Big Push and Network Externalities: Evidence from Québec's EV Charging Market*

Hugo Cordeau†

January 26, 2026

Please click here to obtain the latest version of this paper.

Abstract

Industries with strong complementarities and network externalities may exhibit multiple equilibria. I examine whether direct government provision of network infrastructure can remedy this inefficiency. Using car registration and charging transaction data from 2017–2023, I estimate a structural model of vehicle demand and charging behavior, where consumers choose an EV and derive utility from undertaking OD trips given the charging network. I decompose the Big Push into two components: (i) a coordination effect—shifting the economy from the low-adoption to the high-adoption equilibrium—and (ii) an externality effect—shifting from the high-adoption to the Big Push equilibrium, in which the government moves first. I find that the coordination effect increases EV adoption by 140%, while the externality effect raises adoption by an additional 50%. The externality effect stems from placing socially valuable charging stations that would otherwise be privately unprofitable. It expands the network from 300 to 1,000 stations and yields net crowd-in overall: public provision crowds in private investment in the core while crowding it out at the periphery. Overall, the policy is effective, generating $11 billion in consumer value and $2.7 billion in climate benefits from a $185 million investment. This paper provides empirical evidence that direct public provision of network infrastructure can resolve coordination failures and that government provision can complement, rather than crowd out, private investment.

Keywords: Network Externalities, Coordination Failure, Industrial Policy, Electric Vehicles, Charging Infrastructure

*I am deeply grateful to my supervisors, Yanyou Chen, Eduardo Souza-Rodrigues, Andrew Steck, and Stephan Heblich, for their invaluable guidance. I owe special thanks to Victor Aguirregabiria, Mitsuru Igami, Nathaniel Baum-Snow, Jeffrey Sun, Zi Yang Kang, and James Macek for their insights and support. I also thank the participants of the Harvard Salata Institute, the LSE Environment Week, the AERE Summer Conference, the UEA's North American Meeting, the CREEA Emerging Scholar Workshop, the IO Workshop, and the Microeconomics Brown Bag at the University of Toronto for their valuable feedback and inspiration. I am especially grateful to my discussants, Damian Capelle and Jevan Cherniwchan, for their valuable feedback on this paper. Lastly, I gratefully acknowledge the support of FRQSC, the Climate Positive Energy initiative, and the School of Cities for funding my research.

†Department of Economics, University of Toronto (hugo.cordeau@mail.utoronto.ca)

1 Introduction

Network externalities make industrial policy design particularly challenging, as inadequate policy design risks locking industries into inefficient equilibria. To address this issue, Rosenstein-Rodan (1943) introduced the notion of a "Big Push," highlighting the need for large-scale, state-led investment to overcome coordination failures. This paper studies the effectiveness of such government intervention in the electric vehicle (EV) market. While existing research primarily examines demand-side policies, such as EV subsidies, much less is known about the role and effectiveness of supply-side measures, particularly on how charging providers respond to policy incentives. A central concern with such state-led intervention is the potential crowding out of private investment. This raises important questions: How does public network provision affect EV adoption and private station entry, and should governments invest directly in charging infrastructure—and, if so, where?

To shed light on these questions, this paper is the first to study direct network provision in a network industry with strong market complementarities. There are two main challenges in studying this issue. First, the consumer decision to purchase an EV extends beyond vehicle characteristics to include the quality of the charging network, a complementary market. Second, a firm's investment in charging infrastructure generates spatial spillovers that benefit other market participants without capturing the extent of these benefits.[1]

To address these challenges, my model must capture both consumers' responses to the existing charging network and the strategic reactions of private firms to network expansion. The existing literature does not provide a framework that incorporates these features. My paper fills this gap by developing a spatial equilibrium model in which firms decide where to invest and consumers derive utility from performing trips. Given the network quality, consumers then decide whether to purchase an EV model or instead choose a gasoline vehicle.

I apply the model to the EV market, where the charging network generates a classic chicken-and-egg problem: consumers adopt EVs only if a reliable fast-charging network is in place, while private charging providers invest only if adoption is high enough to ensure profitability. There are two key market failures. First, coordination failures can trap the market in two privately viable NE: a low- or high-adoption one. Second, private providers

[1] In contrast to Garg (2025), whose contribution is primarily theoretical and methodological—developing algebraic geometry tools in the context of historical industrial zoning provision—my work studies an evolving industry with strong network externalities and market complementarities. I focus exclusively on direct network provision and its impact on crowding in and crowding out of private firms. Springel (2021) also conducts related work, focusing on the non-neutrality of subsidies between charging stations and EV demand. The main distinctions of my approach are i) a focus on network provision, ii) explicit consideration of the spatial nature of the charging network, and iii) an analysis of how direct network provision affects private-firm investment and EV adoption.

fail to internalize the externalities they generate—(i) network-formation externalities that enhance the value of existing stations, and (ii) cross-market externalities that stimulate EV demand—capturing only their private returns. My paper addresses both failures. I decompose the Big Push total effect into a Coordination Component and an Externality Component.[2]

Economists have long recognized that vertical integration can internalize externalities, including network externalities. In the case of the EV market, this takes the form of incompatible charging standards, predominantly driven by Tesla. Yet Li (2019) shows that, although compatibility reduces business stealing and lowers the private return to entry, it also makes investments strategic complements: rivals' stations raise one's own payoff. On net, firms invest more under compatibility. Despite this, little evidence exists on which policies most effectively promote the deployment of compatible networks. This paper highlights the distributional nature of compatibility, as it makes cheaper non-Tesla models more attractive.

I study the Québec EV market, where the government implemented a Big Push by mandating the state-owned utility to build a province-wide, publicly accessible fast-charging network (the Electric Circuit). A salient feature of this policy is that the placement of charging stations is not endogenous to EV adoption, but instead driven by public mandate to connect cities. To achieve this, the government allowed cost recovery through regulated electricity tariffs. Beginning in 2015, the network connected cities along major corridors and then expanded outward. Although implemented differently from the U.S. NEVI Formula Program—which allocates US$5 billion in subsidies for DC fast-charger deployment—the two programs share core objectives centered on spatial connectivity. NEVI has been criticized for its slow rollout. One factor that may explain this pace is the mandate that stations be located along designated highway corridors, with a maximum spacing of about 50 miles (80 km). My model speaks directly to the relevance of such spatial constraints by informing the value of this placement approach.

I utilize five data sources in my analysis. First, I use data from the Electric Circuit network, a transaction-level dataset covering all public charging sessions between 2017 and 2023. Each record contains information on the consumer (ZIP code, vehicle model), the charging station (`ConsumerID`, `StationID`, geographic coordinates, opening date), the energy delivered (kWh), the timestamp of the charging event, the state of the battery at the start and end of charging, and the price paid. By linking consumer ZIP codes to the spatial location of charging stations and along with the state of the battery at the start and end of charging, I construct origin–destination trips. Second, I complement the public network

[2]I present an intuitive toy model that explains why this multiplicity arises and how direct network provision can alleviate the market failure; I illustrate the mechanism visually in Section A.2.

data with a dataset on private charger entry. Specifically, I web-scraped the ChargeHub website to obtain, for each charging station, the number of chargers, the plug standards, and the opening date. Third, I use vehicle registration data from Québec's public insurance agency (SAAQ) covering 2011–2022. The dataset includes every licensed vehicle in the province, with characteristics such as brand, model, year, weight, and the owner's city of residence. Fourth, I complement the registration data with a dataset on vehicle characteristics. Specifically, I web-scraped specifications from *The Car Guide*, including MSRP (CAD), weight, horsepower, battery capacity (kWh), range, brand, model, and trim. I then merge this dataset with the registration data and interact it with EV subsidy policies to compute consumer prices, assuming 100% pass-through. Combining these two datasets yields EV adoption by model for each city and year. Finally, I use OpenStreetMap Routing Machine (OSRM) to estimate routing distances between each pair of cities and between each charging station and each city. All datasets span at least 2017–2022, with several extending beyond. The reduced form analysis uses 2015–2022 registration data, model estimation relies on the 2017–2023 period when transaction-level charging data are available, and finally the counterfactual simulations use 2022 vehicle technologies to reach the steady-state level of adoption.[3]

I first provide causal evidence on EV adoption in response to charging network deployment by exploiting the public network's concentric rollout. Using 2015–2022 city-level registration data, I implement a staggered difference-in-differences design. Cities north of Québec City are classified as connected once the charging network has no gap greater than 50 km remains between them and the network core. The 50 km threshold follows Springel and Li (2025), who study charging network spillover. This classification yields connected (treatment) and not-yet-connected (control) cities. I find that newly connected cities experience a statistically significant average treatment effect of a 120% increase in EV adoption, with dynamic treatment effects rising monotonically over time, consistent with the durability of vehicles. Robustness checks using a 100 km cutoff yield nearly identical results.

I then develop a structural model that captures the decision-making of (i) consumers, who choose whether to adopt an EV, and (ii) charging providers, who decide whether and where to invest under a free-entry condition. To capture how network quality shapes adoption, the model also incorporates (iii) travel and charging behavior, which summarizes the expected value of using the charging network for residents of a given origin with a given vehicle model; network quality varies by vehicle range. This object bridges the two markets:

[3]A more rigorous representation is that I map the city of origin o, the vehicle model j, the chosen charging station s, and the state of the battery at the start and end of the trip b into an origin-destination pair and a station choice: $f_j(o, s, b) = (o, d, s)$. Further details are provided in Section A.1.

consumers derive utility from vehicle purchase, but their adoption decisions depend on network quality. Crucially, an output of consumers' travel and charging behavior is a mapping from EV adoption, trip-taking, and charging choices into charging demand. Firms observe this demand and earn profits from realized charging activity.

In the road network, each city (about 100) and potential charging region (about 400) is represented as a node in a graph, with roads forming edges that connect nodes. Consumers make discrete choices over destinations given their origin (ZIP code), destination preferences, and the quality of the charging network. An outside option is also available: consumers may forgo the trip altogether if it is infeasible or perceived as too risky or anxiety-inducing.

Charging firms take the public first-mover build-out as given and simultaneously make entry decisions based on sunk entry costs and the discounted stream of expected profits, in the spirit of Springel (2021). The entry model follows Bresnahan and Reiss (1991). The number of charging stations in a market is the equilibrium outcome of a complete-information entry game in which the installed EV base determines EV adoption and charging demand. Charging demand implies an annual market potential for charging sessions in each region, with per-station demand equal to regional demand divided by the number of active stations. A firm provides a charging station at a node only when doing so yields positive profits. Each charging provider operates its own network. To avoid unstable equilibria and to focus on the research question, this paper restricts attention to two market structures: (i) a set of atomistic single-station owners and (ii) a legal monopoly over the entire charging network. In these cases, the firm either captures all or none of the (charging) network externalities.[4]

In sum, EV demand and charging-station supply are jointly shaped by network externalities: consumers' adoption rises with the coverage and quality of the charging network, while station entry and investment increase with the EV user base. Investments are strategic complements across nodes but substitutes within nodes. I treat Québec as a small open economy and take vehicle technology as exogenous, abstracting from the manufacturing sector. This is not a major limitation: Remmy (2025) shows that, when manufacturers endogenously set range and price as a function of charging network quality, network coverage substitutes for shorter vehicle range and EV markups decline as the network expands. Therefore, incorporating manufacturers would amplify the feedback between network expansion and consumer adoption. I abstract from this dimension to focus on the research question—how public

[4]The model could be extended in two ways. First, by allowing for network incompatibility, where competing providers deploy proprietary standards that fragment the charging network Li (2019). Second, by incorporating an ω-network, where ω represents the share of network externalities captured by a given firm. In this case, the firm internalizes ω of the network benefits when making investment decisions, while the remaining externalities spill over to atomistic competitors. The monopoly case is equivalent to $\omega = 1$, while the single-station monopolist corresponds to $\omega = 0$, capturing only its private value.

provision of network infrastructure influences consumer behaviour and (charging) firm investment.

The estimation proceeds in three stages: EV choice (demand), destination choice, and charging-station choice, and is implemented in reverse order. First, conditional on an origin-destination pair (o, d) and vehicle model j. I construct the feasible station set implied by vehicle range and project the state of charge along the path. Using the realized charging event, I estimate a maximum-likelihood model of charging-station choice. The identification comes from preferences over when to charge, with consumers targeting stations that leave them near 25% SOC at the start of charging. This step delivers an origin-destination charging quality that represents the ex-ante convenience of completing the trip given the network topology. Second, I estimate a destination-choice model for each origin city o, where utility depends on destination-specific preferences, travel distance, and charging quality. I use the city of origin and vehicle range to construct feasible destination sets, and I exploit observed trips to estimate the relevant parameters via maximum likelihood. These estimates imply a network-quality index that aggregates charging convenience across destinations for each origin, interpreted as the expected value of using the charging network. Finally, I use network-quality index as a sufficient statistic for charging access in the EV demand stage (estimated via BLP), alongside prices and product characteristics, to quantify how network provision affects EV adoption.

The estimation reveals four main facts. First, consumers find it costly to charge away from their preferred battery level. Second, improvements in the charging path increase the number of trips undertaken on a given origin-destination route. Third, the estimated destination preferences are internally consistent: preferred destinations deliver positive utility, while the disutility of longer trips is negative but statistically insignificant. This pattern is consistent with the idea that charging-path quality, rather than distance per se, governs long-trip behavior. Importantly, these estimates are conditional on observing a charging trip and therefore capture long-distance travel rather than commuting, for which distance would likely play a larger role. Fourth, consumers respond positively to improvements in the charging network. The results are robust across multiple specifications.

To provide economic intuition, I perform a set of back-of-the-envelope simulations. I numerically evaluate the marginal value of an additional 100 km of range for vehicle models in the 2016 and 2022 charging network, the latter being better developed. I also compute the (maximum) marginal value of adding a charging station under the 2016 and 2022 network configurations. Finally, I assess the interaction between vehicle range and network quality using the 2022 charging network.[5]

[5] Coming shortly.

On the supply side, firms make free-entry decisions in the spirit of Bresnahan and Reiss (1991). I assume linear entry costs of USD $250,000 per charger. I do not estimate this cost; instead, I use engineering-based cost estimates and the entry cost assumption which lies around Li (2019) estimate. I fix the per-session price at $15 and assume a $10 price-cost margin, consistent with prevailing market prices; the crown corporation's pricing is also effectively regulated around this level.[6] [7]

Using the estimated parameters and 2022 EV technology, I simulate steady-state equilibrium outcomes. As consumers and firms best respond to each other, two stable and privately viable Nash equilibria emerge under laissez-faire: a Low equilibrium, in which both sides adopt little, and a High equilibrium, in which consumers and firms adopt moderately. This outcome is not socially optimal, as many charging stations are socially valuable but privately unprofitable. I therefore construct a third equilibrium—used as the baseline— the Big Push, in the spirit of a Stackelberg game in which the government moves first. Specifically, the government installs the realized 2022 public network; given this public network, consumers and private charging firms then re-optimize, yielding a new Nash equilibrium. I decompose the total Big Push effect into two components: a Coordination Component, capturing the shift from the Low to the High equilibrium, and an Externality Component, capturing the additional gains from moving from the High equilibrium to the Big Push.

Coordination Component. A central role of the Big Push is to ensure that the market converges to the high-adoption equilibrium. Moving from the Low to the High equilibrium increases EV adoption by 139%, yielding an average per-vehicle willingness-to-pay (WTP) gain of about $10,300. This transition also generates substantial climate benefits, valued at roughly $1.4 billion under a social cost of carbon (SCC) of $185 per ton. This component captures the coordination effect—how resolving the chicken-and-egg problem between consumers and firms shifts the economy from the low-adoption trap to the high-adoption state.

Externality Component. Beyond coordination, the Big Push generates additional gains by connecting crucial nodes. Moving from the High equilibrium to the Big Push raises per-vehicle WTP by an additional $5,900 and yields a further 52% increase in EV adoption. The corresponding climate benefits amount to about $1.3 billion. These gains stem from strategic public investment in high-connectivity nodes that are socially valuable but privately unprofitable, generating crowd-in effects at the network core and improving overall efficiency.

These findings corroborate Rosenstein-Rodan (1943)'s insight that Big Push policies can

[6] https://lecircuitelectrique.com/en/costkWh/

[7] https://www.hydroquebec.com/business/customer-space/rates/rate-br-fast-charge-stations-billing.html

overcome coordination failures. Consistent with this mechanism, I show that Québec's realized Big Push both resolves the coordination failure and crowds in private firms. Moreover, the intervention induces additional EV adoption at an implied cost of $12-$26 USD per ton of carbon abated, indicating that direct public network provision can strategically complement private investment by internalizing key externalities.

Related Literature

This paper contributes to three strands of literature: (i) EV adoption; (ii) industrial policy in industries with network externality; and (iii) the option value of infrastructure. First, I contribute to the literature on EV markets with a focus on network externalities. Prior work studies how subsidies and complementary infrastructure affect EV adoption (Pavan, 2017; Fournel, 2023; Li, 2019; Springel, 2021; Remmy, 2025; Springel and Li, 2025; Fang et al., 2025).In related work, Li (2019) show that plug standardization increases consumer surplus and sales while reducing redundant infrastructure build-out by altering firms' investment incentives. Their question is whether charging networks should be compatible or incompatible; they find that compatibility raises adoption. My question instead is how a compatible network should be built and how policy can overcome the coordination failure that arises when standards are compatible. To answer this, I develop a rich spatial model in which consumers derive utility from completing trips, and network topology determines the feasibility and convenience of long-distance travel; this trip-driven demand links the placement of stations to EV adoption and firms' entry incentives. A complementary paper, Fournel (2023), examines the effectiveness of EV subsidies on EV adoption and models the charging provider as a local planner. I depart from Fournel (2023) by studying private firms' entry responses to improvements in network quality, asking whether public provision induces sufficient EV adoption to crowd in private investment. A closely related contribution is Springel (2021), which studies the non-neutrality between charging-station subsidies and consumer subsidies. While Springel focuses on subsidy non-neutrality, I focus on how charging firms and consumers respond to public network provision. My theoretical contribution is to develop a rich spatial model in which consumers derive utility from completing trips and therefore value the charging network through the expected utility it delivers. My empirical contribution is to show that direct public network provision crowds in private firms and generates additional EV adoption, at an implied cost of $12-$26 USD per ton of carbon abated, where the lower bound reflects the total "Big Push" effect and the upper bound isolates the externality channel.

