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

Hugo Cordeau†

July 28, 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 government provision of network infrastructure can remedy the resulting inefficiency. Using vehicle registration and charging transaction data from Québec, I estimate a structural model of EV demand in which consumers value the charging network through the origin–destination trips it supports. I decompose the effect of a large government network deployment—a Big Push—into a coordination component, shifting the market from the low- to the high-adoption equilibrium, and an externality component, capturing stations that remain privately unprofitable even under full coordination. The Big Push raises EV adoption by 67 %, with coordination contributing 31% and externalities a further 27%. It expands the network from 249 to 921 chargers, crowding out private investment in the core while crowding it in at the periphery, and generates $484 million in consumer value and $88 million in climate benefits for $22 million of government investment. This paper provides empirical evidence that direct government provision of network infrastructure can resolve coordination failures.

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

Industries with strong complementarities and network externalities may exhibit multiple equilibria. Rosenstein-Rodan (1943) introduced the notion of a "Big Push," arguing that large-scale, state-led investment is needed to escape such coordination failures. This paper studies the effectiveness of one such intervention in the electric vehicle (EV) market. Beyond the policy's effectiveness on EV adoption, a key concern with such governmental intervention is that it affects private investment. I therefore ask three questions. First, how does government network provision affect EV adoption and private station entry? Second, how do EV purchase subsidies interact with direct deployment? Namely, does it complement the network deployment, by widening the coordination gap, or does it substitute for it, by resolving the coordination failure? Third, how does the presence of a vertically integrated (VI) firm like Tesla affect the returns to government deployment, and how does government deployment impact the VI investments?

To shed light on these questions, this paper is the first to study a Big Push Policy 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 upon which consumers perform trips once they purchased an EV. Second, a firm's investment in charging infrastructure generates spatial spillovers that benefit other market participants without capturing the full extent of these benefits.

To address these challenges, my model must capture both consumers' responses to the existing charging network and the strategic response of private firms to the network structure. 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 where to go and which EV model to purchase or instead choose a gasoline vehicle.

I apply the model to the EV market, where charging infrastructure creates a classic chicken-and-egg problem: consumers adopt EVs only if a reliable fast-charging network is already in place, and private providers invest only if adoption is high enough to cover entry costs. This complementarity generates two distinct sources of underprovision. First, because best responses are weakly increasing in the aggregate, the equilibrium set is bounded by a least and a greatest element, hereafter the Low and the High equilibrium, and the market may coordinate on the Low Vives (2005). Second, even at the High equilibrium, providers internalize only their own returns: they ignore both the network-formation externality, since

a new station raises the return to other stations, and the cross-market externality, since charging investment raises vehicle demand and adoption raises charging revenue in turn. The Big Push acts on both margins, and I decompose its total effect accordingly: the move from Low to High measures the coordination component, and the move from High to the Big Push equilibrium measures the externality component.

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 government 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 show that spatial placement is crucial and that consumers greatly value stations in (otherwise) charging deserts.

I utilize five data sources in my analysis. First, I use data from the Electric Circuit network, a transaction-level dataset covering all government 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 along with the state of the battery at the start and end of charging, I construct origin–destination trips. Second, I complement the government network 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 government 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.[1]

I first provide causal evidence on EV adoption in response to charging network deployment by exploiting the government 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 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 spillovers. 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 of three sets of agents: (i) consumers, who choose whether to adopt an EV and which one; (ii) fringe charging providers, who choose whether and where to invest under free entry; and (iii) the VI firm, which chooses where to invest in its proprietary network. Consumers solve a three-stage problem by backward induction: they choose a vehicle, then a destination, then where to charge. The last two stages take place on a graph whose nodes are cities (about 100) and charging locations (about 400), with roads as edges. Their continuation value is a log-sum inclusive value, the ex ante expected utility from travel and charging before taste shocks are realized. Because it depends on range, this measure of network quality varies across origins and vehicle models.

Network quality bridges the two complementary markets, vehicles and charging. On the demand side, it enters vehicle choice, so adoption depends on the vehicle itself and on the network a given range makes accessible. On the supply side, the same travel and charging behavior maps adoption into charging demand at each location, and that demand determines where investment is profitable. Fringe charging providers earn profits from realized charging activity, while the VI firm invests to capture the additional vehicle sales that a denser

[1] 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.3.

proprietary network generates.

