity. The effect on the rival’s own sales, and on total adoption, is conditional: both hinge on whether cross-brand substitution is strong enough relative to the two firms’ margins, which yields a testable threshold rather than a sign prediction. Whether an asymmetric franchise restriction helps or harms aggregate adoption is therefore an empirical question, and the model tells us exactly which parameter decides it.

We answer it with weekly electric vehicle registrations from the Washington State Department of Licensing, covering battery electric, plug-in hybrid, and hybrid vehicles, together with clean-energy investment, lagged gasoline prices, average household earnings, and weather covariates. Rather than impose a parallel-trends assumption we cannot defend, we treat counterfactual construction as a prediction problem in the spirit of the machine-learning synthetic control literature: we train on the pre-2014 period using Virginia—which permitted direct sales throughout and matches Washington closely on population, median age, median rent, and age and marital structure—as a donor series alongside the covariates, then predict the post-2014 path. We compare linear regression, ridge, support vector regression, gradient boosting, and random forests on held-out pre-period validation data and carry forward the best performer.

The results are asymmetric in a way the model anticipates. Aggregate weekly electric vehicle adoption rose by roughly seven vehicles relative to the counterfactual (95% CI 1.6 to 12.4). That aggregate gain decomposes sharply: adoption of the exempted manufacturer’s vehicles rose by about 12.4 per week (95% CI 8.4 to 16.5), an increase on the order of 18 percent, while adoption of all other electric vehicles *fell* by roughly nine per week (95% CI -13.9 to -3.9), a decline of about nine percent. A Bayesian structural time-series analysis using **CausalImpact**, which models seasonality explicitly and rests on different assumptions than the synthetic control, recovers the same sign and a comparable magnitude for total adoption. Adding an indicator for the first Model S release and estimating in logs leaves the aggregate conclusion intact.

The contribution is to establish that sales-channel regulation is a first-order determinant of clean technology diffusion, and that its aggregate and distributional effects can point in opposite directions. A restriction that looks anticompetitive on its face—one firm exempted, its rivals constrained—accelerated total adoption in this setting, because the exempted firm’s direct channel converted demand more efficiently than the dealer channel converted the demand it was handed. That is an uncomfortable finding for both sides of the policy debate, and it is the kind of finding only an empirical treatment can produce. For policymakers weighing similar carve-outs, and

for manufacturers currently lobbying for direct-sales rights, the practical implication is that the aggregate case and the competitive case must be argued separately, because the evidence supports them differently.

Establishing that the channel matters raised an immediate follow-on question. If the legal structure of distribution determines who captures demand, what determines whether demand exists in a given neighborhood at all? That question moved me from the sales channel to the physical infrastructure that makes an electric vehicle usable, and to the household characteristics that decide how much public infrastructure is worth.

Charging Infrastructure, Income, and the Design of a Public Network

In “Charging Infrastructure, Income, and Electric Vehicle Adoption: Evidence from Washington State,” also with Nur Sunar and Jayashankar M. Swaminathan, I treat public charging deployment as what it is operationally: a network design problem.

The motivating decision belongs to a planner disbursing a fixed infrastructure budget across tens of thousands of candidate sites. Two choices dominate. The first is technology: a DC-fast port costs several times more than a Level 2 port but charges an order of magnitude faster, so the mix determines both how much coverage a budget buys and which driving patterns it serves. The second is geography: densify existing metropolitan clusters where installation is turnkey and demand is proven, or extend coverage into the gaps between clusters where installation is costlier and demand is thinner. Facility location models in operations management take exactly these technology- and geography-specific adoption elasticities as inputs. The empirical literature, which has established that charging availability matters in aggregate, has estimated technology mix, spatial structure, income, and housing type largely in isolation. The inputs the models need did not exist. This paper estimates them.

I assemble a quarterly panel of 454 Washington ZIP codes from 2010 through 2019—16,798 ZIP-quarter observations—merging state registration records split by battery electric and plug-in hybrid, Department of Energy charging station records geocoded to ZIP and classified by connector and facility type, American Community Survey demographics including median household income and share of multiunit housing, and regional retail gasoline prices. The window deliberately ends before 2020, so that the pandemic and the Bipartisan Infrastructure Law do not contaminate the estimates. Three features of the data—strong dynamic persistence in adoption, endogenous siting, and unobserved ZIP heterogeneity correlated with both regressors and outcomes—rule out OLS, static fixed effects, and static instrumental variables, so the primary estimator is Blundell–Bond system GMM with collapsed instruments and shallow lag depth, with ZIP and year-quarter fixed effects and Windmeijer-corrected standard errors clustered on ZIP.