Second, I provide contemporaneous empirical evidence on coordination failures in industrial-

policy settings by studying a measurable, real-world Big Push. In industries with network externalities and strong complementarities, multiple equilibria may arise, creating the potential for coordination failures. To address this, Rosenstein-Rodan (1943) argued for large-scale, state-led investment—a "Big Push"—to shift the economy toward the high-adoption equilibrium. A central concern, is that such intervention may crowd out private investment. My main contribution to this literature is to show that direct public network provision can both (i) resolve the coordination failure and (ii) crowd in private firms. Recent work seeks to isolate the coordination channel from fundamentals. Garg (2025) develop an algebraic-geometry framework to guide equilibrium selection and classify factual and counterfactual equilibria; applied to Indian industrial zones, they find that zones increase the probability of escaping a low-industrialization equilibrium by 38%. I complement their approach by studying equilibrium switching in an evolving industry using a novel empirical model and by quantifying the welfare gains from direct public provision of network infrastructure. While Garg (2025) focus on industrial zones as a legal and institutional intervention, I focus on the physical build-out of a compatible charging network and its general-equilibrium implications. Other approaches, including Choi and Shim (2024) and Sturm (2025), study how subsidies can address coordination failures. Choi and Shim (2024) show that a one-time technology subsidy in 1970s South Korea resolved coordination failures, shifting the economy to a high-industrialization equilibrium and increasing heavy manufacturing output by 8.6%. Sturm (2025) propose a super-Pigouvian policy that preserves decentralized price regulation while preventing coordination failure and, in a structural application to South Korea's heavy and chemical industry, modestly improves welfare relative to the worst Pigouvian equilibrium. My paper differs from existing work in two key ways. First, I study direct network provision rather than price-based mechanisms or equilibrium-selection criteria. Second, I exploit a realized Big Push and show that, although imperfect, the local authority's build-out effectively resolved the coordination failure. My theoretical contribution is to develop a multiple-equilibria model in which direct public network provision can alleviate a coordination failure and shift the market out of a low-adoption equilibrium. My empirical contribution is to show that the realized Big Push achieves precisely this: direct public provision both (i) resolves the coordination failure and (ii) crowds in private firms. Importantly, the spatial pattern of crowd-in underscores the complementarity between public provision and private entry: public investment strengthens the profitability of core connector nodes, where private firms expand, while partially displacing private entry at the periphery.

Third, I connect network economics to the option-value literature by treating the fast-charging network as a real option that provides insurance and flexibility even when it is not used. Weisbrod (1964) first formalized non-use value in collective goods, while Arrow

and Fisher (2013) emphasized quasi-option value—the value of preserving opportunities under uncertainty and irreversibility. In this setting, network availability along potential trips increases the value of EV ownership by reducing downside risk from contingencies (battery depletion, detours, etc.), even if drivers rarely charge en route; this non-use value is therefore only weakly monetized in observed charging transactions. Building on Dixit and Pindyck (1994), the option value is especially large in rural areas that "bridge" regions and complete paths, where small additions to connectivity generate large gains in route value. My theoretical contribution is to model charging infrastructure as an option value: access to fast charging along potential routes raises EV utility even when drivers do not realize charging, so the network's value can exceed what is implied by observed usage. My empirical contribution is twofold. First, my reduced-form evidence shows that continuous corridor coverage by fast-charging stations increases EV adoption. Second, in the structural model I quantify this mechanism by allowing route-level fast-charging access to raise utility even without realized charging, and I derive an origin-destination route-connectivity measure based on expected, rather than realized, trip patterns. I then show that it implies a substantial willingness to pay for network connectivity.

The paper proceeds as follows. Section 2 reviews the institutional background and policies affecting EV adoption, network deployment and data. In Section 3, I show that consumers respond to improvements in the charging network through increased EV adoption. Exploiting the network's concentric expansion in a staggered-treatment design, I estimate dynamic treatment effects and present empirical facts on network usage that motivate the model assumptions. Section 4 presents the model, the equilibrium definition, and the solution algorithm. Section 5 describes the estimation procedure and reports the results. Section 6 presents the counterfactual analysis and discussion. Section 7 concludes.

2 Institutional Background & Data

In 2012, the province of Québec, through state-owned Hydro-Québec, launched the Electrique Circuit, the first public charging network in Canada. Its mission was simple: enable province-wide EV travel "so any citizen can drive anywhere in Québec with ease." In 2018, Québec enacted Bill 184, which (i) authorizes the government to set public fast-charging tariffs by regulation and (ii) allows the *Régie de l'énergie*, when setting electricity distribution rates, to include the revenues required to operate the public fast-charging service and to weigh stated economic, social, and environmental objectives. This framework effectively permits cost recovery from the broader rate base during the expansion phase. In parallel,

Hydro-Québec has extended coverage to all regions of the province, including remote corridors. Policies of this type are often described as a "Big Push," intended to accelerate early EV adoption.[8]

The public fast-charging network also mitigates two key frictions: (i) early availability and (ii) plug compatibility. During the study period, most Electric Circuit DC fast-charging sites provided both CCS and CHAdeMO connectors, and Tesla owners could access these public chargers using a CCS adapter.[9] [10] [11]

Compatibility reduces technology risk for consumers by allowing them to choose vehicles based on car characteristics rather than plug standards. Québec's policy is unusual in this regard. The closest parallels are China, Germany, and Norway. However, China and Germany are large economies with substantial domestic vehicle manufacturing, which makes it harder to isolate the effect of charging-network provision on adoption. Norway, in turn, combined network expansion with exceptionally generous EV incentives, including large tax exemptions for EV purchasers and road-fee waivers, which confounds the role of infrastructure. Québec offers a cleaner setting: it has a dispersed, North American-style geography with harsh winters, which raises the barriers to EV adoption and increases the relevance of fast-charging coverage. As a result, Québec is an ideal laboratory to study how direct public network provision affects EV adoption and market structure.

2.1 Network Evolution

Table 1 shows that the fast-charging network expanded rapidly from 2016 to 2023. The number of charging sites increased from 132 to 1,419 (more than tenfold). Most of this expansion came from public deployment: public sites rose from 40 to 895 (about 22 times more), compared with Tesla sites increasing from 40 to 302 (about 8 times more) and private sites from 52 to 222 (about 4 times more). Overall, public provision accounted for about two-thirds of all new sites, versus about 20% from Tesla and 13% from private operators. As a result, the network became much more public over time: the public share increased from 30% in 2016 to 63% in 2023, while Tesla's share fell from 30% to 21% (after peaking in 2018) and the private share declined from 39% to 16%.

[8] Bill 184 (2018, c.25): An Act to promote the establishment of a public fast-charging service for electric vehicles; amendments to the Hydro-Québec Act and the Act respecting the *Régie de l'énergie*.

[9] Electric Circuit connector guidance: stations may offer both CCS and CHAdeMO, although typically only one DC connector can be used at a time.

[10] CHAdeMO is primarily used by Nissan and Mitsubishi, while CCS (Combined Charging System) is used by manufacturers such as Chevrolet, Hyundai, Kia, Volkswagen, BMW, Ford, and Audi.

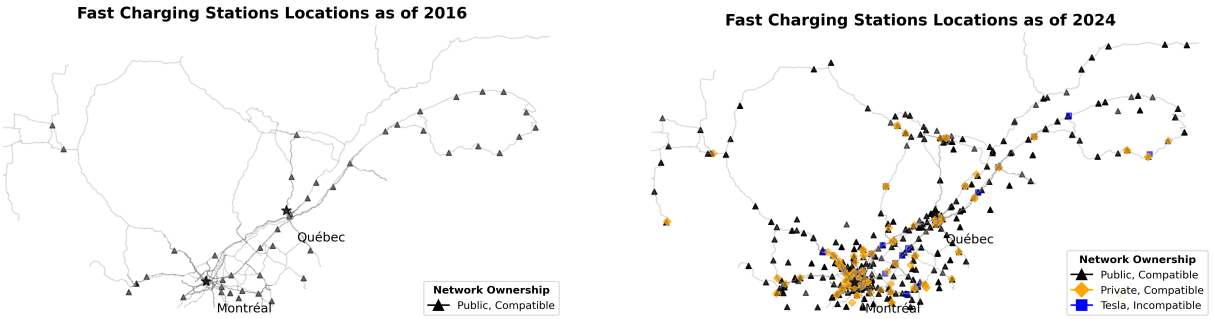
[11] Tesla vehicles can use CCS fast chargers via a CCS1 adapter; beginning in 2023-2025, Tesla has also started opening select Superchargers to non-Tesla vehicles through NACS/Magic Dock, which post-dates the period studied here.

Table 1: Stations by Network Type, 2016—2023

Year	All	Public	Private	Tesla
2016	132	40	52	40
2017	232	68	72	92
2018	400	111	93	196
2019	596	238	110	248
2020	802	415	123	264
2021	1013	585	156	272
2022	1217	741	180	296
2023	1419	895	222	302

Notes: All = Public + Private + Tesla.

Figure 1 shows the rapid expansion—and extensive coverage—of the public fast-charging network as early as 2016, and how this network continued to grow, becoming much denser by 2024. Interestingly, there has been a surge in private fast-charging providers. However, these networks differ from the public ones, as they are denser near urban centers—likely alleviating congestion. From this figure alone, it remains unclear if this is a crowd-in or crowd out story.



(a) Public Fast-Charging Network: 2016 (b) EV Network Location 2024

Figure 1: Public Fast-Charging Network (2016) and EV Network Location (2024)

2.2 Data

This paper merges panel data on vehicle registrations, charging infrastructure (through web scraping), and transactional charging usage. The core unit of geography in the reduced-form analysis is the city, which provides the most disaggregated level of observation. For the structural model, I aggregate to the level of the *Municipalités Régionales de Comté* (MRC), which are clusters of municipalities. All geographic coordinates are recorded in EPSG:4326.

Vehicle registrations provide car fleets from 2011–2022 across more than 1,000 Quebec municipalities. The dataset reports, for each city and year, every vehicle with a license plate, including its model, year, curb weight, and (from 2017 onward) fuel type. [12]

Two steps are required to make the data suitable for analysis. First, identifying EV adoption prior to 2017 is essential for measuring the early effects of network deployment, yet fuel type is not recorded before that year. I therefore train a random forest classifier using model, brand, and weight to impute whether a vehicle is an EV. Because the models of 2016 are nested in the car fleet of 2017, the algorithm achieves an accuracy rate of 97%, enabling consistent EV share measures over the full sample period.

Second, to estimate discrete choice models, car-level characteristics are needed. I web-scraped The Car Guide to obtain brand–model–trim attributes, including MSRP, range, and battery size. To improve matching accuracy, I employed the ChatGPT API to reconcile scraped trims with the registration data, conditional on model, brand, and weight.[13] [14]

Charging infrastructure is observed annually at the station level, with precise latitude and longitude. Because no comprehensive dataset is fully up to date, I use ChargeHub to reconstruct the rollout of the network, treating the date of the first user comment as the station's opening date. These data are central to identifying supply-side parameters. In particular, the period after 2020 sees substantial private-sector entries, and it remains an open question whether public investment crowds in or crowds out private investment. I use the inferred opening dates of private stations to estimate parameters governing private-sector entry decisions.[15] [16]

Transactional charging data are observed at the charger level from 2017–2024. The inventory covers roughly 7,000 stations (about 1,300 fast chargers and the remainder slow) with rich station metadata (in-service date, nameplate power in kW, operator, and precise latitude-longitude). Each charging session records a `ConsumerID` and `StationID`, energy delivered (`kWh`)), price paid, vehicle model, the consumer ZIP code, and battery state-of-charge at the onset and offset of the charging event. This unique combination of high-frequency transactions and detailed infrastructure attributes directly supports the model estimation. I define a "station" as a colocated set of posts operated by a single provider at one address; co-located providers are treated as distinct stations.

Charging behavior comes from session-level records on the public network from 2017–2024.

[12] https://www.donneesquebec.ca/recherche/dataset/vehicules-en-circulation.

[13] https://www.guideautoweb.com/en/.

[14] This approach yields more accurate matches than standard fuzzy string matching.

[15] https://chargehub.com/en/

[16] I validate this approach against public-sector deployments, for which reliable commissioning dates are available

The dataset covers over 100 million sessions in total, including about 3 million fast-charging sessions. Each record includes an anonymized consumer identifier, vehicle model, the home postal code (3-digit FSA), the station identifier, timestamp, energy delivered (kWh), duration, billed amount, and battery state-of-charge at the start and end of the session. I use station coordinates to locate the charge and the consumer FSA to proxy home location; I then compute trip distances and characterize within- and out-of-home-area charging. Session data are linked to the station inventory through station IDs and to the registration panel through the harmonized make–model dictionary.

Policy variables encode Quebec's provincial EV rebates and Canada's federal iZEV program at the vehicle–year level. I implement the statutory rules directly as piecewise functions of nominal MSRP, vehicle type, and battery characteristics. In Quebec, the provincial rebate program began in 2017 with generous support for BEVs and PHEVs, subject to MSRP caps that tightened over time. The federal iZEV program was introduced in 2019 and remains in effect, with eligibility and rebate amounts varying by vehicle class and range. Table 2 summarizes the structure of these programs.

Table 2: EV Purchase Incentives

Program	**Years**	**Ineligible above MSRP**	**BEV**	**PHEV**
Quebec	2017—2019	BEV: > \$125k; PHEV: ≥\$75k	up to \$8k	up to \$8k
Quebec	2020—2021	≥\$60k	\$8k	up to \$8k
Quebec	2022—present	≥\$60k	\$7k	up to \$5k
Federal iZEV	2019—2021	≥\$55k	\$5k	up to \$5k
Federal iZEV	2022—present	Cars: ≥\$65k; SUVs/pickups/vans: ≥\$70k	\$5k	up to \$5k

Note: Amounts are nominal and shown as simple maxima. Quebec 2020—present PHEV amounts vary by kWh. For more details, see Table A.3.

Caps apply to nominal MSRP as defined in the program rules. Rebate amounts are applied in the analysis as price offsets at the model–year level.

3 Descriptive Evidence

The objective of this section is twofold: (i) to assess whether EV adoption increases when the charging network improves, and (ii) to document how consumers use the charging network. In Section 3.1, I exploit the network's concentric expansion and implement a staggered difference-in-differences. This analysis provides evidence that consumers respond to improvements in charging-network quality through increased EV adoption. In Section 3.2, I present

the empirical facts, together with the relevant figures and graphical intuition on how consumers use, and derive value from, the charging network. This evidence provides intuition for the model and its key assumptions, and clarifies the rationale for adopting them. Finally, in Section 3.3, I state the assumptions explicitly and explain why they are reasonable in light of the empirical facts.

3.1 Reduced-Form Evidence

This paper posits that consumers care about the charging network. This section provides evidence for this mechanism. Identification follows Springel and Li (2025), who document a sharp increase in charging station usage once the maximum gap between charging stations along a route falls below 50 km, at which point routes are considered connected. While Springel and Li (2025) document an increase in charging station usage along connected routes—the intensive margin of EV use—it has not been shown that such connectivity leads to increased EV adoption—the extensive margin. This section provides such evidence. Using the same threshold as in their paper, I classify cities as connected (treatment) or not-yet-connected (control). The implementation exploits the concentric south-to-north expansion of the public fast-charging network, focusing on Route 138 north of Quebec City[17]. In the spirit of Springel and Li (2025), a city is considered connected if the maximum distance to the nearest charging station along the relevant corridor is at most 50 km. It is otherwise classified as disconnected (i.e., not yet connected). The identifying assumption is that once the network becomes connected, residents gain access to the corridor and are more likely to adopt EVs.

This setting provides an ideal case study given the gradual northward rollout of the charging network between 2016 and 2022. Because the bulk of the population resides south of Quebec City, the expansion generates a natural staggered rollout well suited to a difference-in-differences framework. I exploit this northward expansion to construct treated and not-yet-treated groups with similar characteristics. Crucially, network deployment is policy driven: earlier-treated cities are not selected based on inherent adoption propensity but rather on their proximity to the network core. While proximity is not random, control cities lie only slightly farther along the same corridors. As a result, differential adoption patterns can be plausibly attributed to network connectivity rather than underlying city characteristics. Thus, while both the charging network and EV adoption expand over time, this connectivity measure—which classifies cities as connected or disconnected—isolates the causal effect attributable to changes in network value.

[17] See Figure A.2 in the appendix for the evolution of the network.

I estimate dynamic treatment effects using the staggered DiD estimator of Callaway and Sant'Anna (2021). Let $k = t - T_c$ denote event time relative to city c's first treatment date T_c. I report coefficients from the event–study regression

$$\log(MS_{c,t}) \;=\; \beta_0 \;+\; \sum_{k \neq -1} \delta_k \, \mathbf{1}\{k\}_{c,t} \;+\; \alpha_c \;+\; \gamma_t \;+\; \varepsilon_{c,t},$$

where $MS_{c,t}$ is the EV market share in city c and year t; $\mathbf{1}\{k\}_{c,t}$ equals 1 if city c is k periods from treatment in year t. For $k < 0$ (leads), δ_k tests for pre-trends; for $k > 0$ (lags), δ_k traces post-treatment dynamics. Following Callaway and Sant'Anna (2021), identification uses not-yet-treated cities as controls at each event time. City fixed effects α_c absorb time-invariant heterogeneity; year fixed effects γ_t absorb common shocks. I interpret δ_k as ATT(k) relative to $k = -1$ and cluster standard errors at the city level.

Using the Callaway and Sant'Anna (2021) estimator with city and year fixed effects, I estimate that connecting a city with gap no wider than 50 km raises log(EV share) by 0.788. In levels, this corresponds to $\exp(0.788) = 2.20$, or about a 120% increase relative to not-yet-treated controls. As Springel and Li (2025) also contain a 100 km estimator, I perform a sanity check in column (2) and find a near exact result, underlying the robustness of the result. Note that, as every 100 km is also a 50 km gap and that there is solely 40 potential cities to be treated, having a similar result is expected.

Table 3: Average Treatment Effect, Network Connection

	(1) 50 km	(2) 100 km
log(EV Share)	0.788*	0.788*
	(0.3367)	(0.3285)
City Fixed Effects	Yes	Yes
Year Fixed Effects	Yes	Yes
N	490	490
Note:	Stars: * 95% CI excludes 0	

Because pre-connection EV shares were small and vehicles are durable, adoption need not translate into large short-run changes. I therefore estimate dynamic treatment effects to examine the adjustment path. Figure 2 reports these results. Estimation is doubly robust and uses not-yet-treated cities as controls at each event time.

The figure highlights two patterns. First, there is no evidence of pre-trends: EV adoption is flat prior to treatment. Second, adoption rises sharply after treatment, with the estimated

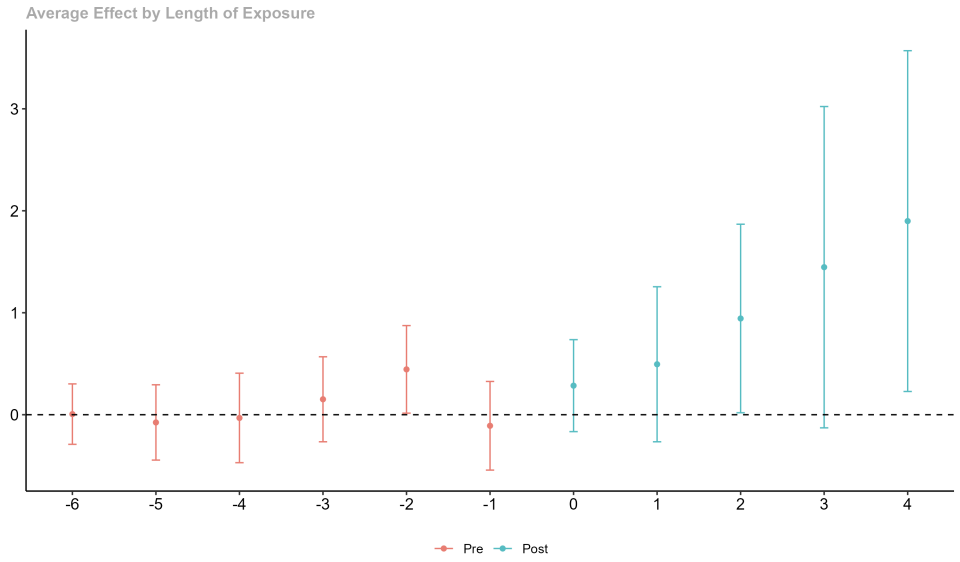


Figure 2: Dynamic Treatment Effects of Charging Network Connection on EV Adoption

effect increasing monotonically over time. The post-treatment estimates are noisy, however, with only the second and fourth periods significantly different from zero.