I now embed the EV demand response in firms' profits. The investment game is sequential and solved by backward induction. The government moves first: its network is set by policy rather than chosen strategically, so it does not respond to the private actors. The VI firm moves second. Observing the government build-out, it chooses its proprietary network to maximize profits from EV sales rather than from charging sessions, and as a Stackelberg leader it anticipates the fringe's free-entry response. The fringe moves last. Observing both the government and the VI networks, it enters wherever entry is profitable and earns profits from charging sessions, with the free-entry condition following Bresnahan and Reiss (1991). I restrict attention to two market structures: (i) a set of atomistic single-station owners together with the VI firm, and (ii) the same set of owners without the VI firm. The first corresponds to the observed market, in which the VI firm internalizes part of the network externality because its chargers raise its own vehicle sales. Comparing the two isolates the contribution of the VI firm to EV adoption and fringe investment.

To reduce the dimensionality, I adopt the local manager assumption of Aguirregabiria and Ho (2012). Rather than choosing the full network to maximize firm-wide profit, each candidate station is evaluated by a local manager who takes the current network as given and builds the station if it raises firm profit. As a Stackelberg leader the manager anticipates the fringe's free-entry response. Costs are incurred at the station while revenue accrues through vehicle sales, so the manager compares the margin on the EVs the station induces against its cost. Managers are called in order of distance to Montréal, starting with the nearest candidate node and expanding outward. The assumption replaces a choice over 2^N networks with N binary decisions taken in a fixed order.

The fringe problem is solved analogously. The entry game features complementarities through the adoption channel, since entry raises EV adoption and therefore charging demand at every node, but entry at neighbouring nodes remains locally substitutable, so the game is closer to a game of weak complementarity. I therefore compute candidate extremal equilibria by iterating best responses from the two extreme networks, in the spirit of Vives (2005). The Low equilibrium is initialized with no private station at any node, the High with a station at every node. Each iteration has two steps. First, every station that is unprofitable given the current network is removed. Second, the nearest unbuilt node is considered and is built if entry is profitable given the network that remains. The iteration proceeds outward until no profitable entry is left. The two initializations sustain different levels of EV adoption and hence different levels of charging demand, which is what allows them to converge to distinct fixed points.

In sum, EV demand and charging-station supply are jointly shaped by network exter-

nalities: 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 globally but substitutes locally. I treat Québec as a small open economy and take vehicle technology as exogenous, abstracting from the manufacturing sector as in Li (2019). 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.

The demand estimation proceeds in three stages: EV choice, 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 route. Using realized 2023 charging events, I estimate a maximum-likelihood model of charging-station choice. Identification comes from consumers' preferences over charging timing, with consumers targeting stations that leave them with approximately 27% state of charge (SOC) at the onset of charging. This step delivers an origin-destination charging quality measure that captures the expected utility of completing a trip given the network topology.[2] Second, I estimate a destination-choice model for each origin city o, where utility depends on destination-specific preferences and charging quality. I use the origin city and vehicle range to construct feasible destination sets and exploit observed destination choices to estimate realized 2023 destination preferences via maximum likelihood. These estimates generate a network-quality index that summarizes the ex-ante expected utility of completing an origin-destination trip given the charging network and vehicle model. Finally, I combine EV adoption data from 2014–2022 with the implied network-quality measures and estimate the EV demand model via GMM. The specification includes prices and product characteristics, with battery size, lithium prices, and exchange rates used as instruments.

The demand estimation reveals five main findings. First, consumers face a cost when charging away from their preferred battery level. Charging events are concentrated around a state of charge of approximately 27%, which is preferred to charging at either 50% or 10%. Second, both detours from the main route and the location of charging stations on one side of the road or the other impose costs on consumers (as in Houde, 2012). Third, charging deserts are particularly valuable: having no competing charging station within a

[2] I restrict the analysis to trips involving a single charging stop to reduce computational complexity. This restriction should not substantially affect the objective of constructing a sufficient statistic for network quality, as the single-stop procedure captures the relevant variation in network accessibility.