Four findings organize the paper. First, an aggregate complementarity: the port elasticity is 0.067 ($p = 0.043$), so doubling public ports in a ZIP-quarter is associated with roughly a five to seven percent increase in registrations, conditional on lagged adoption and demographics.

Second, and more useful operationally, a technology reversal. Per port, DC-fast dominates (0.149 versus 0.055 for Level 2). Per dollar, the ranking inverts: because DC-fast costs five to ten times more, Level 2 delivers roughly 2.5 to 5 times more adoption per dollar wherever dwell time is not binding. The estimated elasticities imply a breakeven cost ratio of 2.71, well below observed cost ratios, which yields a clean site-level decision rule—DC-fast for corridors where dwell is measured in tens of minutes, Level 2 for destination charging where dwell is measured in hours.

Third, coverage beats density. HDBSCAN clustering shows Washington’s network is organized as compact, widely separated clusters—mean inter-centroid distance of 270 km against a mean cluster diameter of 43 km—and a direct panel test that decomposes new ports into within-cluster additions and inter-cluster gap-fills favors gap-filling. Notably, no single facility type dominates: total port availability, not siting mix, drives the aggregate effect.

Fourth, and I think most important, three demand-side moderators condition all of the above. Adoption is concave in income with a candidate plateau near \$156k, so the marginal returns to infrastructure are not where the early-adopter distribution would suggest. Multiunit housing is the sharpest and most precisely estimated moderation in the paper: the interaction with port availability is -0.450 ($p < 0.001$), implying that moving from an entirely single-family to an entirely multiunit ZIP roughly eliminates the infrastructure elasticity. In multiunit-heavy neighborhoods, public ports substitute for absent residential charging rather than complementing it, which points to curbside and multiunit-targeted programs rather than more public stations. And battery electric vehicles are roughly three times more infrastructure-elastic than plug-in hybrids (0.162 versus 0.056), and more income- and housing-elastic as well, which argues for technology-differentiated allocation.

I also probe the income mechanism directly rather than asserting one. Concavity is consistent with substitution—higher-income households buying electric instead of gasoline vehicles—or with saturation, buying more of both at a decreasing rate as garage capacity binds. Separate-equation estimates directionally favor substitution; imposing shared persistence across a stacked two-equation system recovers symmetric concavity and the substitution reading disappears. An interaction of income with the electric–gasoline relative price is precisely estimated in both equations and points to an affordability-to-preference transition, but discriminates the mechanism only weakly. I report that the data cannot separate the two, and note that the operational implication is invariant to which dominates: pair infrastructure with income-targeted subsidies. Under substitution, subsidies close the affordability wedge; under saturation, they move more households onto the increasing-returns portion of the concave response.

Future Research

Two extensions follow directly from work in progress. The franchise law paper points toward heterogeneity by buyer income: high-income buyers plausibly select expensive vehicles and respond weakly to rebates, so the distributional consequences of a channel restriction may differ sharply from its aggregate effect. Splitting the registration data by vehicle price tier and manufacturer segment tests this, and connects the channel result to the income concavity I estimate in the infrastructure paper. The infrastructure paper points toward multi-state replication: because franchise regimes and charging networks vary independently across states, the two mechanisms can be estimated jointly, and the coverage-versus-density result can be tested where geography is not Washington’s.

More broadly, I want to extend this agenda to the next intermediary layer. Increasingly, the party standing between a policy and a household is an algorithm—a routing system, a pricing engine, a recommendation surface, a rebate-eligibility screen. These systems are designed under objectives that need not align with the environmental or distributional goal the policy encodes, and like franchise statutes they are rarely revisited once written. My technical background positions me well to study them empirically, and my current work supplies the template: identify the structure standing between intent and outcome, characterize whose incentives it encodes, and measure what it does to the people on the far side of it.

Earlier Work

Before entering the doctoral program I contributed to research outside operations, most substantially as a co-author on a genetics study published in *PLOS ONE* (Makhanova et al., 2017), which dissected the genetic architecture of atherosclerosis using quantitative trait locus analyses across three round-robin intercrosses of apolipoprotein E-null mice. My contributions were computational—formal analysis, software, and visualization for genome-wide scans and interaction tests. I include it not as part of my current agenda but as early evidence of what I bring to collaborations: the ability to take ownership of the analytical and computational layer of an empirical project in an unfamiliar domain, which is the same capability my dissertation rests on.