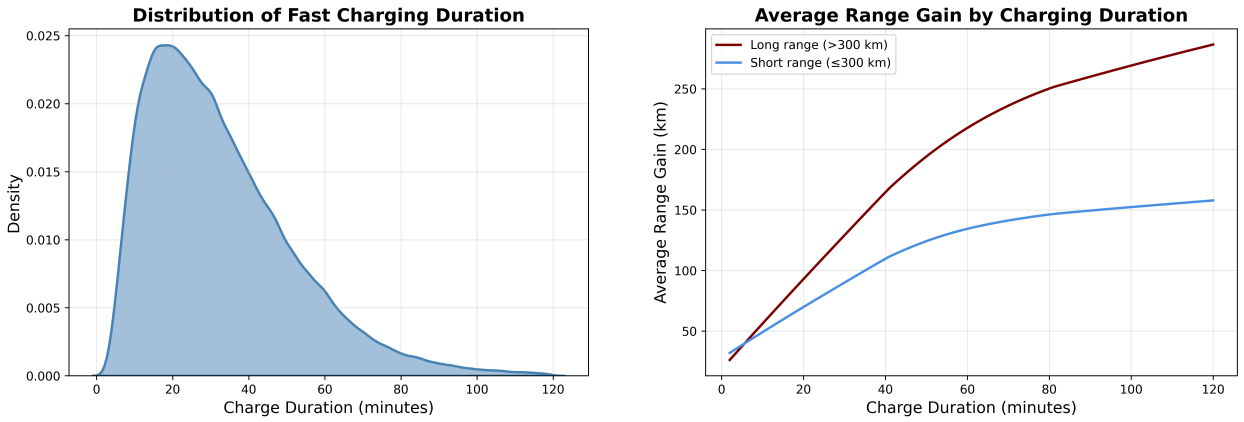
Evaluating the impact of a charging network on an outcome as complex as EV adoption is inherently non-trivial. The reduced-form estimates in this section should be read primarily as evidence that the relationship between network availability and consumer adoption is positive and economically meaningful, rather than as a basis for precise point estimates. The structure of the network makes the definition of treatment and control groups hazardous, with only a handful of suitable cities for the natural experiment. Moreover, the estimates capture (i) the extensive margin of being connected to a rich charging network on (ii) a subset of the population that might not be representative. Therefore, the coefficient external validity is weak. In particular, the reduced form cannot inform policymakers about the marginal value of an additional charger beyond those that connect peripheral cities to the network. These limitations underscore the need for a more elaborate framework. A credible model must jointly address the endogenous determination of EV adoption and charging station deployment. This estimate is not used in the structural model; rather, it underscores the need for, and relevance of, the structural analysis.

The next part of this section and the following develops such a framework. I first present key empirical facts in Section 3.2 along with its implied modelling assumption in Section 3.3 and then present a structural demand model of vehicle choice, the associated estimating equation, and the supply-side model of network entry in Section 4. Together, these components provide the basis for the paper's main results.

3.2 Empirical Facts

Empirical Fact 1: Charging is costly. This may be obvious, but it is useful to quantify what "costly" means in this setting. Here, "costly" primarily refers to time costs. Using the 2023 transaction-level charging data, Figure 3 reports (left) the distribution of charging-session duration and (right) the implied range gained as a function of time spent charging. The left panel shows that most sessions last between 15 and 60 minutes, with substantial mass around 20-30 minutes. The right panel plots charging duration on the x-axis against cumulative kilometres added on the y-axis, illustrating the time required to obtain meaningful additional range.

While these descriptive patterns do not enter the model directly, they motivate two modeling choices. First, because charging is time-consuming, I restrict attention to trips that involve at most a single charging stop. Second, the presence of sizeable time costs supports the clustering of sessions into "on-the-way" and "on-the-way-back" charging: I assume drivers charge only when needed to complete the trip (given the return-home constraint).



(a) Charging-session duration

(b) Range gained vs. charging time

Figure 3: Charging is time-consuming: duration and implied range gains (2023)

Empirical Fact 2: Over 90% of performed trips required only a single fast charge. Table 4 shows that, among trips that involve fast charging, the vast majority are completed with a single charging stop: 91.3% for short trips (<300 km) and 91.5% for long trips (≥300 km). Trips requiring two stops are relatively uncommon (about 7% in both range classes), and trips with more than two stops are rare (1-2%). The similarity across range classes suggests that multi-stop charging is not disproportionately concentrated among longer trips in this sample. Given that charging is costly, and charging twice is even more so, this pattern is unsurprising. What is more striking is the near invariance across range classes, which is consistent with either (i) consumers selecting into vehicle range classes that match their trip

needs, or (ii) consumers adjusting their realized trip patterns after purchase, undertaking only the trips that are worthwhile given the disutility of charging.

Table 4: Share of Unique Trips by Number of Stops and Range Class

Stops (per trip)	**Short (<300 km)**	**Long (≥300 km)**
1 stop (%)	91.3	91.5
2 stops (%)	7.0	7.2
>2 stops (%)	1.7	1.3

Empirical Fact 3: Consumers prefer to initiate fast charging around 25% battery remaining. A key empirical fact is that drivers tend to begin fast-charging sessions with roughly 25% battery remaining and rarely end above 80%. The latter pattern is largely mechanical: charging power declines sharply at high states of charge, so pushing beyond 80% is time-inefficient. The onset pattern is more nuanced. It plausibly reflects both (i) behavioral costs—starting a session at very low battery is psychologically costly due to range anxiety – and (ii) technological considerations—arriving with a lower battery can reduce total charging time because charging is faster at low states of charge.

Importantly, this onset behavior looks similar across vehicles with different driving ranges: I find no clear evidence that longer-range models (e.g., those with more than 300 km of rated range) initiate charging at systematically different battery levels. This suggests that the 25% target is primarily driven by behavioral frictions rather than range per se. The 300 km cutoff is used only as a descriptive threshold for visualization; the model itself uses continuous vehicle range.

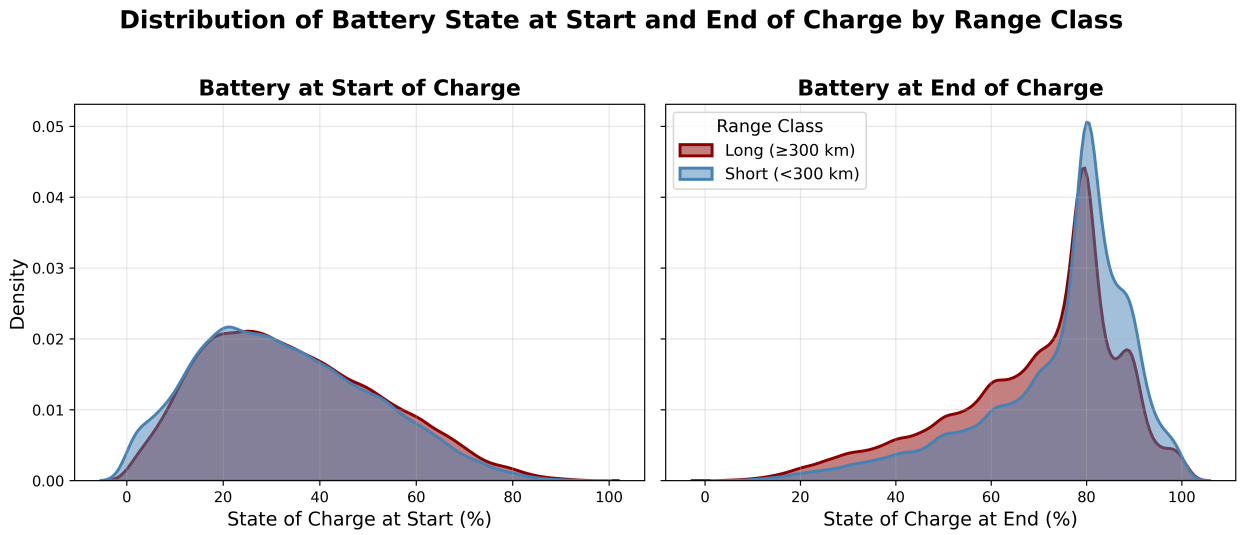


Figure 4: Distribution of state of charge at the start (left) and end (right) of fast-charging sessions (2023)

Empirical Fact 4: Consumers prefer to return home with *at least* 20% battery remaining. A subset of fast-charging events can be linked to subsequent local (slow) charging using `ConsumerID`. In these cases, a fast charge is followed within the same 24 hours by a short-range local charge away from their city of origin. This sequencing provides direct evidence on the battery level at which drivers arrive near their final destination, a quantity that is otherwise unobserved in the fast-charging data alone. Figure 5 plots the inferred SOC upon arrival for this subsample. The distribution indicates that drivers typically arrive with at least about 20% battery remaining.

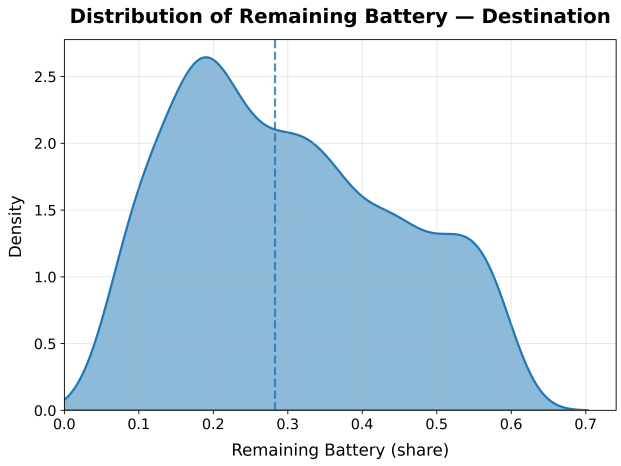


Figure 5: Inferred state of charge upon arrival near the final destination for trips that are followed by a local on-street charge

Empirical Fact 5: Over 80% of single-stop fast-charging events occur on the way back home. I observe the battery state of charge at the beginning and end of each fast-charging session. Restricting attention to single-stop trips, the charging event can occur either on the outbound leg (on the way to the destination) or on the return leg (on the way back home). I leverage the `ConsumerID` panel structure and drop consumers who fast-charge more than once within a 24-hour window, so that each retained day corresponds to a unique single-stop event. I then classify each event as outbound versus return using the SOC at the start of the charge and simple feasibility logic (see Appendix Section A.1 for details). The key intuition is that, because fast charging is time-costly, drivers typically charge only when needed to complete the trip. In particular, conditional on the SOC at the end of the charge, I compute the projected arrival SOC at home—the SOC the driver would have upon reaching home if they were returning immediately after charging. When the projected arrival SOC at home is implausibly high, the event is more consistent with an outbound leg (the driver has not yet reached the final destination). Conversely, when the projected arrival SOC at home is moderate, the event is consistent with return charging. Using this classification, I find that over 80% of single-stop fast-charging events occur on the way back home.

The left-hand panel shows the distribution of projected arrival SOC at home (computed conditional on the end-of-charge SOC). The distribution is bimodal: one mode corresponds to charges that would imply an implausibly high SOC upon arriving home (outbound events), whereas the other corresponds to charges that imply a reasonable arrival SOC at home (return events). The separation is consistent with Empirical Fact 4 and the idea that drivers target charging onset around 25% SOC. The right-hand panel reports the corresponding distribution conditional on classifying the event as outbound (see Appendix Section A.1 for further validation).

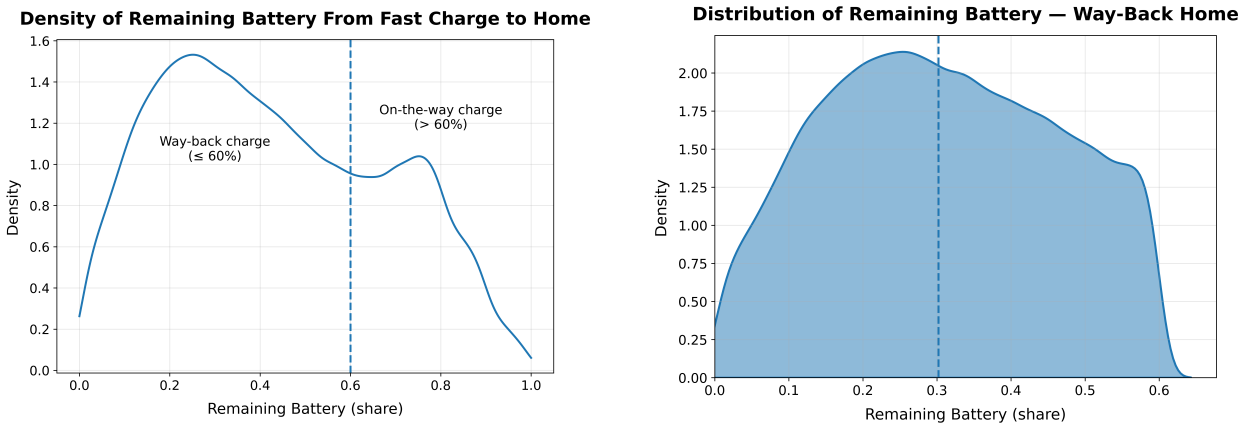


Figure 6: Distribution of projected arrival SOC at home, computed conditional on the end-of-charge SOC

3.3 Implied Assumptions

The empirical patterns in Section 3.2 imply a set of assumptions that I use to (i) define feasible choice sets and (ii) interpret observed fast-charging decisions as preference-revealing rather than constraint-driven.

First, because fast charging is time-consuming, I restrict attention to trips that involve at most a single fast-charging stop. This captures the dominant margin in the data and avoids modelling a high-dimensional multi-stop routing problem. The identifying logic is that a well-functioning one-stop charging network provides a reliable proxy for the more complex n-stop network, in the sense that it largely preserves the ranking of network quality across origin-destination pairs. Additional intermediate charging opportunities can be viewed as a positive transformation of the feasible set: they weakly expand reachability and reduce effective travel costs, but they do not overturn the ranking of destinations or the structure of optimal routing.

Second, I focus on *way-home* (return-leg) fast-charging events. I assume the consumer leaves home with a 100% battery, reaches a destination, and then fast-charges on the return leg before coming back home; the sequence is origin-destination-station, charge, station-origin. This restriction matches the empirical prevalence of return charging and enables destination inference using feasibility logic anchored at the home origin (see Appendix Section A.1). Intuitively, a driver would not fast-charge if returning home without charging were feasible, and would be unlikely to charge to a high state of charge in that case (consistent with the bimodal pattern in Figure 6). Conversely, when the battery level at charging onset is too low to reach home without charging, and the battery level at unplugging implies arrival at home within the preferred range, the event is more likely to be a *way-home* trip.

Third, I assume that drivers have a preferred battery level at which to initiate fast charging, and that deviations from this target are costly. Consistent with the observed clustering of start-of-charge SOC, I treat 25% as the preferred onset level and model charging initiation away from 25% as a utility cost.

Finally, I assume that arriving (back home) with low remaining battery is costly and that drivers typically maintain an arrival buffer of about 20% SOC. I use this regularity to bound the feasible destination set $\mathcal{D}_{odjt}$: I restrict attention to destinations that can be reached and return home. Operationally, I cap feasible round-trip distance at approximately 1.6 times the vehicle's range, which corresponds to starting at 100%, charging once from 10% to 80%, and returning home with about 10% SOC.

4 Model

The reduced-form findings of Section 3.1 indicate that EV consumers respond to the deployment of the charging network. To further understand the impact of the charging network under various policy scenarios and to control for its endogenous nature, I introduce a spatial equilibrium model. The model has two salient features: it accounts for how consumers use the network and generates a charging-network quality measure that explains EV adoption, while simultaneously modeling the charging-station investment decision. Fast-charging firms make free-entry decisions. They do not compete on price or quantity; instead, realized charging demand is split evenly across stations at a given node, taking rivals' entry decisions as given. Following Bresnahan and Reiss (1991), firms enter until expected profits are driven to zero, and no firm enters when expected profits are negative.

In this section, I first present the EV demand functions. I then describe how I construct the network-quality variable and detail its modeling. I also introduce the profit and cost functions. Finally, I discuss the equilibrium concepts.

4.1 EV Demand

Consumers play a two-stage game. First, they choose their car models; second, they perform trips to desired destination and choose where to stop along the way. We condense the second stage into a network quality variable, Λ_{ojt}, that varies by city and time, and model type through range (and incompatible plugs when it applies).

Consumer demand follows the discrete-choice framework of Berry et al. (1993). In each period t and city o, consumers arrive to purchase a vehicle. Consumer i's conditional indirect utility from choosing a car model j is given by:

$$U_{ijot} = \delta_{ojt} + \mu_{ijot} + \epsilon_{ijot} \tag{1}$$

where δ_{ojt} is the mean utility of product j in city o at time t, common to all consumers, and μ_{ijot} captures consumer-specific deviations from mean utility due to heterogeneous preferences. The idiosyncratic term ϵ_{ijot} is assumed to follow a Type I Extreme Value distribution. All alternative options—defined as models belonging to the Battery Electric Vehicle (BEV)—are treated as inside goods. The outside option is defined as purchasing a conventional gasoline vehicle, with its utility normalized to zero.

The mean utility is specified as:

$$\delta_{ojt} = \alpha \cdot p_{jt} + \beta \cdot \Lambda_{ojt} + \xi_{ojt} \tag{2}$$

where p_{jt} denotes the price of vehicle j, Λ_{ojt} represents the quality of the charging network for consumer in city o for model j. Unobserved vehicle-specific attributes and demand shocks are represented by ξ_{ojt}.

The individual deviation μ_{ijot} reflects consumer i's heterogeneous tastes and is modeled as:

$$\mu_{ijot} = \sigma_p \cdot p_{jt} \cdot \nu_{ip} + \sigma_\Lambda \cdot \Lambda_{ojt} \cdot \nu_{i\Lambda} \tag{3}$$

where ν_{ip} and $\nu_{i\Lambda}$ are independent standard normal deviates. Random coefficients allow for heterogeneity in price sensitivity, charging network valuation, and preferences for range.