10 km radius nearly doubles the likelihood of a charging event. Fourth, while the origin-destination distance initially reduces the likelihood of selecting a destination, this effect disappears once destination preferences are controlled for.[3] Fifth, charging-network quality significantly affects EV adoption decisions. Specifically, a 10% improvement in charging-network quality increases EV adoption by 12%.

On the supply side, I estimate the fixed cost of charging-station entry using observed private entry decisions and model-implied profits. In the spirit of Bresnahan and Reiss (1991), I exploit the equilibrium conditions implied by entry and use observed entry timing to recover cost bounds. I then use the estimated entry cost to infer the price cost margin (PCM) per EV of the VI firm required to rationalize the observed network expansion taken, using the 2022 market outcome. The model is calibrated so that simulated EV adoption stock matches the 2022 observed EV stock.

The estimated entry cost is $\widehat{F}$ = \$30,236 per charger, consistent with engineering estimates for fast-charging deployment; a \$150,000 charging site amortized to an annuity factor of about 5, which corresponds to a payback horizon of five years. This is at the low end of the assumptions used in industry: NREL E-FAST (8 years, 10%) implies an annuity factor of 5.33, Argonne CHECT (15 years, 12%) implies 6.81, and RMI (10 years, 8%) implies 6.71. The implied price-cost margin of about \$10,000 per vehicle is consistent with the markups reported by Remmy (2025), who estimates an average markup of about 6,770 euros (approximately \$8,000) in the German automotive industry.

Using the estimated parameters and 2022 EV technology, I simulate steady-state equilibria in which consumers and firms best respond to one another. The game exhibits weak strategic complementarities and admits multiple equilibria, so I report the Low equilibrium, in which adoption and investment are both minimal, and the High equilibrium, in which consumers and firms coordinate on the largest network the market can sustain. Neither is socially optimal, because many stations are socially valuable but privately unprofitable. A third equilibrium, the Big Push, represents the government intervention. The government installs the network realized in 2022, and taking that network as given, consumers, the vertically integrated firm, and the fringe re-optimize to a new Nash equilibrium. I decompose the total Big Push effect into a coordination component, the shift from the Low to the High equilibrium, and an externality component, the further gain from the High equilibrium to the Big Push.

The causal impact of the Big Push. The deployment raises EV adoption by 67% relative

[3]Given the available data, the analysis conditions on long-distance trips that require a charging event. Therefore, the result should be interpreted as showing that, conditional on undertaking such a trip, distance does not significantly affect destination choice.

to the Low equilibrium, at a cost of $22.4 million. Coordination and externalities contribute in roughly equal measure: the total effect factors as 1.31×1.27, where 1.31 is the shift from the Low to the High equilibrium and 1.27 is the further shift from the High equilibrium to the Big Push. The two channels operate differently. Coordination leaves the total number of chargers roughly unchanged but reallocates them. The number of active stations more than doubles, and the fringe gains share at the expense of the vertically integrated firm. The externality channel instead expands the network, raising the number of chargers fourfold and crowding out part of the fringe. The impact is heterogeneous across space; the externality effect benefits cities just outside the network core the most. Welfare gains are $484 million in consumer surplus, $144 million in manufacturer surplus, and $88 million in climate benefits.

This gives quantitative support to the classic Big Push argument of Rosenstein-Rodan (1943), that coordinated investment can move an economy out of a low-level equilibrium. Québec's realized Big Push does exactly that in a network industry: it moves the market out of the low-adoption trap. The Big Push turned out to be an effective policy: it has an implied cost of $47 per ton of CO_2 abated, well below the social cost of carbon.

Subsidies as complements or substitutes. Because charging is a two-sided market, demand-side support feeds back into station deployment. I therefore repeat the exercise with the real-world EV purchase subsidy in place. The two instruments are complements in absolute terms: a fixed deployment converts roughly 4,500 more vehicles once demand is subsidized, or nearly 30% more. In proportional terms they are substitutes, because the subsidy raises the baseline enough that a larger absolute gain is a smaller relative one. This is consistent with Springel (2021) on the non-neutrality of subsidies and on the diminishing marginal value of the charging network.