Consumers choose the vehicle that gives them the highest utility. Market shares are determined by integrating individual choice probabilities over the distribution of consumer heterogeneity:

$$s_{ojt} = \int \frac{\exp(\delta_{ojt} + \mu_{ijot})}{1 + \sum_{k=1}^{J} \exp(\delta_{kot} + \mu_{ikot})} \, dF(\nu_i) \tag{4}$$

The vector of parameters to be estimated is $\theta = (\alpha, \beta, \sigma_p, \sigma_\Lambda)$. We recover $\xi_{ojt}(\theta)$ as the residual term implied by observed market shares, following the inversion approach in Berry et al. (1993).

4.2 Network Quality and Trip Demand

This subsection maps the charging network $\mathcal{K}_{ojt}$ into a micro-founded model-based measure of network quality, $\Lambda_{ojt} \equiv \mathbb{E}[\mathcal{K}_{ojt}]$, the ex-ante expected value of access to the charging network for a consumer of vehicle model j in origin o at time t. Constructing Λ_{ojt} requires (i) a model of origin-destination trip demand and associated preference heterogeneity, and (ii) a model of the value that the charging network delivers for each origin-destination route. The key point is that "network quality" is not captured by simple charger availability: the same station can be pivotal for some routes and irrelevant for others. Network quality is therefore inherently route-specific and intertwined with origin-destination demand and charging preferences patterns.

The section proceeds in three layers. First, I microfound route-level charging utility by having a discrete-choice model of charging-station selection along a given route (Section 4.2.1). Even when multiple stations are feasible, they deliver heterogeneous utility, so some stations are systematically more valuable than others for a given origin-destination pair, up to an EVT1 preference shock. Second, I summarize the implied charging value for each origin-destination pair and incorporate it into a destination-choice framework (Section 4.2.2), in which consumers choose among destinations based on origin-destination characteristics and the associated charging path. Third, I aggregate these objects into the ex-ante expected

value of the network (Section 4.2.3).

Under this structure, Λ_{ojt} provides a microfounded mapping from the charging network $\mathcal{K}_{ojt}$ to an economically meaningful index of network quality. This mapping is essential because stations differ in their contribution to network value. For example, if origins and destinations are distributed along a continuum of locations and consumers prefer to initiate charging around a particular state of charge, then a spatially uniform network can dominate a network that is highly concentrated in urban centers: uniform coverage reduces detours and feasibility gaps across many routes, whereas urban concentration leaves large portions of the route space poorly served.

4.2.1 Charging-Station Choice

This subsection specifies the consumer's charging decision within an (o, d) trip. Conditional on a trip from origin o to destination d and vehicle model j, consumer i chooses a charging station s from the feasible set $\mathcal{S}_{odjt}$ to maximize indirect utility. The feasible set is determined by technological constraints (battery capacity and range) and by the spatial configuration of the charging network (See Section 5.1.1 for details on $\mathcal{S}_{odjt}$).

Utility from choosing station s on the (o, d) route is

$$V^c_{iodsj} = \eta\, b_{odsj} \;+\; \varepsilon^c_{iodsj}, \qquad s \in \mathcal{S}_{odjt} \tag{5}$$

The term η captures range anxiety as a preference over the state of charge at which charging is initiated. Finally, ε^c_{iodsj} is an EVT1 taste shock, implying a logit choice rule conditional on feasibility.

The specification b_s for range anxiety is a quadratic loss around a preferred charging-onset state of charge b_s:

$$b_{odsj} = \left(\mathrm{SOC}^{\mathrm{start}}_{odsj} - \overline{\mathrm{SOC}}\right)^2 \tag{6}$$

I set $\overline{\mathrm{SOC}} = 0.25$ to match a salient empirical regularity: consumers disproportionately initiate charging when the battery is around 25% of capacity. In addition, the descriptive evidence indicates that charging sessions rarely continue beyond 80% of capacity. Accordingly, I impose an upper bound of 80% on the end-of-charge state of charge, reflecting the sharp decline in charging speed at high states of charge and the fact that few drivers choose to charge past this point. This implies that, for each vehicle model and range type, a given network delivers heterogeneous value. This matters because Λ is a key determinant of EV model choice.

In this model, conditional on an origin-destination pair (o, d) and a single charging stop, the primary dimension along which feasible stations differ is the state of charge at the start

of the session, $\text{SOC}^{\text{start}}_{iodsj}$. Through (6), $\text{SOC}^{\text{start}}_{iodsj}$ governs both (i) charging efficiency, since charging speed varies with the state of charge, and (ii) the usable post-charge range to complete the trip and return home, given the 80% cap on end-of-charge SOC. When the trip requires little additional energy, starting near 25% preserves fast charging while limiting range-anxiety costs. When the trip requires substantially more energy, starting near 25% remains efficient and, relative to starting later, allows the driver to add more usable range before reaching the 80% stopping rule.

In principle, one could also model the state of charge upon arrival at the final destination. In practice, this is empirically difficult because it is highly correlated with $\text{SOC}^{\text{start}}_{iodsj}$, the end-of-charge SOC, and the return-home constraint. I therefore focus on $\text{SOC}^{\text{start}}_{iodsj}$, which is sufficient in the single-stop setting. One could also include an explicit detour cost. In the current specification, the feasible choice set restricts attention to stations that require a detour of at most 20% additional distance relative to the shortest path.[18]

Charging Path Ex-ante Value. Remarkably, despite its simplicity, the charging model generates substantial heterogeneity in route-level charging convenience. For a given trip (o, d) and vehicle type j, each station differs in the utility they provide (with no added value only when $s \notin \mathcal{S}_{odjt}$). This heterogeneity is central because consumers choose destinations while anticipating the charging opportunities available along each origin-destination route, and subsequently choose vehicle models as a function of charging-network quality, equivalently, the convenience of using the network given a vehicle model and city of origin.

In particular, consumers observe the deterministic component of station utility when planning a trip, but the EVT1 taste shock is realized only at the time of travel and captures idiosyncratic, transitory motives (e.g., an unexpected need to stop for water or a washroom break) that are unobserved by the researcher. I therefore summarize the charging component of an origin-destination trip by its *ex-ante* value, which can be embedded directly in the destination-choice problem.

To model destination choice, I treat charging as an *inside-the-trip* discrete choice that enters the generalized cost of traveling from origin o to destination d. For a given (o, d) pair and vehicle range j, the consumer anticipates choosing among feasible charging options along the route. Under the same Type-I extreme value assumption used in the charging-station choice model, this lower-level problem yields a log-sum (inclusive value) that summarizes the attractiveness of available charging opportunities.

Formally, let $\bar{V}^{c}_{odjst}$ denote the mean utility of selecting charging option s for trips between (o, d) for car model j (net of the idiosyncratic EVT1 shock), and let $\mathcal{S}_{odjt}$ be the corresponding

[18] An explicit detour-cost term will be incorporated in future versions.

feasible set. The *ex-ante* value of the charging subproblem is

$$OD_{odjt} \equiv \mathbb{E}\left[\max_{s\in\mathcal{S}_{odjt}} \left\{\bar{V}^c_{odjst} + \varepsilon_{odjst}, V_0 + \varepsilon_0\right\}\right] = \log\left(1 + \sum_{s\in\mathcal{S}_{odjt}} \exp\left(\bar{V}^c_{odjst}\right)\right), \tag{7}$$

up to the usual normalization constant. Intuitively, OD_{odjt} increases when the route offers more feasible charging options, or when existing options provide higher mean utility (e.g., better timing relative to the preferred state of charge). I normalize the outside option to $V_0 = 0$, which captures the case in which the driver chooses not to fast-charge, or equivalently when no fast-charging option is feasible, $\mathcal{S}_{odjt} = \emptyset$. One can interpret this outside option as relying on an alternative mode (e.g., borrowing another gasoline vehicle or taking the train). In that case, the charging network delivers no value for the (o, d) trip, so the inclusive value reduces to $\log(1) = 0$. Each charger therefore affects network value either on the extensive margin, by making additional origin-destination pairs feasible, or on the intensive margin, by providing more (and perhaps better) options along feasible routes. In the destination-choice equation, OD_{odjt} thus enters as a single index summarizing the expected charging convenience of traveling from o to d for vehicle type j.

4.2.2 Destination Choice

This subsection specifies the consumer's destination decision. Conditional on residing in origin city o, charging network $\mathcal{K}_t$ and owning vehicle model j, consumer i chooses a destination d from the feasible set $\mathcal{D}_{ojt}$ to maximize indirect utility. The feasible set $\mathcal{D}_{ojt}$ contains destinations that can be completed with at most one charging stop given the vehicle's range and the battery constraints embedded in the charging problem (See Section 5.1.2 for details on $\mathcal{D}_{ojt}$).

Utility from choosing destination d is

$$V^d_{iodj} = \lambda\, d_{od} \;+\; \zeta_d \;+\; \rho \cdot OD_{odj} \;+\; \varepsilon^d_{iodj}, \qquad d \in \mathcal{D}_{ojt} \tag{8}$$

The term d_{od} denotes origin-destination distance and λ captures the marginal disutility of distance. The destination term ζ_d captures baseline destination attractiveness common to all consumers in o with vehicle type j. The term OD_{odj} is the charging-path inclusive value defined in (7); it summarizes the *ex-ante* convenience of completing the (o, d) trip with one charging stop, given the charging network and model j. I normalize its coefficient to $\rho = 1$, so that OD_{odj} enters destination utility in the same units as the lower-level systematic utilities (as an expected-max utility term), consistent with standard inclusive-value logic. Finally,

ε^d_{iodj} is an EVT1 taste shock, implying a logit choice rule over destinations conditional on feasibility.

The key modeling implication is that destination utility depends not only on distance and destination attractiveness, but also on the expected quality of charging along the route. All else equal, destinations whose charging opportunities align better with the consumer's preferred charging onset generate a higher OD_{odj} and are therefore more likely to be chosen.

4.2.3 Network Quality

This subsection defines the microfounded, model-based measure of network quality, Λ_{ojt}, which summarizes the ex-ante value of access to the public charging network for a consumer residing in origin o, owning vehicle model j, at time t. The key challenge is aggregation: the network delivers value at the *route* level, through the charging opportunities available along a specific origin-destination trip, yet consumers evaluate vehicles *before* knowing which long trip they will need to undertake. I therefore construct Λ_{ojt} as the expected charging convenience across potential destinations, using origin-specific trip-demand weights that are independent of the charging network and reflect consumers' baseline trip-pattern preferences.

The charging network affects EV adoption primarily by relaxing charging frictions for long trips. Prior to purchasing a vehicle, a consumer in origin o faces uncertainty about which destination d will arise. Some destinations are relevant and frequent (e.g., nearby regional centers), while others are rare. A network that improves charging convenience on frequently traveled routes is more valuable than a network that improves convenience on rarely traveled routes. A meaningful index of network quality must therefore weight route-level charging value by the likelihood that each destination is realized.

Crucially, these weights should capture *baseline trip demand* rather than demand that is itself altered by the charging network. Otherwise, the network would enter both through route-level charging value and through the destination weights, mechanically amplifying network effects. I therefore define destination weights using only non-charging determinants of destination choice: destination attractiveness and travel distance.

Let $\mathcal{D}$ denote the set of candidate destinations. Prior to vehicle purchase, I model destination uncertainty by assuming that consumer i draws a finite set of potential trip destinations, $\mathcal{J}_{io} = \{d_{i1}, \ldots, d_{iM}\}$, with elements drawn i.i.d. from an origin-specific distribution $P_0(d \mid o)$. This interpretation is useful because it makes clear how aggregation works: the realized ex-post value of the network depends on the random set $\mathcal{J}_{io}$, while Λ_{ojt} is its ex-ante expectation under $P_0(\cdot \mid o)$. In particular, by the law of large numbers, the empirical frequency of destinations in $\mathcal{J}_{io}$ converges in expectation to $P_0(d \mid o)$, so that integrating over draws of $\mathcal{J}_{io}$ yields destination weights equal to $P_0(d \mid o)$.

I parameterize $P_0(d \mid o)$ using a standard logit structure based origin-destionation preference, excluding the charging component resulting in (i) a destination fixed effect ζ_d, capturing baseline destination attractiveness common across origins, and (ii) origin-destination distance d_{od}, entering with marginal disutility $\lambda < 0$. The resulting origin-specific baseline destination distribution is

$$P_0(d \mid o) = \frac{\exp(\zeta_d + \lambda d_{od})}{\sum_{d' \in \mathcal{D}} \exp(\zeta_{d'} + \lambda d_{od'})} \tag{9}$$

This choice of $P_0(d \mid o)$ is convenient for three reasons. First, it provides a disciplined, origin-specific weighting that assigns higher mass to nearby and intrinsically attractive destinations. Second, it is consistent with the discrete-choice structure and can be estimated in the destination-choice step. Third, by excluding charging convenience, it yields weights that are invariant to the charging network, avoiding double counting when constructing network quality.[19]

Network quality index. For each origin-destination pair (o, d), vehicle model j, and time t, the charging network determines the ex-ante value of charging, OD_{odjt}, defined in (7). I aggregate these route-level objects into a single measure of network quality as the expected charging quality weighted by destinations preferences:

$$\Lambda_{ojt} \;\equiv\; \mathbb{E}[\mathcal{K}_{ojt}] = \sum_{d \in \mathcal{D}} P_0(d \mid o)\, OD_{odjt} \tag{10}$$

Intuitively, Λ_{ojt} is high when the network provides convenient charging opportunities for the destinations that are most likely to arise for residents of o. Because OD_{odjt} depends on feasibility and on the distribution of charging utility across stations on the route, Λ_{ojt} captures both extensive-margin improvements (more feasible routes) and intensive-margin improvements (better-timed charging opportunities) delivered by the network. More precisely, if no charging station is feasible for a given origin-destination pair, i.e. $\mathcal{S}_{odjt} = \emptyset$, then $OD_{odjt} = \log(1) = 0$. Conversely, adding any feasible charging option increases the route value. In particular, if $s^\star \notin \mathcal{S}_{odjt}$ becomes feasible, then OD_{odjt} increases. The marginal gain is larger when $s^\star$ delivers higher mean utility: a well-located station along the route adds more value than a poorly located (or infeasible) one.

[19] Moreover, the framework can be extended to allow the baseline destination distribution $P_0(d \mid o)$ to be consumer-specific. For example, one could leverage observed transactions to construct a consumer-specific destination choice set $\mathcal{J}_i$, and treat $P_0(d \mid o)$ as a prior that is updated using usage data as the likelihood. Alternatively, one could allow preferences to be heterogeneous across consumers by introducing idiosyncratic variation in the draws from $P_0(d \mid o)$.

4.2.4 Charging Demand

While the first stage of the demand estimation links the charging network $\mathcal{K}$ to EV adoption through network quality, $\mathcal{E}(*)$, consumer trip and charging choices also matter on the supply side through implied charging demand. Let $\mathcal{C}(\mathcal{K})$ denote the mapping from the charging network $\mathcal{K}$ to charging demand.

Intuitively, $\mathcal{C}(\mathcal{K})$ aggregates i) consumers' EV adoption, ii) destination choices and iii) charging-station choices to deliver the volume of charging implied by a given network and its associated EV fleet. This aggregation collapses complex trip and charging patterns into a single demand scalar.[20]

Formally, consider a charging node r. Node r generates charging demand only if it hosts at least one station, i.e. $n_r \geq 1$. Let $s \in r$ denote a station located in node r. Let $\mathcal{E}_{oj}$ denote the EV stock of model j in origin city o, and let M denote the number of trips per vehicle (or, equivalently, the expected number of trips generated by each EV) over the relevant period. The implied charging demand at node r is

$$Q_r(\mathcal{K}, \mathcal{E}) \;=\; \sum_{o} \sum_{j} \mathcal{E}_{oj}(\mathcal{K})\, M \sum_{d \in \mathcal{D}_{oj}} P(d \mid o, j) \sum_{s \in r} P\big(s \mid o, j, d\big), \tag{11}$$

where $P(d \mid o, j)$ is the probability that an EV owner of model j originating from o selects destination d, and $P\big(s \mid o, j, d\big)$ is the probability that the trip (o, d) charges at station s. Summing over $s \in r$ aggregates the probability mass of charging at any station located in node r. Note that Equation (11) is used only for the counterfactual analysis. Because the counterfactual is a static game, I suppress time variation and drop the t subscript.

4.3 Supply

Charging firms take the public first-mover build-out as given and, simultaneously, make entry decisions based on sunk entry costs and the discounted stream of expected profits in the spirit of Springel (2021). The entry model follows Bresnahan and Reiss (1991). The number of charging stations in a market is the equilibrium outcome of a complete-information entry game in which the installed base of EVs determines market size and charging demand. Demand implies an annual market potential for charging sessions in region r, denoted Q_r, and a per-station demand $q_r = \frac{Q_r}{n_r}$ that declines with the number of active stations. Given per-session revenue net of variable costs and a fixed operating cost, entry continues until the marginal entrant earns zero expected profit.

[20] EV adoption, $\mathcal{E}(\mathcal{K})$, is implied by the network, $\mathcal{K}$

4.3.1 Charger Profit Function

The number of charging stations in a market is the equilibrium outcome of a complete-information entry game with free entry. A charging station's entry decision depends on its sunk entry cost and expected per-period profits. Non-recurring fixed costs include the cost of charging equipment and labor associated with installation. Per-period profits are determined by charging demand and markups. Because the charging network generates two types of externalities—within the charging network itself and in the EV market—I distinguish two types of firms. [21] [22] [23] [24]

First, a *single-owner* firm operates a single station, representing a benchmark monopolistic economy without network externalities. Second, a *monopolist* firm (m) owns multiple stations across space and internalizes charging-network externalities.

$$\pi_r^{so} = \underbrace{[p - c] \cdot q_r^{so}}_{\text{Per-period profit}} - \underbrace{F}_{\text{Sunk fixed cost}} \tag{12}$$

$$\pi^m = \sum_{r \in R} \Big([p - c] \cdot Q_r - F \Big) \tag{13}$$

In estimation, the supply side is modelled under the single-owner assumption unless precise otherwise.

4.3.2 Entry Decision

I follow Bresnahan and Reiss (1991) where firm makes free entry decision, that is, they enter when their expected payoff is amount the entry cost. The single owner station has one strategy entering with $n_r^{so} = 1$ or not.

$$\pi_r^{so} = \mathbf{1}\{n_r^{so} = 1\} \Big((p - c)\, q_r - F \Big) \tag{14}$$

On the other hand, the monopoly behaves differently. As the sole player, it internalizes the externalities associated with network deployment. Namely, EV adoption depends not only on the network $\mathcal{E}(\mathcal{K})$ but also on the charging flows $\mathcal{C}(\mathcal{K}, \mathcal{E})$. As such, the quantity of charges Q_r is directly related to the constructed network.

[21] The discounted stream of expected profits is obtained from steady-state per-period profits, assuming a seven-year revenue horizon and a 15% discount rate.

[22] https://www.mckinsey.com/features/mckinsey-center-for-future-mobility/our-insights/can-public-ev-fast-charging-stations-be-profitable-in-the-united-states

[23] https://docs.nrel.gov/docs/fy23osti/85654.pdf

[24] Li (2019) has a flow cost of ≈ $32,000 USD per station, per quarter.

The local Manager Argument. The strategy is not solved through joint optimization. Instead, the monopoly delegates decisions to fully informed local managers in each region, who maximize network profit. Managers act sequentially, beginning in the network core (Montréal) and expanding progressively outward toward the periphery.

$$\frac{\partial \pi_r^m}{\partial n_r} = \frac{\partial}{\partial n_r}\left(\sum_{r\in\mathcal{R}}\Big((p-c)Q_r - F\Big)\right) \tag{15}$$

A new charging station is built in region r whenever the marginal increase in network profit exceeds the fixed cost F, that is,

$$\frac{\partial \pi_r^m}{\partial n_r} \geq F.$$

Each local manager perform this decision sequentially. The game is iterated once no local manager or consumer benefits from changing its strategy.

4.4 Equilibrium

Let $\mathcal{R}$ denote regions (nodes). A charging network is $\mathcal{K} = \{k_r\}_{r\in\mathcal{R}}$, where $k_r \in \mathbb{N}$ is the number of chargers in region r. Let $\Lambda(\mathcal{K})$ be the (vector of) network-quality sufficient statistic generated by consumers' expected maximum value of performing a single trip given charging network structure $\mathcal{K}$. Let $\mathcal{E}(\Lambda)$ denote the resulting EV adoption and model shares (think EV stock) from the vehicle-demand problem given network quality Λ. Let $\mathcal{C}(\mathcal{K}, \mathcal{E})$ denote the charging flow that maps the EV stock $\mathcal{E}(\Lambda)$ and network quality $\mathcal{K}$ jointly into charging demand. Let $\mathcal{C}_r(\mathcal{K}, \mathcal{E})$ denote total charging demand in region r implied by charging behavior, and assume that demand is split evenly across chargers within a region. Let $\pi_r(\mathcal{K}, \mathcal{E})$ denote the expected per-charger profit in region r, given a per-charger fixed cost F and variable cost $c_r(\cdot)$.

An equilibrium is a triple $(\mathcal{K}^*, \Lambda^*, \mathcal{E}^*)$ such that:

1. **Trip/charging consistency:** $\Lambda^* = \Lambda(\mathcal{K}^*)$, i.e., network quality equals the sufficient statistic generated by consumers' optimal trip/charging behavior given the realized network.

2. **Vehicle-demand optimality:** $\mathcal{E}^* = \mathcal{E}(\Lambda^*)$, i.e., consumers' EV adoption are optimal given network quality. (Vehicle prices are taken as exogenous.)

3. **Free-entry equilibrium:** for each region $r \in \mathcal{R}$, the number of chargers k_r^* satisfies

the integer free-entry conditions

$$\pi_r(\mathcal{K}^*, \mathcal{E}^*) \geq 0 \quad \text{if } k_r^* \geq 1, \qquad \pi_r(\mathcal{K}^* + \mathbf{e}_r, \mathcal{E}^*) \leq 0,$$

where $\mathbf{e}_r$ denotes adding one additional charger in region r (holding all other regions fixed). Equivalently, the realized network is a free-entry outcome among symmetric potential entrants: no firm can increase expected profit by unilaterally adding a charger, and any charger that operates earns non-negative expected profit.

Conditions (1)–(3) jointly impose belief consistency: the network used to form beliefs and compute Λ and $\mathcal{E}$ coincides with the network implied by firms' entry decisions.

Algorithm. Given parameters, I compute $(\mathcal{K}^*, \Lambda^*, \mathcal{E}^*)$ by iterating on the fixed point in $\mathcal{K}$:

1. Initialize $\tilde{\mathcal{K}} \leftarrow \mathcal{K}_{2022}^{obs}$
2. **Consumers' beliefs:** set $\mathcal{K} \leftarrow \tilde{\mathcal{K}}$
3. **Trip/charging stage:** given $\mathcal{K}$, solve consumers' trip/charging problem and compute $\Lambda \leftarrow \Lambda(\mathcal{K})$
4. **Vehicle stage:** given Λ, solve vehicle demand and compute EV adoption $\mathcal{E} \leftarrow \mathcal{E}(\Lambda)$
5. **Firms' entry stage:**
   (a) Set firms' belief $\hat{\mathcal{K}} \leftarrow \tilde{\mathcal{K}}$
   (b) Given $(\hat{\mathcal{K}}, \mathcal{E})$, form implied charging demand $\mathcal{C}(\hat{\mathcal{K}}, \mathcal{E})$ and profits
   (c) Update entry in each region using the free-entry rule, yielding an updated network $\tilde{\mathcal{K}} \leftarrow \mathrm{FE}(\mathcal{E}; \hat{\mathcal{K}})$.
   (d) Iterate the inner loop until the network no longer changes, $\tilde{\mathcal{K}} = \hat{\mathcal{K}}$
6. If $\|\mathcal{K} - \tilde{\mathcal{K}}\|$ is above tolerance, return to step 2; otherwise set $\mathcal{K}^* = \tilde{\mathcal{K}}$, $\Lambda^* = \Lambda(\mathcal{K}^*)$, and $\mathcal{E}^* = \mathcal{E}(\Lambda^*)$.

5 Estimation

This section presents the estimation strategy and the associated identification logic. The empirical analysis proceeds in two linked blocks. The first block focuses on trips and charging behavior and is used to recover the ex-ante expected value of the charging network. The second block uses this object as an input into the EV adoption model. While the two

blocks are estimated separately, they are structurally connected: charging behavior identifies network quality, which in turn enters vehicle demand.

5.1 Network Quality and Trip Demand

I recover the ex-ante expected value of using the charging network, denoted $\mathbb{E}[\mathcal{K}_{ojt}]$, which depends on the spatial configuration of the charging network, the origin city o, and the vehicle model j through its range. Consumers derive utility from undertaking trips to destinations d drawn from a city-specific distribution, $\mathcal{J}_i \sim_{\text{i.i.d.}} P(d \mid o)$. For each realized trip, utility is primarily affected by the quality of the charging network available along the origin-destination pair.

The estimation of network quality proceeds in two steps. First, I estimate charging-behavior parameters at the trip level. Conditional on an observed (o, d) trip, this step recovers the cost of charging away from a preferred target state of charge. The parameter of interest captures how costly it is for consumers to deviate from their preferred charging level, which is centered around 25%. This target emerges as optimal for three reasons.

First, it balances charging efficiency, as marginal charging speeds begin to decline at higher states of charge. Second, it minimizes range anxiety: while lower battery levels would be mechanically more efficient, they are psychologically costly. Third, the closer the charge occurs to the final destination, the higher the battery level upon arrival at home, and the less charging is required given reduced information asymmetry about subsequent travel needs. Notably, even long-range vehicles initiate charging at around 25%, highlighting that this behavior reflects a combination of mechanical and psychological factors rather than vehicle range per se.

Second, I estimate destination choice. Prior to departure, consumers do not condition on a specific charging station, but they do form expectations over the value of the charging network available along each (o, d) route. These expectations can be interpreted as reduced-form preference shocks—analogous to thirst, hunger, or restroom needs—that govern the timing and intensity of charging. Given (i) the destination-specific payoff from reaching d, which is common across consumers, (ii) the travel cost associated with the (o, d) pair, and (iii) the ex-ante expected value of network usage, $\mathbb{E}[\mathcal{K}_{ojt}]$, consumers make a discrete choice over destinations.

Together, charging behavior and destination choices identify $\mathbb{E}[\mathcal{K}_{ojt}]$ from revealed preferences. This object summarizes how the charging network affects ex-ante travel decisions and serves as a sufficient statistic for network quality.

5.1.1 Charging Parameters

This subsection estimates the importance of the battery level at the start of charging. Using all single-charge returning trips observed in 2023, I estimate a maximum likelihood model of charging-station choice. The likelihood compares the realized choice to the feasible set of charging stations, $\mathcal{S}_{odjt}$, which depends on the origin-destination pair (o, d) and the vehicle model j.

The identifying assumption is that the 2023 charging network is sufficiently dense and well connected that, conditional on feasibility, observed station choices primarily reflect preferences rather than binding network constraints. Formally, when a consumer chooses station s over another feasible station $s' \in \mathcal{S}_{odjt}$, this implies that s yields higher deterministic utility than s' up to an idiosyncratic EVT1 preference shock. That is, station s is preferred to s' in expectation, with deviations governed by the extreme-value error structure. Variation in observed choices therefore reveals preferences over *when* and *where* to charge. In particular, this setting allows me to recover consumers' preferred charging state of charge and the cost of deviating from it.

These estimates are central to the construction of network quality. A charging network generates value not only by being available, but also by placing chargers in locations that align with consumers' desired charging points along their trips. Importantly, these preferred locations vary across origin-destination pairs, implying that network quality is inherently route-specific.

Station Choice Set. To reduce dimensionality and to accurately represent the consumer's effective choice set for trip n, I restrict the feasible set of charging stations, $\mathcal{S}_{odjt}$, using a set of economically motivated constraints. First, I require stations to lie along the travel corridor (R1). Second, I limit detours by imposing a maximum detour ratio: the total path length via the station cannot exceed 1.2 times the direct distance between origin and destination (R2). Finally, feasibility is constrained by battery availability: the remaining battery range at the station must be sufficient to reach the station conditional on a max state of the battery of 80% at the end of the charge (MC). Together, these restrictions ensure that $\mathcal{S}_{odjt}$ captures the stations that are physically reachable, behaviorally plausible, and relevant for the consumer's charging decision.

Inferring Destination. I do not observe destinations directly in the data. However, I observe a near-universe of fast-charging transactions at high frequency. I exploit this richness to restrict attention to a subset of charging events that are informative about travel patterns. In particular, I focus on single-charge trips, defined as cases in which a consumer performs no other fast charge within a 24-hour window. These trips plausibly correspond to origin-destination travel completed with a single public charge. The identifying logic is that

a well-functioning one-stop charging network is a sufficient statistic for a more complex n-stop charging network. Additional intermediate charging opportunities constitute a positive transformation of the feasible set: they weakly expand reachability and reduce effective travel costs, but do not alter the ranking of destinations or the structure of optimal routing.

Conditional on this restriction, a consumer may charge either before reaching the final destination (outbound) or on the way back (returning). I therefore cluster charging events into two groups. Consistent with consumers' preferred charging target of approximately 25%, about 80% of observed charges are classified as returning trips. This restriction is central for identification.

For returning trips, I observe both the origin-station distance and the battery state of charge at the start of charging. Together with vehicle-specific range, these objects imply a feasible set of destinations, denoted $\mathcal{D}_{osj}$, given the origin, charging station, and vehicle model. I infer the final destination by selecting the location that best matches the implied residual distance (the "missing kilometers") after accounting for the observed origin-station leg. Additional details on the destination inference procedure are provided in Appendix A.1.

Estimating equation. Conditional on a realized origin-destination pair (o, d), I model the consumer's choice of charging station as a standard discrete-choice problem over a feasible set of stations. For each trip n, let $y_{ns} = 1$ if station s is chosen and $y_{ns} = 0$ otherwise. The feasible choice set, denoted $\mathcal{S}_{odjt}$, includes the stations that are physically reachable and behaviorally plausible for the (o, d) trip under the vehicle's range and battery constraints.

I assume that the indirect utility from choosing station s for trip n is

$$V_{ns}(\theta_s) \;=\; \eta\, b_{ns},\; + \varepsilon_{ns}, \qquad \theta_s = (\eta) \tag{16}$$

where ε_{ns} is i.i.d. Type-I extreme value. Here, η governs the cost of deviating from the preferred battery level at the start of charging, which is centered at 25%. The penalty is specified as a quadratic function, $(b_{ns} - 0.25)^2$, where b_{ns} denotes the battery state of charge upon arrival at station s. This term captures the disutility of charging away from the preferred state of charge, reflecting range anxiety and charge efficiency.

Under the EVT1 assumption, the implied choice probability is the multinomial logit:

$$P_{ns}(s \mid o, d; \theta_s) = \frac{\exp(\eta\, b_{ns})}{\displaystyle\sum_{s' \in \mathcal{S}_{odjt}} \exp(\eta\, b_{ns'})} \tag{17}$$

Parameters are estimated by maximum likelihood using observed station choices:

$$\max_{\theta_s} \log \mathcal{L}(\theta_s) = \sum_{n} \sum_{s \in \mathcal{S}_{odjt}} y_{ns} \log P_{ns}(s \mid o, d; \theta_s) \tag{18}$$

This step recovers how station infrastructure and the battery-state channel shape charging-station choice, holding fixed the origin-destination pair and restricting attention to stations in the feasible set.

Estimation Results. Table 5 reports maximum likelihood estimates from the charging-station choice model. Column (1) presents the main specification and identifies a sharp preference for initiating a fast charge near a target state of charge of 25%. The coefficient η is negative and precisely estimated ($\hat{\eta} = -0.056$, s.e. $= 0.0019$), implying meaningful utility losses when drivers arrive to charge substantially above or below the 25% target. This suggests that drivers either wait to charge until the battery is around 25% or, alternatively, choose to charge earlier rather than reaching very low state of charge. Economically, the estimate indicates that charging decisions are not driven solely by reachability: conditional on feasibility, consumers actively time charging to occur at a preferred battery level, capturing both charging-efficiency considerations and range-anxiety costs.

Overall, these results imply that an effective charging network must offer sufficiently uniform spatial coverage so that, for any origin-destination pair, consumers can not only access a charging station, but can do so at or near their preferred charging level.

Table 5: Result of Charging Behavior Estimation

	(1)
Battery Deviation from 25% at Start, η	-0.056^{***}
	(0.0019)
N Observations	3,806,690

Notes: This table reports maximum likelihood estimates from a charging-station choice model using the full set of single-stop, round-trip journeys observed in 2023. For each origin-destination-vehicle combination, the feasible charging set $\mathcal{S}_{odjt}$ is defined based on vehicle range and network geometry. Candidate charging stations are restricted to those (i) located between the origin and the destination, (ii) inducing a detour of at most 20% relative to the shortest path, and (iii) allowing the vehicle to complete the origin-destination-station-origin trip with strictly positive remaining battery, assuming a maximum charging stop of 80%. I report standard errors based on the inverse of the quasi-Newton Hessian approximation returned by the MLE routine.

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

5.1.2 Destination Parameters

This subsection estimates long-trip destination demand. For each origin city o and vehicle model j, I construct a feasible set of destinations $\mathcal{D}_{ojt}$ using a simple reachability rule. A destination d is included if it can be reached with at most one charge under a conservative range extension, implemented as a reachable factor of 1.6 (i.e., up to 160% of the vehicle's range). This proxy corresponds to a one-stop long trip that starts charging at a low battery level (around 20%) and ends at a high battery level (around 80%).

For each feasible origin-destination pair and model j, I evaluate the associated one-stop route problem, denoted OD_{odj}. This step summarizes how convenient the charging network is for traveling from o to d with model j—in particular, whether the route offers charging opportunities in locations that align with consumers' preferred charging points. The destination choice model then compares the realized destination to this feasible set, attributing systematic destination patterns to both (i) baseline destination attractiveness and (ii) route-specific charging convenience.

These destination parameters quantify how much the charging network shifts long-trip behavior by making certain destinations effectively easier to reach and more comfortable to undertake. In turn, they provide a microfoundation for network quality, $\mathbb{E}[\mathcal{K}_{ojt}]$.

Estimating equation. Conditional on an origin city o and vehicle model j, I model the consumer's destination choice as a standard discrete-choice problem over a feasible set of destinations. For each trip n, let $y_{nd} = 1$ if destination d is chosen and $y_{nd} = 0$ otherwise. The feasible choice set, denoted $\mathcal{D}_{ojt}$, includes destinations that are reachable with at most one charging stop given the vehicle's range and battery constraints.

I assume that the indirect utility from choosing destination d for trip n is

$$V_{nd}(\theta_d) \;=\; \lambda\, d_{od} \;+\; \zeta_d \;+\; OD_{odj}(\theta_s) \;+\; \varepsilon_{nd}, \qquad \theta_d = \big(\lambda, \{\zeta_d\}\big), \tag{19}$$

where ε_{nd} is i.i.d. Type-I extreme value. The term d_{od} denotes origin-destination distance and λ captures the marginal disutility of distance. The destination fixed effect ζ_d absorbs time-invariant differences in baseline destination attractiveness. Finally, $OD_{odj}(\theta_s)$ summarizes the ex-ante expected convenience of completing the (o, d) trip with one charging stop, constructed using the charging-stage parameters θ_s. Intuitively, this term shifts destination utility toward routes whose charging opportunities better align with consumers' preferred charging points.

Under the EVT1 assumption, the implied choice probability is the multinomial logit:

$$P_{nd}(d \mid o; \theta_d) = \frac{\exp(\lambda\, d_{od} + \zeta_d + OD_{odj}(\theta_s))}{\sum_{d' \in \mathcal{D}_{oj}} \exp(\lambda\, d_{od'} + \zeta_{d'} + OD_{od'j}(\theta_s))} \tag{20}$$

Parameters are estimated by maximum likelihood using observed destination choices:

$$\max_{\theta_d} \; \log \mathcal{L}(\theta_d) = \sum_n \sum_{d \in \mathcal{D}_{oj}} y_{nd} \, \log P_{nd}(d \mid o; \theta_d) \tag{21}$$

This step recovers how distance, baseline destination appeal, and route-specific charging convenience jointly shape long-trip destination choice, conditional on feasibility and taking charging behavior as an ex-ante object.

Estimation Results. Table 6 reports maximum likelihood estimates from the destination-choice model. The specification includes a single continuous cost shifter—origin-destination distance—and a full set of destination fixed effects, $\{\zeta_d\}$, which absorb time-invariant differences in baseline destination attractiveness (e.g., amenities, population, employment centers, tourism appeal, and other persistent attributes).

The estimated distance coefficient is negative, $\hat{\lambda} = -0.0282$, consistent with standard spatial demand: conditional on feasibility and destination fixed effects, consumers weakly prefer closer destinations. However, the estimate is imprecise. The standard error is 0.0324, implying that the coefficient is not statistically distinguishable from zero at conventional levels. Economically, this indicates that, within the set of feasible one-stop destinations, the data do not provide strong evidence of systematic distance aversion beyond what is already captured by charging network quality.

Because ζ_d is included, the distance coefficient is identified purely from within-destination variation across origins and vehicle models (through differences in d_{od} across o for a given d and variation in feasibility across j). The inclusion of $\{\zeta_d\}$ is consequential: it ensures that $\hat{\lambda}$ is not mechanically picking up that some destinations are both popular and geographically closer to many origins.

There are two main reasons why $\hat{\lambda}$ can be weakly identified in this setting. First, the feasible choice set $\mathcal{D}_{ojt}$ is defined by a one-stop reachability rule, which truncates the support of distances by construction and can compress residual distance variation among alternatives. Second, the model conditions on the ex-ante route object OD_{odj}, so a substantial portion of systematic route-level variation that otherwise correlates with distance is absorbed by charging convenience. Both features reduce the incremental explanatory power of d_{od} among feasible destinations.