I bring a novel contribution: the subsidy does not resolve the coordination failure but widens it, nearly doubling the gap between the Low and High equilibria. The cost of landing on the wrong equilibrium, or equivalently the value of a mechanism that selects the right one, rises with demand-side support rather than falling. In other words, demand side incentives increase the Big Push value.

The value of VI. Finally, I shut down Tesla's charger investment, so that the firm sells vehicles but cannot build a proprietary network. Because the integrated firm internalizes the effect of its chargers on its own vehicle sales, it invests where the fringe will not, and its network is worth most where coordination fails. Two results follow. Low equilibrium adoption falls from 23,785 to 21,330 vehicles when Tesla cannot build, a loss of 2,455 vehicles or 10.3%, against losses of 3.6% in the High equilibrium and 2.6% under the Big Push. The measured effect of the Big Push is correspondingly larger absent Tesla's network, 17,346 vehicles (a factor of 1.81) rather than 15,906 (1.67), because the state the market starts from

is worse. VI thus works mainly as insurance against the worst equilibrium; its contribution in the High and Big Push states is roughly a third as large. This finding points in the opposite direction from Li (2019), who shows that mandating compatibility raises EV sales, though the two counterfactuals remove different objects and are not directly comparable.

Related Literature

This paper contributes to three strands of literature: (i) EV adoption and charging infrastructure, (ii) industrial policy and coordination failures, and (iii) entry games with multiple equilibria.

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) shows that plug standardization increases consumer surplus and EV sales while reducing redundant infrastructure build-out by altering firms' investment incentives. I differ from Li by studying how the Big Push affects EV adoption and private investment. My model also contains a VI firm, but its role is to show how the Big Push crowds that firm in or out, and how vertical integration lets it internalize part of the network externality.

A complementary paper, Fournel (2023), examines the effectiveness of EV subsidies. I depart from Fournel by studying private firms' entry responses to improvements in network quality, asking whether government provision crowds private investment in or out. A closely related contribution is Springel (2021), who studies the non-neutrality between charging-station subsidies and consumer subsidies. Springel asks how the two instruments differ; I ask how charging firms and consumers respond to government network provision. I recover her two central findings, the non-neutrality of subsidies and the diminishing marginal value of charging stations, and I add a third margin. Demand-side support widens the coordination failure instead of resolving it: the gap between the Low and High equilibria nearly doubles once the subsidy is in place. The cost of landing on the wrong equilibrium, and equivalently the value of a mechanism that selects the right one, therefore rises with the subsidy.

Beyond the charging margin, Barwick et al. (2025) quantify EV industrial policy on the manufacturing side, estimating a dynamic model of the global battery supply chain in which learning by doing interacts with government subsidies. They study the upstream technology lever and I study the complementary-infrastructure lever, so the two papers bracket the supply-side alternatives to demand-side subsidies. Closest in policy spirit, Cole et al. (2023) simulate U.S. federal spending on EV tax credits and charger construction and

find that charging-station subsidies are highly cost-effective, with abatement costs below the social cost of carbon (SCC). Infrastructure enters their framework as an aggregate stock. I microfound where infrastructure creates value, through trip completion, endogenize the private entry response, and show that the below-SCC abatement result survives in a full spatial equilibrium with multiple equilibria.

My theoretical contribution is 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 government network provision raises EV adoption and crowds private investment in or out depending on location, at an implied cost of US$47 per ton of CO_2.