This result is neither surprising nor problematic. In this setting, the primary determinants of whether a long trip is undertaken are the destination-specific components $\{\zeta_d\}$ and the charging-path value OD_{odj}, which already internalize much of the effective cost of distance. Intuitively, the destination fixed effects $\{\zeta_d\}$ should be interpreted as answering the following question: *conditional on undertaking a long trip that requires a charging stop, which feasible destinations deliver the highest net value relative to their travel cost?* The full set of destination fixed effects is reported in Appendix Table A.2.

Table 6: Results of Destination Estimation

	(1)
Origin-Destination Distance, λ	-0.0282
	(0.0324)
ζ_d	Yes
Observations	686,421

Notes: This table reports maximum likelihood estimates from a destination choice model using the full set of single-stop, round-trip journeys observed in 2023. Each observation is an origin-destination journey for a given car model, and the choice set consists of feasible destinations. Feasibility is determined by the charging network quality measure D_{ojt} and origin-destination characteristics. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Although the distance coefficient is imprecise, the estimation still serves its primary purpose: recovering the destination-side components needed to map the charging network value. In particular, the destination fixed effects capture baseline destination demand, while the route object OD_{odj} links charging convenience to destination choice. This varies across network quality, while the (o, d) value remain constant across network. Together with the charging-station choice estimates, destination choices therefore help identify $\mathbb{E}[\mathcal{K}_{ojt}]$ from revealed preferences. This object summarizes how the charging network affects ex-ante travel decisions and serves as a sufficient statistic for network quality. In the next subsection, I incorporate $\mathbb{E}[\mathcal{K}_{ojt}]$ into the EV adoption model, where it is the primary parameter of interest.

5.2 EV Demand

The estimating equation follows the discrete–choice framework of Berry et al. (1993). It implies the standard logit:

$$\log S_{ojt} - \log S_{ot}^{O} = \delta x_{jt} - \alpha p_{jt} + \beta \Lambda_{ojt} + \mu_{ojt},$$

where S_{omt} denotes the market share of model m in city o in year t, including only battery electric vehicles (BEVs). The outside option, S_{ot}^{O}, corresponds to the share of non–EVs. In this specification, x_m collects vehicle characteristics and controls, p_{jt} is the price of model m, and Λ_{omt} measures the charging–network quality relevant to consumers in city o considering model m. The consumer-vehicle-specific taste shock, μ_{iomt}, represents unobserved factors affecting utility, assumed to be i.i.d. Type-I extreme value distributed. The outside option is choosing a non-EV, with its indirect utility normalized to $U_{io0} = \mu_{io0}$. The charging network quality measures are described in detail above. Consumer i chooses vehicle j if $U_{ijo} \geq U_{ij'o}$ for all j'. The market share for vehicle j is given by the closed-form expression:

$$S_{ojt} = \frac{\exp(\delta x_{jt} - \alpha p_{jt} + \beta \Lambda_{ojt})}{\sum_{j't} \exp(\delta x_{j't} - \alpha p_{j't} + \beta \Lambda_{j'ot})} \tag{22}$$

The estimation exploits three sources of variation. First, the network varies across vehicle models: for a given city and network, vehicle range interacts with desired trip lengths and overall network quality. Second, the charging network varies across cities, as each car model faces a distinct network quality within a given city. Third, the charging network evolves over time, with public provision being the primary driver. This time variation is crucial for identification, as public station deployment improves network quality in ways that are not explained by implicit profitability.

Consumers pay a purchase price p_j, which is equal to the MSRP minus government subsidies:

$$p_{jt} = \text{MSRP}_{jt} - \text{Province Subsidy}_{jt} - \text{Federal Subsidy}_{jt}$$

MSRP and manufacturer discounts are the same across all markets and vary only by vehicle models and time. Provincial subsidies vary across vehicle models, provinces, and time, while federal subsidies vary across vehicle models but remain constant over time during the study period.[25]

Because vehicle prices are endogenously chosen by manufacturers, I use standard BLP-style instruments that exploit variation in both product characteristics and production costs. The first set of instruments is based on the sum of own-brand and rival product characteris-

[25]This definition of the purchase price assumes full pass-through of subsidies to consumers. While full pass-through is imposed rather than estimated in our specification, this assumption is consistent with prior empirical evidence in vehicle markets showing that consumer subsidies are largely passed through to buyers. For example, Sallee (2011) estimates subsidy incidence and finds that federal tax incentives for the Toyota Prius were fully captured by consumers in his setting. Muehlegger and Rapson (2022) estimate pass-through of EV subsidies and find that consumers capture approximately 73-84% of the subsidy in baseline specifications, and they cannot reject full pass-through once policy endogeneity is addressed. Similarly, Busse et al. (2006) estimate pass-through of manufacturer rebates and show that consumers capture a substantial share (roughly 70-90%) of customer-directed incentives. Taken together, this evidence suggests that assuming full pass-through provides a reasonable approximation in the context we study.

tics, which shifts relative markups while remaining exogenous to unobserved demand shocks. To strengthen identification, I augment these with cost shifters. Specifically, I construct a composite index of key input prices—including lithium—that captures fluctuations in global commodity markets affecting marginal costs but unrelated to local consumer demand. I also interact this price with the exchange rate movements provide an additional exogenous source of variation by influencing the cost of imported vehicles and components without directly affecting consumer preferences. Together, these instruments introduce supply-side variation in vehicle prices that is plausibly orthogonal to unobserved demand shocks and thus suitable for consistent estimation.

Estimation Results. Table 7 reports logit demand estimates using city-year-model observations and a rich set of fixed effects. Across all specifications, the price coefficient is negative and precisely estimated, implying economically meaningful price sensitivity.

The network-quality index, $E[K_{ojt}]$, enters positively in columns (1), (2) and (3) and is statistically significant in columns (1) and (2), indicating that improvements in charging network quality increase BEV demand. The magnitude is stable across the baseline specifications. In column (3), where $E[K_{ojt}]$ is interacted with range, the point estimate on $E[K_{ojt}]$ becomes imprecise, consistent with limited independent variation once the interaction is included and with the correlation between network quality and range-related travel capability.

Range enters with a negative coefficient in columns (2) and (3), but this should not be interpreted as consumers disliking range. The main reason is collinearity: $E[K_{ojt}]$ is constructed using range-dependent feasibility, so Range mechanically co-moves with the network-quality regressor, leaving little independent variation in Range once $E[K_{ojt}]$ is included. In addition, brand and high-dimensional time-varying fixed effects absorb much of persistent product quality and common shocks, so the remaining range coefficient is identified from residual within-cell variation.

Columns (1.1)-(1.3) assess robustness to alternative fixed-effect structures. The estimated coefficient on $E[K_{ojt}]$ is sensitive to which components are absorbed, but it remains positive and economically meaningful across specifications. The estimate is large and precisely estimated when identification draws on variation across city-years and/or across models (column (1.1)). It attenuates and becomes statistically insignificant once ξ_{jt} (Product×Year fixed effects) absorb nearly all within-model-year variation, leaving identification to rely primarily on cross-city differences within a given model-year (column (1.2)). This implies that cross-city differences within a model-year are not driving the results; instead, the effect is identified from changes in charging-network quality over time. Column (1.3) clarifies the key source of identification: within a given city and vehicle type (e.g., SUV, sedan), changes

over time in local network quality. This indicates that consumers respond to changes in network quality; once Product×Year fixed effects are included, we recover estimates close to the baseline specification (1). Overall, the evidence suggests that the network effect is identified mainly from within-city, within-type changes in network quality over time, rather than from time-invariant city characteristics or common year shocks.

Table 7: Result of EV Demand Estimation

	Alternative Specifications			FE Decomposition of Spec. (1)		
	(1)	(2)	(3)	(1.1)	(1.2)	(1.3)
Price	-0.023^{***}	-0.016^{***}	-0.016^{***}	-0.045^{***}	-0.020^{***}	-0.056^{***}
	(0.004)	(0.004)	(0.004)	(0.003)	(0.004)	(0.004)
$E[K_{ojt}]$	1.406^{***}	1.627^{***}	0.273	2.786^{***}	0.247	6.736^{***}
	(0.281)	(0.322)	(7.401)	(0.590)	(0.230)	(0.994)
Range (per 10 km)		-0.046^{**}	-0.040^{**}			
		(0.019)	(0.021)			
$E[K_{ojt}]\times$ Range (per 10 km)			0.020			
			(0.122)			
Brand Fixed Effects	Yes	Yes	Yes	Yes	Yes	–
ξ_{jt} (Product×Year FEs)	Yes	Yes	Yes	–	Yes	–
Market FEs × Type × Year	Yes	Yes	Yes	–	–	Yes
Supply Instruments	Yes	Yes	Yes	Yes	Yes	Yes
N (Observations)	2,536	2,536	2,536	2,536	2,536	2,536

Notes: Each observation is a city-year-model (2016-2022). The regression is restricted to the BEV; the outside option represents the alternative options (PHEV, Hybrid and ICE). Columns (1)-(3) report alternative demand specifications: (1) includes Price and $E[K_{ojt}]$; (2) adds Range (per 10 km); (3) adds $E[K_{ojt}]\times$Range. Columns (1.1)-(1.3) decompose specification (1) by varying the absorbed fixed effects: (1.1) absorbs firm fixed effects only; (1.2) absorbs Type×Year×Market fixed effects only; (1.3) absorbs firm and Product×Year fixed effects. Standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

6 Counterfactual

The baseline is the Big Push. The government moves first by expanding the charging network, after which charging firms take the public build-out as given and simultaneously make entry decisions based on sunk entry costs and the discounted stream of expected profits. My primary counterfactual removes government intervention and instead computes the laissez-faire outcome. Under laissez-faire, two stable and privately viable Nash equilibria emerge. I decompose the Big Push into two components: (i) a *Coordination Component*, which captures the shift from the Low- to the High-adoption equilibrium under laissez-faire, and (ii)

an *Externality Component*, which captures the additional shift from the High-adoption equilibrium to the Big Push equilibrium induced by the government's first-mover investment.

Coordination (High vs. Low). The coordination component isolates the gain from equilibrium selection under laissez-faire: it measures the effect of the market converging to the High- rather than the Low-adoption Nash equilibrium. EV adoption rises by 139%, network quality increases by 1.06, and per-vehicle WTP rises by about $10,277. In aggregate, consumer surplus increases by $7,167M and climate benefits add $1,429M (SCC=$185). This is the pure "chicken-and-egg" channel: the outcome changes because expectations and best responses coordinate on the High equilibrium.

Externalities (Big Push vs. High). The externality component measures the incremental gain from public first-mover investment once the market is already in the High-adoption equilibrium. Relative to the High equilibrium, the Big Push raises network quality by 0.61 and increases EV adoption by a further 52%, yielding an additional $5,913 in per-vehicle WTP. Consumer surplus rises by $3,991M and climate benefits by $1,290M. These gains come from public provision at high-connectivity nodes that are socially valuable but privately unprofitable, which modestly crowds in entry at the core (14 stations) while displacing some entry at the fringe.

Together, the two counterfactual experiments separate *equilibrium selection* from *network externalities*. The coordination component captures large extensive-margin gains from moving from Low to High adoption, while the externality component captures intensive-margin gains from improving network topology within the High state by internalizing unpriced spillovers at connector nodes. Summing both components yields the total Big Push effect. Unlike laissez-faire, the Big Push requires public outlays ($185M gross, approximately $60M net after charging revenues) but delivers substantially larger consumer and climate benefits.

Table 8: Counterfactual Results: Big Push Total Effect Decomposition

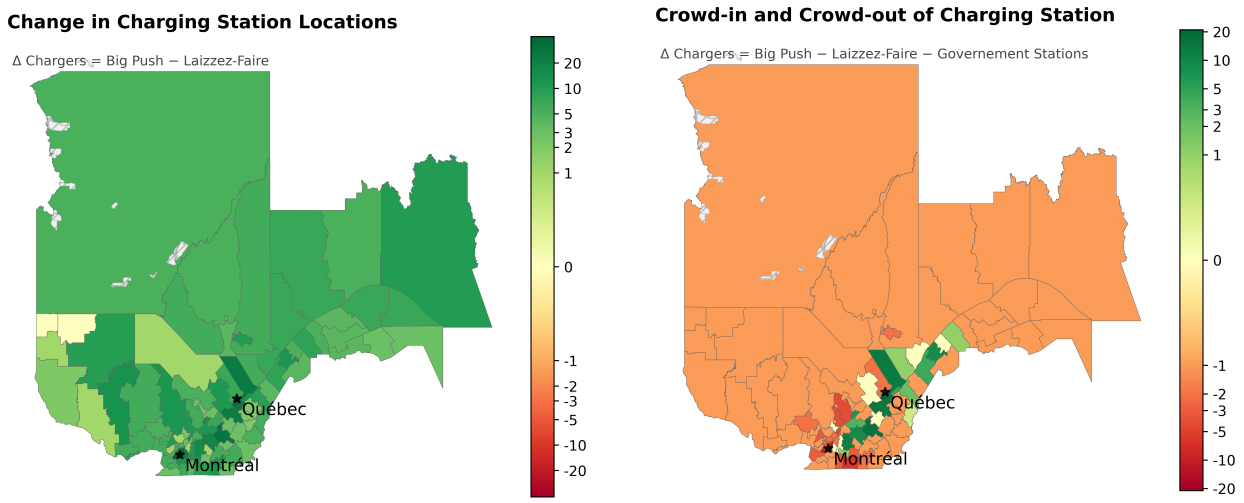
	(1) **Coordination** *High Vs. Low*	(2) **Externality** *Big Push Vs. High*	(3) **Total Effect** *Big Push Vs. Low*
Policy Outcome			
Δ EV Adoption	139%	52%	264%
Δ Network Quality	1.06	0.61	1.68
WTP (in USD)	*($10,277)*	*($5,913)*	*($16,191)*
Δ Stations	329	755	1084
Crowd-in	-	14	14
Benefits *(in $M USD)*			
Consumer Benefits	7,167	3,991	11,158
Private Profit	15	1.5	16.5
Climate Benefits			
SCC @ 185	1,429	1,290	2,715
SCC @ 50	386	349	734
Investment Cost	-	185	185

Note: The total Big Push is divided into two components: coordination and externality. Under Laissez-Faire, two stable equilibria emerge: a *(High, High)* state (both adopt) and a *(Low, Low)* state. Column (1) shows the gain from moving from the degenerate Low equilibrium to the High equilibrium (the coordination component). Column (2) isolates the externality component, capturing the additional gains from moving from the High equilibrium to the Big Push. Beyond simply shifting the equilibrium toward the High state, the Big Push also addresses network externalities. Column (3) reports the total effect of moving from the Low equilibrium directly to the Big Push scenario. Per-vehicle willingness-to-pay (WTP) gains from improved network quality are about $10,300 under the coordination component (High vs. Low) and $5,900 under the externality component (Big Push vs. High). This decomposition highlights two distinct channels: the coordination effect, which ensures convergence to the high-adoption equilibrium, and the externality effect, which captures the remaining uninternalized spillovers. Together, these channels imply a total per-vehicle gain of about $16,200 when moving from the Low equilibrium directly to the Big Push.

While Table 8 quantifies the overall effectiveness of the Big Push, it does not explain why the policy performs well. The key mechanism is that the government places stations at spatially important connector nodes that are socially valuable but privately unprofitable and therefore would not be built in equilibrium. Figure 7 isolates this mechanism by com-

paring the Big Push to the laissez-faire High equilibrium, so that the difference reflects the externality component.

The left panel (Figure 7a) maps the change in the total number of chargers. It shows a broad expansion in coverage: the Big Push adds stations at locations that are privately unprofitable but socially valuable because they increase network connectivity. The right panel shows how the private sector responds to both (i) public investment and (ii) the induced increase in EV adoption. Green indicates crowd-in (more private entry) and red indicates crowd-out (less private entry). The pattern is stark: private entry is crowded in at the network core, where higher utilization and improved connectivity raise profitability, but crowded out at the periphery, where public provision reduces the returns to new private stations. Overall, the figure highlights a spatially differentiated response: public investment complements private entry at central nodes while displacing it in fringe areas.

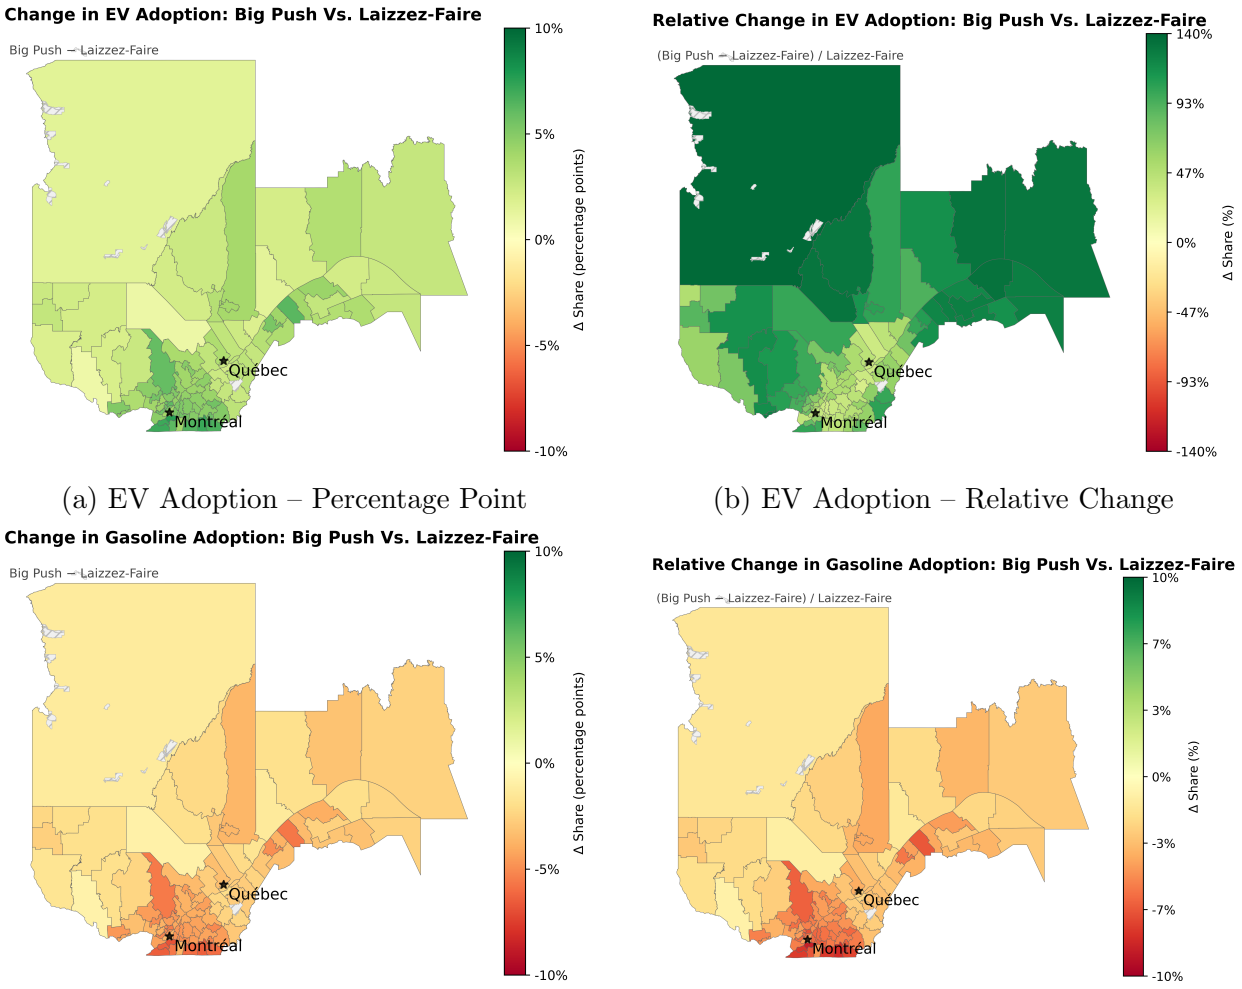


(a) Change in Number of Charging Stations

(b) Crowd-In vs. Crowd-Out Effects

Figure 7: Changes in the Charging Network Structure: Big Push vs. Laissez-Faire

This network has a substantial effect on EV adoption and also alters its spatial distribution. While Table 8 reports aggregate changes in the EV fleet, Figure 8 illustrates spatial heterogeneity in EV and gasoline-vehicle adoption. The figure reports both percentage-point and relative changes. EV adoption increases across the entire map, with the largest percentage-point gains concentrated in the network core (the Greater Montréal area). By contrast, relative gains are strongest in the periphery, where little private investment would otherwise occur. Thus, by extending the network to its outer edges, the Big Push increases adoption in the urban core but, in relative terms, has its largest impact on rural communities.



(a) EV Adoption – Percentage Point

(b) EV Adoption – Relative Change

(c) Gasoline Adoption – Percentage Point

(d) Gasoline Adoption – Relative Change

Figure 8: Changes in Vehicle Adoption under Big Push vs. Laissez-Faire

These findings corroborate Rosenstein-Rodan (1943)'s insight that Big Push policies can overcome coordination failures. Consistent with this mechanism, I show that Québec's realized Big Push (i) resolves the coordination failure by shifting the market from a Low- to a High-adoption equilibrium, (ii) crowds in private firms, with a spatial pattern that underscores complementarity between public provision and private entry (crowd-in at core connector nodes and partial crowd-out at the periphery), and (iii) induces additional EV adoption at an implied cost of \$12-\$26 USD per ton of carbon abated (the lower bound reflects the total Big Push effect, while the upper bound isolates the externality channel).

More broadly, my empirical contributions are threefold. First, reduced-form evidence shows that continuous corridor coverage by fast-charging stations increases EV adoption. Second, in the structural model I quantify this mechanism by allowing route-level fast-charging access to raise utility even without realized charging, and I derive an origin-

destination route-connectivity measure based on expected, rather than realized, trip patterns; this implies a substantial willingness to pay for network connectivity. Third, I provide direct evidence that spatially targeted public provision can strategically complement private investment by crowding in private entry while extending coverage to socially valuable but privately unprofitable nodes.

7 Conclusion

Network industries can exhibit multiple equilibria when adoption and investment are strategic complements. In the context of EV charging, this paper shows that a spatially targeted "Big Push" can move the market out of a low-adoption trap and toward a high-adoption equilibrium. The central lesson is not merely that more infrastructure raises adoption, but that where it is deployed matters: placing stations at high-connectivity nodes materially improves the value of the network.

The results highlight two distinct roles for public policy. First, public provision can serve as an equilibrium-selection device, coordinating consumers and firms on a high-adoption outcome. Second, even conditional on high adoption, government investment can raise welfare by internalizing network spillover by deploying stations at socially valuable connector locations that private firms would not build because they are privately unprofitable. Importantly, this intervention need not crowd out private investment. Instead, public provision can be complementary to private entry: by strengthening connectivity at the network core, it raises utilization and profitability where firms expand (crowd-in), while partially displacing entry at the fringe where returns are lower.

For policy design, the implication is that industrial policy in network markets should be spatially and topologically targeted. Jurisdictions with public ownership can implement direct provision, while decentralized settings can use geographically differentiated subsidies or temporary exclusivity rights to replicate the same incentives. Overall, the findings provide empirical support for Rosenstein-Rodan (1943)'s insight that Big Push policies can overcome coordination failures and unlock high-adoption equilibria in network industries.

References

Arrow, K. J. and A. C. Fisher (2013). Environmental preservation, uncertainty, and irreversibility. In *Economics of Natural & Environmental Resources (Routledge Revivals)*, pp. 479–488. Routledge.

Berry, S. T., J. A. Levinsohn, and A. Pakes (1993). Automobile prices in market equilibrium: Part i and ii.

Bresnahan, T. F. and P. C. Reiss (1991). Entry and competition in concentrated markets. *Journal of political economy 99*(5), 977–1009.

Busse, M., J. Silva-Risso, and F. Zettelmeyer (2006). $1,000 cash back: The pass-through of automanufacturer promotions. American Economic Review 96(4), 1253 – –1270.

Callaway, B. and P. H. Sant'Anna (2021). Difference-in-differences with multiple time periods. *Journal of econometrics 225*(2), 200–230.

Choi, J. and Y. Shim (2024). Industrialization and the big push: Theory and evidence from south korea. Working Paper, IMF.

Dixit, A. K. and R. S. Pindyck (1994). *Investment under uncertainty*. Princeton university press.

Fang, H., M. Li, L. Wang, and Z. Yang (2025). High-speed rail and china's electric vehicle adoption miracle. Technical report, National Bureau of Economic Research.

Fournel, J.-F. (2023). Electric vehicle subsidies: Cost-effectiveness and emission reductions.

Garg, T. (2025). Can industrial policy overcome coordination failures? theory and evidence. Working Paper, MIT.

Li, J. (2019). Compatibility and investment in the us electric vehicle market. *Unpublished manuscript, MIT*.

Muehlegger, E. and D. S. Rapson (2022). Subsidizing low-and middle-income adoption of electric vehicles: Quasi-experimental evidence from california. *Journal of Public Economics 216*, 104752.

Pavan, G. (2017). Green car adoption and the supply of alternative fuels. *Unpublished manuscript*.

Remmy, K. (2025). Adjustable product attributes, indirect network effects, and subsidy design: The case of electric vehicles. Working Paper.

Rosenstein-Rodan, P. N. (1943). Problems of industrialisation of eastern and south-eastern europe. *The economic journal 53*(210-211), 202–211.

Sallee, J. M. (2011). The surprising incidence of tax credits for the toyota prius. *American Economic Journal: Economic Policy 3*(2), 189–219.

Springel, K. (2021). Network externality and subsidy structure in two-sided markets: Evidence from electric vehicle incentives. *American Economic Review 13*(4), 393–432.

Springel, K. and J. Li (2025). Entry and coordination in the u.s. electric vehicle charging industry. *Working Paper*.

Sturm, J. (2025). How to fix a coordination failure: A super-pigouvian approach. Working Paper, MIT.

Weisbrod, B. A. (1964). Collective-consumption services of individual-consumption goods. *The Quarterly Journal of Economics 78*(3), 471–477.

A.1 Destination Inference

I infer each trip's destination by combining (i) the charging station where the event occurs, (ii) the driver's home origin, and (iii) the battery state of charge (SOC) observed at the start of charging. The key idea is that, for trips involving a single charging stop, the SOC at the start of the *returning* charge tightly pins down the total distance driven since departure, which in turn restricts the set of plausible destinations.

Step 1: Restrict attention to single-charge trips

I focus on trips that entail a single public fast charge within a 24-hour window. Operationally, I drop cases where the same `ConsumerID` charges twice (or more) within 24 hours. The identifying logic is that a well-functioning one-stop charging network is a sufficient statistic for a more complex n-stop charging network. Additional intermediate charging opportunities act as a positive transformation of the feasible set: they weakly expand reachability and reduce effective travel costs, but do not overturn the ranking of destinations or the structure of optimal routing.

Step 2: Classify outbound versus returning charges

Even among single-charge days, a charge can occur either on the way to the destination (outbound) or on the way back home (returning). I classify the charge using a reachability logic based on SOC and distances:

- If, at the start of charging, the driver has enough remaining range to reach home, and the station is close to home, this charge is likely *outbound*: the driver tops up early rather than charging only when necessary.
- If, at the start of charging, the driver has too little remaining range to reach home, the vehicle must have already reached its destination. This pattern is most consistent with the *returning* leg: the driver has reached the destination earlier and is now charging to complete the trip back home.

This classification matters because the returning-trip charge contains the most direct information on total distance driven since the last high-SOC state (conceptually, leaving home at or near full battery).

Step 3: Translate start-of-charge SOC into used distance

Assume the driver leaves home with a full battery (normalized to 1). Let km_j denote the vehicle's usable range (in km), implied by its battery capacity. If b_n^{start} is the SOC at the beginning of charge n, then the implied distance driven since departure is

$$\widehat{L}_n = \text{km}_j \left(1 - b_n^{\text{start}}\right). \tag{23}$$

Step 4: Residual distance and feasible destination set

Let o be the origin (home), s the observed charging station, and $\text{dist}(\cdot, \cdot)$ denote the distance metric used in the paper. The origin-to-station distance $\text{dist}(o, s)$ is observed. For a returning-trip charge under a single-stop structure, the remaining distance not accounted for by the home-to-station segment corresponds to the destination component of the round trip. I define the implied destination radius as

$$r_n = \frac{\widehat{L}_n - \text{dist}(o, s)}{2}. \tag{24}$$

Intuitively, $\widehat{L}_n$ measures total distance since leaving home; subtracting $\text{dist}(o, s)$ leaves the part attributable to reaching the destination and then returning toward home; dividing by two maps that round-trip component into a one-way distance scale.

This yields a feasible destination set $\mathcal{D}_n$ consisting of destinations that are (i) consistent with the remaining distance implied by SOC and (ii) farther from the origin than the charging station:

$$\mathcal{D}_n = \{d : \ \text{dist}(d, s) \leq r_n, \ \text{dist}(o, d) \geq \text{dist}(o, s)\}. \tag{25}$$

Step 5: Point inference by best matching the implied distance

Within $\mathcal{D}_n$, I assign the inferred destination as the location whose origin-to-destination distance best matches the implied one-way distance:

$$\hat{d}_n = \arg\min_{d \in \mathcal{D}_n} |r_n - \text{dist}(s, d)|. \tag{26}$$

Overall, the approach uses SOC as a travel meter. The single-charge restriction limits itinerary complexity; the outbound/returning classification determines when SOC has a clear structural interpretation; and the residual-distance mapping converts SOC into a disciplined feasible set and a transparent destination assignment rule.

Validation of the trip classification

A subset of fast-charging events can be independently linked to subsequent local on-street charges using `ConsumerID`. These cases provide direct evidence on the battery level at which drivers arrive near their final destination, a quantity that is otherwise unobserved. In this subsample, fast charging is followed by a short-range local charge away from the home origin, revealing the realized, and thus desired, remaining battery upon arrival.

Figure A.1 compares the distribution of remaining battery levels at the fast charger implied by this fast-to-local charging sequence (left panel) with the distribution recovered from the clustering-based classification of returning trips (right panel). While the two distributions are not identical, they are very similar in shape and closely aligned at the median, which lies around 30%. This correspondence provides validation that the clustering procedure correctly identifies returning trip. Importantly, this validation relies on a minimal subset of the data and is not used in estimation.

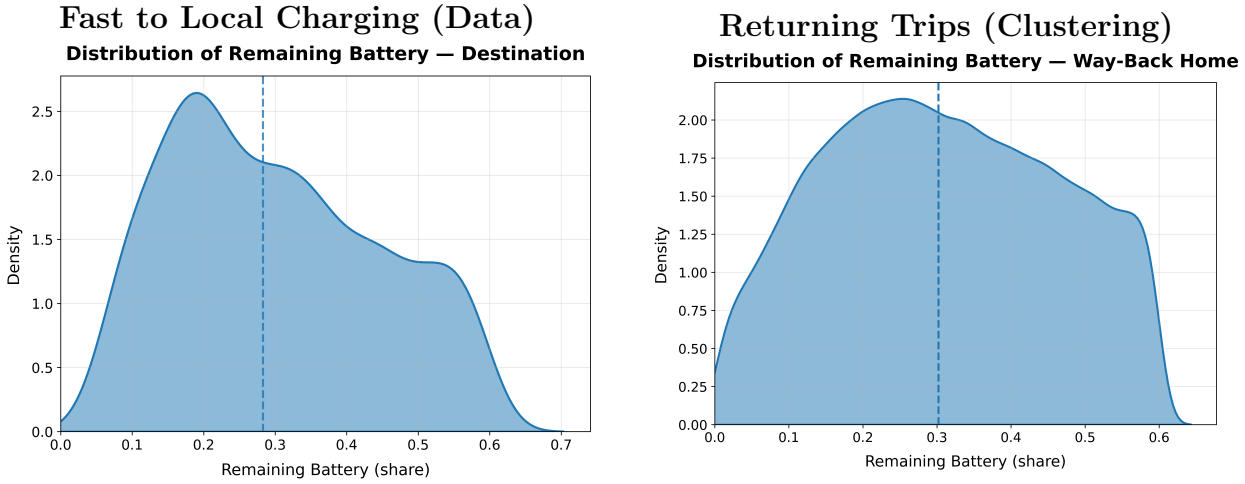


Figure A.1: Validation of returning-trip classification using fast-to-local charging sequences

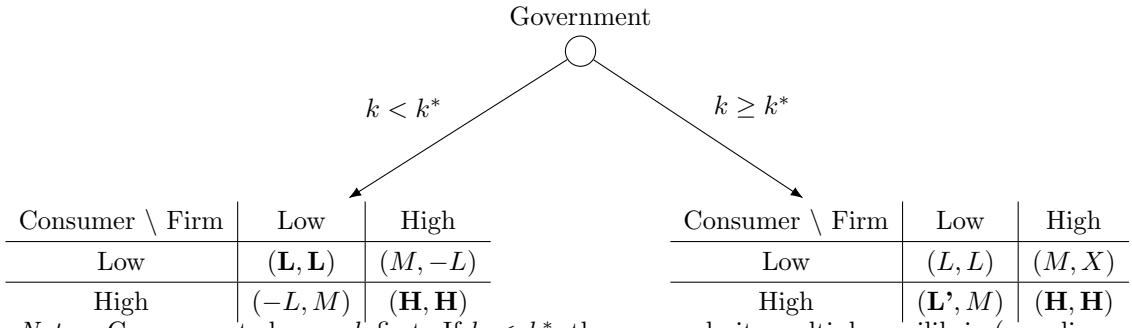
A.2 Two Stable Nash-Equilibrium

Toy Model. A key feature of the paper is the existence of two stable Nash equilibria under laissez-faire. A central aspect of the Big Push is to ensure that firms coordinate on the high-adoption equilibrium. This subsection introduces a simple toy model to illustrate the intuition that minimal public intervention can shift the economy toward the high-adoption path. The objective is not public provision per se, but rather to nudge the equilibrium away from the low-adoption trap. I also visually demonstrate the stability of these two point down the line.

In essence, the interaction between EV consumers and fast-charging firms resembles a coordination game (akin to a prisoner's dilemma). Consumers may choose High or Low adoption, and firms may also choose High or Low investment. Without Loss of generality, let's assume the payoff to be ordered as follows $0 < L < M < H$. Payoffs are given by the following matrix, where the row player is the Consumer and the column player is the Firm:

Consumer \ Firm	Low	High
Low	$(\mathbf{L}, \mathbf{L})$	$(M, -L)$
High	$(-L, M)$	$(\mathbf{H}, \mathbf{H})$

In bold are the respective best responses. From this, it is straightforward that there are two Nash equilibria. However, a key role of government provision can be to coordinate the outcome toward the high-adoption path – not by directly addressing the network externality, which is the other focus of this paper, but by eliminating the consumer's low-adoption option from the best-response set. To illustrate, consider the scenario in which the government moves first.



Consumer \ Firm	Low	High
Low	$(\mathbf{L}, \mathbf{L})$	$(M, -L)$
High	$(-L, M)$	$(\mathbf{H}, \mathbf{H})$

Consumer \ Firm	Low	High
Low	(L, L)	(M, X)
High	$(\mathbf{L'}, M)$	$(\mathbf{H}, \mathbf{H})$

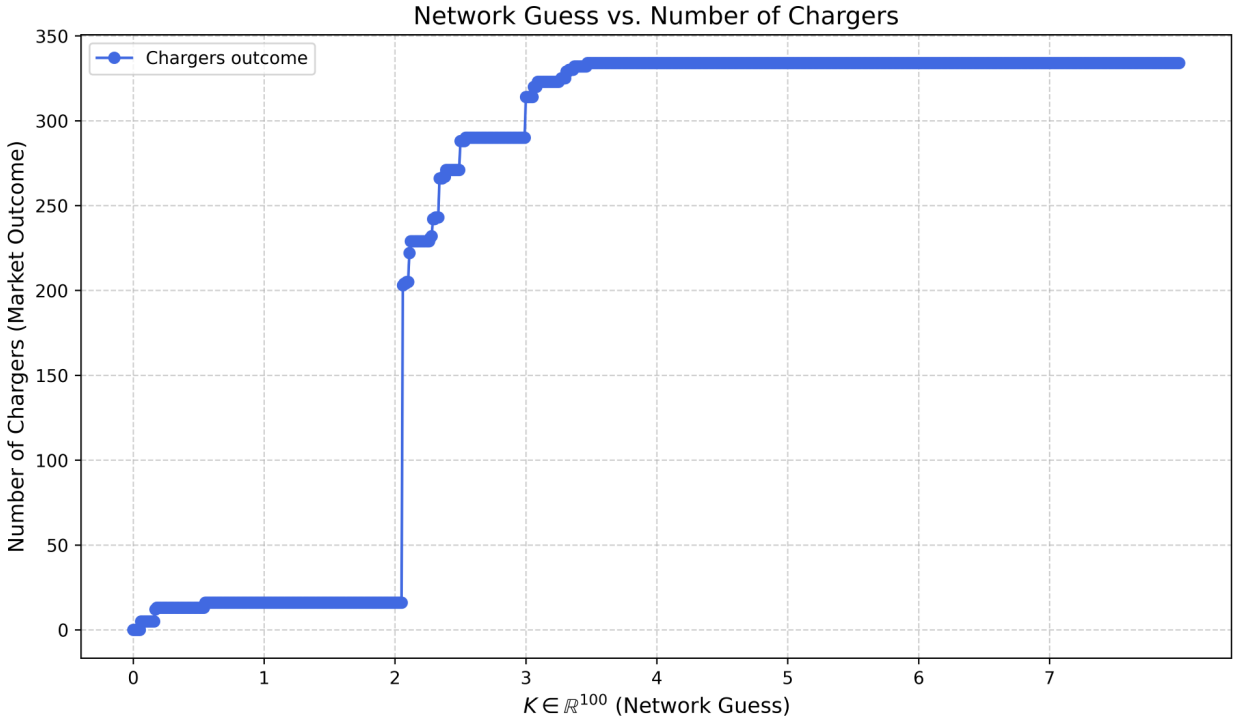
Notes: Government chooses k first. If $k < k^*$, the game admits multiple equilibria (coordination problem). If $k \geq k^*$, consumer adoption $High(k)$ becomes dominant, making (High, High the sole equilibrium.

Graphical Demonstration. While the previous discussion establishes the possibility of multiple Nash equilibria at a conceptual level, it remains to be shown that such multiplicity arises in the quantitative model. To demonstrate this, I exploit the path dependence of the fixed-point algorithm described in Section 4.4. Specifically, I show that the model converges to two distinct stable equilibria depending on the initial network configuration. In plain English, this means that the market can settle into two different equilibria depending on consumers' and firms' initial beliefs about the charging network.

I proceed by initializing the algorithm at a sequence of increasingly dense networks. Starting from a minimal configuration, $\mathcal{K}_1^{(1)} = (1, 0, 0, \ldots, 0)$, I then consider progressively denser initial conditions, $\mathcal{K}_1^{(2)} = (1, 1, 0, \ldots, 0)$, continuing in this fashion until reaching a uniform initial network $\mathcal{K}_2^{(0)} = (2, 2, 2, \ldots, 2)$. For each initial condition, I iterate on

the equilibrium mapping until convergence using the procedure described in Section 4.4. I increase the initial density further until additional chargers no longer affect the limiting outcome.

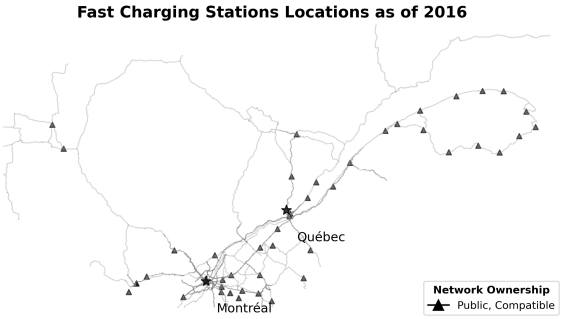
This exercise reveals two basins of attraction: initial conditions below a threshold converge to a low-adoption equilibrium characterized by sparse private entry and limited EV adoption, while sufficiently dense initial networks converge to a high-adoption equilibrium with extensive private entry and widespread EV adoption. The existence of these distinct convergence paths provides constructive evidence of multiple stable Nash equilibria in the model and illustrates how small public interventions can shift the economy from the low-adoption basin to the high-adoption equilibrium.



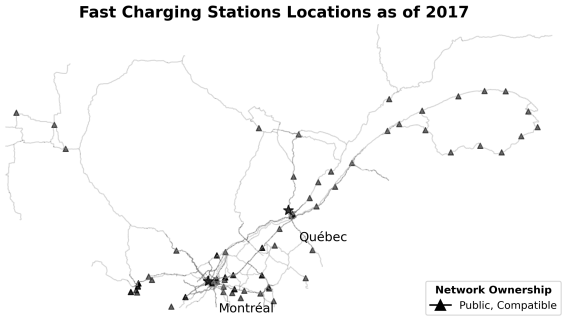
In this figure, the horizontal axis represents the initial network belief, where $\mathcal{K} = 1$ denotes a vector of ones of dimension 100. The vertical axis reports the resulting equilibrium outcome, measured by the total number of charging stations. Despite increasingly optimistic initial beliefs about network density, the equilibrium outcome does not rise beyond a certain point: additional stations become privately unprofitable, highlighting the role of network externalities. At the same time, when initial beliefs are sufficiently low, the market converges to a low-adoption equilibrium with very limited charging entry. These patterns illustrate the existence of two stable steady-state equilibria, with convergence depending on initial beliefs.

A.3 Network Evolution

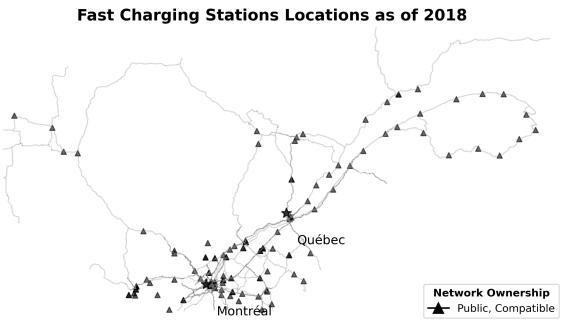
Figure A.2: Network Evolution



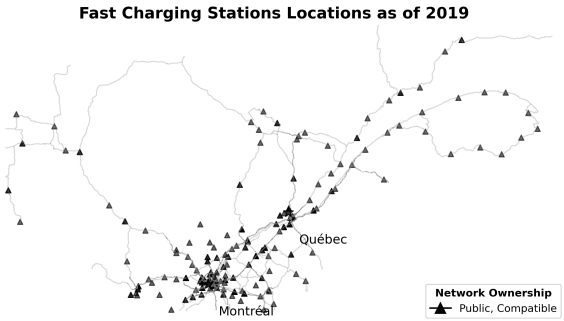
(2016)



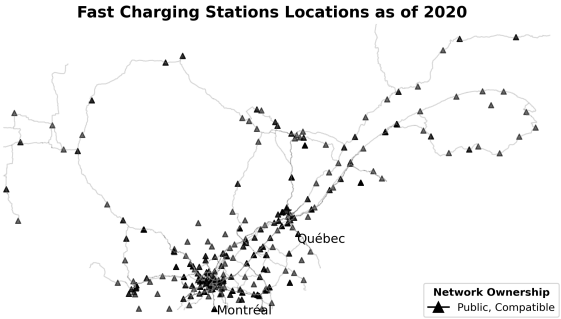
(2017)



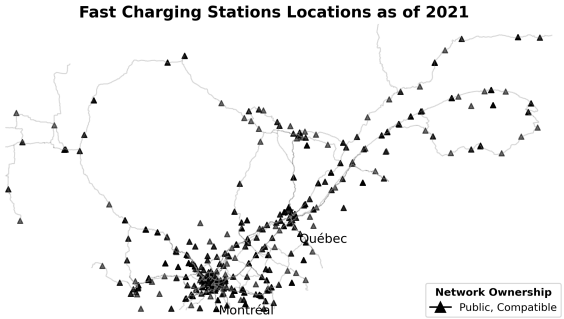
(2018)



(2019)



(2020)



(2021)

A.4 Alternative counterfactual

Table A.1: Counterfactual Results: Relative to Laissez-Faire

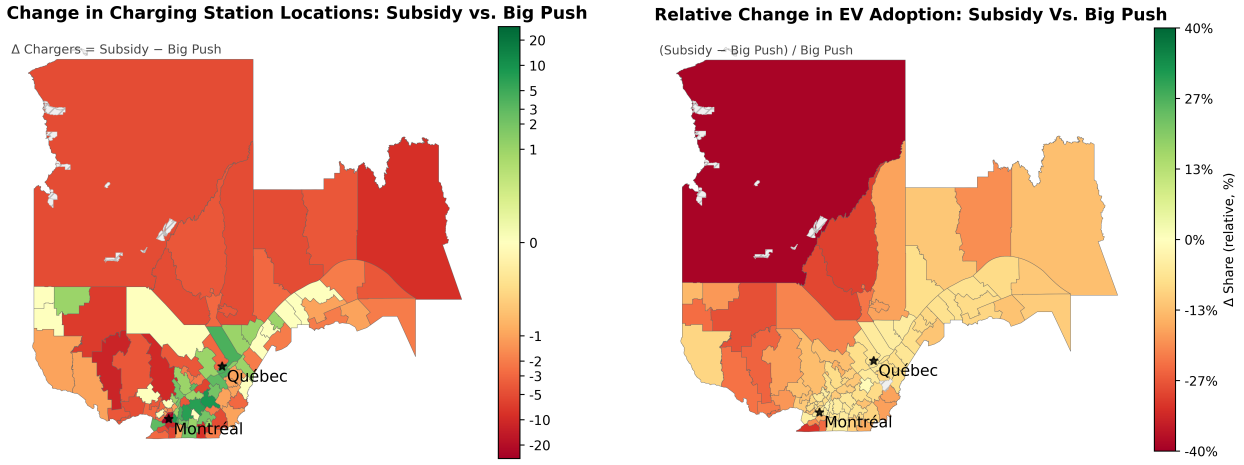
	Big Push	**Subsidy**	**Monopoly**
Policy Outcome			
Δ EV Adoption	52%	39%	47%
Δ Network Quality			
Rawlsian	34 %	26 %	23 %
Utilitarian	32 %	25 %	27 %
Δ Stations	755	586	183
Crowd-in	14	-	-
Benefits *(in $M USD)*			
Consumer Benefits	6,072	4,190	3,970
Private Profit	2	51	128
Climate Benefits			
SCC @ 185	1,290	970	1,156
SCC @ 50	350	260	312
Policy Cost	60.5	57.5	0

Note: All values are reported relative to the Laissez-Faire baseline. The per-vehicle gains from improved network quality are approximately $9,000 under the Big Push, $6,700 under the subsidy, and $6,000 under the monopoly. The average profit per private charging station is about $115,000 under the Big Push, $87,000 under the subsidy, and $700,000 under the monopoly. Each EV is assumed to save 30 tons of carbon over its lifetime.

The subsidy provides a financial incentive of similar magnitude to the Big Push. It improves network quality by about 25% and generates a 40% increase in EV adoption. Each EV gains about $6,700 in additional network value relative to the Laissez-Faire, translating into $4.2 billion in consumer benefits and roughly $1 billion in climate benefits. Unlike the Big Push, the subsidy strongly stimulates private entry: it crowds in 586 new stations with an average profit of $87,000 each, yielding $51 million in total private profits. This makes it the most favorable policy for private investors. However, the policy cost remains high, at $57.5 million, since the government shoulders the subsidy on each new installation.

The monopoly scenario also offers a financial incentive of about the same order as the

Big Push. It improves network quality by roughly 25% and increases EV adoption by 45%. Each EV benefits from about $6,000 in network value, producing $4 billion in consumer gains and just over $1 billion in climate benefits. Unlike the subsidy, however, the monopoly yields fewer new stations—183 in total—but at higher average profitability, around $700,000 per station, generating $128 million in private profits. Because the monopolist internalizes all network externalities, it deploys strategically to maximize overall network profitability, balancing expansion with the risk of cannibalization. Importantly, the monopoly entails no direct fiscal cost for the government, making it fiscally attractive despite somewhat lower consumer gains compared to the Big Push.



(a) Change in Network Structure

(b) Spatial Change in EV Adoption

Figure A.3: Subsidy vs. Big Push: Network Structure and EV Adoption

To better understand the role of each policy, it is useful to compare Rawlsian and utilitarian measures of network quality. The Rawlsian metric assigns equal weight to each region, while the utilitarian metric weights regions by population. Unsurprisingly, the monopoly performs best on the utilitarian margin, as the monopolist maximizes network value by concentrating chargers where demand is highest. By contrast, public provision delivers the highest Rawlsian quality, reflecting its mandate to ensure broad geographic coverage; it is, by far, the most effective policy for connecting peripheral regions. Figure A.3 illustrates the divergence between private and public incentives by comparing the spatial allocation under the subsidy with the Big Push outcome. Despite deploying fewer total stations, the subsidy channels investment disproportionately toward the network core, leaving peripheral regions with fewer chargers than under public provision. In the right panel of Figure A.3b, we observe that the displacement of stations under the subsidy has a disproportionate effect on the network periphery, particularly in less profitable areas. Because these regions are often rural and lower-income, the allocation of chargers raises redistributive concerns: public provision

expands access more evenly, while subsidy-driven entry tends to reinforce existing spatial inequalities.

Finally, an interesting result is the effectiveness of the monopoly. By deploying only a limited number of stations—about 180—it nevertheless raises overall network quality by 25%. This outcome arises because the monopolist invests exclusively in the highest-value sites to maximize charging throughput. As shown in Figure A.4, most of the new stations are concentrated along the corridor between Québec City and Montréal and in high-demand touristic areas such as Gaspésie, Côte-Nord (north of Québec City), and Mont-Tremblant (west of Montréal). The network does not extend further into the periphery beyond what Laissez-Faire alone would produce. Thus, while highly effective at fostering adoption in major urban and tourist regions, the monopoly highlights the distributional consequences of such a policy: it disproportionately benefits the two largest cities, leaving rural communities underserved. By contrast, the Big Push extends coverage more evenly across space and, as a result, delivers larger welfare gains overall.

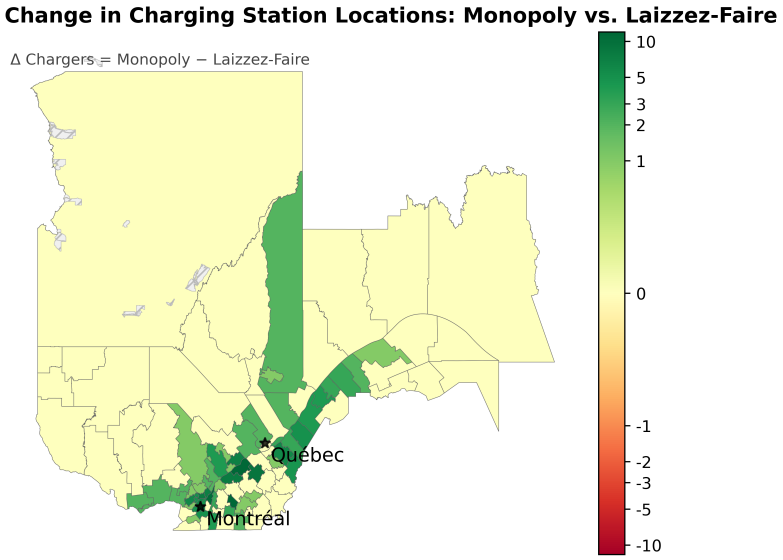


Figure A.4: Change in Network Structure: Monopoly vs. Laissez-Faire

A.5 Destination Fixed Effects

Table A.2: Destination fixed effects (full set)

Rank	Destination	$\hat{\zeta}_d$
1	Québec	7.339
2	Sainte-Agathe-des-Monts	6.092
3	Trois-Rivières	5.778
4	Longueuil	5.742
5	Gatineau	5.738
6	Lévis	5.389
7	Rivière-du-Loup	5.336
8	La Malbaie	5.162
9	Sherbrooke	5.115
10	Saguenay	5.060
11	Magog	5.009
12	Sainte-Adèle	4.983
13	Drummondville	4.832
14	Rawdon	4.759
15	Stoneham-et-Tewkesbury	4.736
16	Rimouski	4.704
17	Saint-André-Avellin	4.595
18	Cowansville	4.563
19	Boischatel	4.556
20	Chandler	4.518
21	Sorel-Tracy	4.492
22	Val-d'Or	4.476
23	Gaspé	4.447
24	Baie-Saint-Paul	4.437
25	Saint-Félicien	4.365
26	Blainville	4.353
27	Baie-Comeau	4.337
28	Saint-Hyacinthe	4.331
29	Mont-Laurier	4.329
30	Louiseville	4.247
31	Sainte-Anne-des-Monts	4.181

Continued on next page

Table A.2: Destination fixed effects (full set)

Rank	Destination	$\hat{\zeta}_d$
32	Saint-Jérôme	4.157
33	Joliette	4.143
34	Chambly	4.008
35	Matane	4.003
36	Vaudreuil-Dorion	3.989
37	Sept-Îles	3.961
38	Carleton-sur-Mer	3.866
39	Victoriaville	3.853
40	Alma	3.822
41	Forestville	3.818
42	Saint-Jean-sur-Richelieu	3.808
43	Rouyn-Noranda	3.796
44	Shawinigan	3.782
45	Granby	3.778
46	Saint-Georges	3.762
47	Châteauguay	3.719
48	La Tuque	3.686
49	Maniwaki	3.672
50	Amqui	3.647
51	Thetford Mines	3.519
52	Montmagny	3.518
53	Saint-Raymond	3.509
54	Terrebonne	3.488
55	Amos	3.475
56	Lac-Mégantic	3.439
57	Longue-Pointe-de-Mingan	3.398
58	Mont-Joli	3.343
59	Salaberry-de-Valleyfield	3.299
60	Dolbeau-Mistassini	3.251
61	La Pocatière	3.229
62	Lachute	3.206

Continued on next page

Table A.2: Destination fixed effects (full set)

Rank	Destination	$\hat{\zeta}_d$
63	Saint-Eustache	3.129
64	Plessisville	3.117
65	Bécancour	3.050
66	Ville-Marie	3.022
67	Témiscouata-sur-le-Lac	2.970
68	Lavaltrie	2.870
69	L'Islet	2.786
70	New Richmond	2.674
71	Saint-Honoré	2.665
72	Beauceville	2.639
73	La Sarre	2.595
74	Saint-Apollinaire	2.553
75	Saint-Lin–Laurentides	2.513
76	Sainte-Marie	2.458
77	Saint-Rémi	2.362
78	Trois-Pistoles	2.305
79	Nicolet	2.192
80	Val-des-Sources	2.012
81	Val-des-Monts	1.896
82	Saint-Tite	1.848
83	Mansfield-et-Pontefract	1.839
84	Lac-Etchemin	1.831
85	Repentigny	1.814
86	Cookshire-Eaton	1.797
87	Coaticook	1.783
88	Saint-Pierre-de-l'Île-d'Orléans	1.710
89	Mirabel	1.630
90	Saint-Henri	1.367
91	Windsor	1.353
92	Marieville	0.840
93	Ormstown	0.693

Continued on next page

Table A.2: Destination fixed effects (full set)

Rank	Destination	$\hat{\zeta}_d$
94	Notre-Dame-du-Mont-Carmel	0.667
95	Montréal	0.000
96	Sainte-Julie	0.000
97	Laval	0.000
98	Mistissini	0.000
99	Chibougamau	0.000

A.6 Other

Table A.3: Summary of Quebec and Federal EV Purchase Incentives

Program	**Years**	**MSRP cap**	**Rebate structure**
Quebec	2017–2019	Full $<$\$75k; partial \$75–125k; $>$ \$125k excl.	BEV: \$8k if MSRP $<$\$75k; \$3k if \$75–125k. PHEV: \$8k ($\geq$15 kWh) or \$4k (7–$<$15 kWh); ineligible if MSRP $\geq$\$75k.
Quebec	2020–2021	\$60k	BEV: \$8k. PHEV: \$8k ($\geq$15 kWh); \$4k (7–$<$15 kWh); \$500 (4–$<$7 kWh).
Quebec	2022–present	\$60k	BEV: \$7k. PHEV: \$5k ($\geq$15 kWh); \$2.5k (7–$<$15 kWh); \$300 (4–$<$7 kWh).
Federal iZEV	2019–2021	\$55k	BEV: \$5k. PHEV: \$5k ($\geq$15 kWh); \$2.5k ($<$15 kWh).
Federal iZEV	2022–present	Cars: \$65k; SUVs/pickups/vans: \$70k	BEV: \$5k. PHEV: \$5k (rated electric range $\geq$50 km); \$2.5k otherwise.

Notes: Caps and amounts are in nominal dollars. "Excl." indicates excluded (ineligible). Ranges use half-open intervals where indicated (e.g., 7–$<$15 kWh).