Second, I provide contemporary empirical evidence on coordination failures in industrial-policy settings by studying a measurable, realized Big Push. In industries with network externalities and strong complementarities, multiple equilibria can arise, creating the potential for coordination failures. Rosenstein-Rodan (1943) argued that large-scale, state-led investment—a "Big Push"—is needed to shift such economies toward the high-adoption equilibrium, an insight formalized by Murphy et al. (1989). A central concern with such intervention is that it may crowd out private investment; my main contribution to this literature is to show that direct government network provision can both (i) resolve the coordination failure and (ii) perform partial crowd in of private firms at strategic nodes. Empirical work documents that large-scale interventions can have persistent effects, from place-based programs (Kline and Moretti, 2014) to temporary trade protection (Juhász, 2018). Lane (2025) provides causal evidence that South Korea's heavy and chemical industry drive expanded targeted sectors, with gains that persisted after support was withdrawn and propagated downstream; because the design is reduced-form, however, the counterfactual equilibrium is unobservable, whereas my structural model recovers that object. Liu (2019) derives a sufficient statistic, distortion centrality, that identifies which sectors to target in a distorted production network; my targeting problem is the spatial analogue, identifying high-connectivity nodes of the road network that are socially valuable but privately unprofitable. Recent work seeks to isolate the coordination channel from fundamentals. Garg (2025) develops an algebraic-geometry framework to classify factual and counterfactual equilibria and, applied to Indian industrial zones, finds that zones raise the probability of escaping a low-industrialization equilibrium by 38%. I complement this approach by studying equilibrium switching in an evolving network industry: while the paper analyzes industrial zones as a legal and institutional intervention, I study the physical build-out of a compatible charging network, model the entry response of private firms, and quantify the welfare gains from network provision. Similarly, Choi and Shim (2024) show that a one-time technology subsidy

in 1970s South Korea shifted the economy to a high-industrialization equilibrium, raising heavy manufacturing output by 8.6%, and Sturm (2025) proposes a super-Pigouvian policy that preserves decentralized price regulation while preventing coordination failure. My paper differs from these works in two key ways. First, I study direct provision of quantities rather than price-based mechanisms or equilibrium-selection criteria: subsidies and taxes raise the payoff to the high-adoption equilibrium but need not select it, whereas building the network corrects the failure at its source. Second, I exploit a realized Big Push in which both the government build-out and the private entry response are observed. My theoretical contribution is to develop a multiple-equilibria model in which direct government network provision alleviates the coordination failure and shifts the market out of the low-adoption trap. My empirical contribution is twofold. I show that Québec's realized Big Push achieves precisely this, moving the market beyond the High equilibrium and raising EV adoption by 67%. I further show that demand-side subsidies in a two-sided market may not resolve the coordination failure but widen it, enlarging the gap between equilibria without selecting between them. The realized Big Push does precisely this selecting work, and its causal effect is larger in absolute terms when the subsidy is in place: a thicker demand side raises the return to the deployment even as it deepens the coordination problem the deployment resolves.

Third, I connect this setting to the literature on entry games and equilibrium multiplicity. Following Bresnahan and Reiss (1991), empirical entry models use observed entry decisions to recover the profitability conditions that rationalize market structure. Their identification relies on geographically isolated markets; mine are non-isolated by construction—entry at one node raises EV adoption and hence charging demand everywhere—so entry conditions must hold jointly in a spatial equilibrium, and I exploit the timing of entry along the network's concentric rollout, rather than a cross-section of markets, to bound fixed costs. Methodologically, I build on the theory of games with strategic complementarities: Tarski (1955) guarantees that monotone systems admit Low and High fixed points, and **?** and Vives (1990) show that in supermodular games these extremal equilibria bound all equilibrium behavior and can be computed by iterating best responses from the extremes of the strategy space. Jia (2008) brings these tools to empirical entry, exploiting the supermodularity of a chain's profit in its own store network to make an otherwise intractable combinatorial location problem solvable. My entry game differs in that complementarities operate across independent firms through the adoption channel, while nearby stations remain local substitutes, so the game is not globally supermodular. I therefore adapt the extremal logic: I compute candidate Low and High equilibria by iterating best responses from the empty and full networks and verify ex post that each limit point is a Nash equilibrium. To reduce the dimensionality, I adopt the local manager assumption of Aguirregabiria and Ho (2012):

each candidate station is evaluated in turn, outward from Montréal, by a manager who takes the current network as given, anticipates the fringe's free-entry response, and builds if the margin on the EVs the station induces exceeds its cost, which replaces a choice over 2^N networks with N ordered binary decisions. The fringe is solved analogously, and because entry raises adoption everywhere while neighbouring stations remain substitutable, I follow Vives (2005) and iterate best responses from the empty and full networks to obtain the Low and High extremal equilibria, which sustain different levels of adoption and hence converge to distinct fixed points. My theoretical contribution is a model in which complementarities and network externalities generate multiple equilibria. My empirical contribution is twofold. First, I provide evidence that a Big Push can overcome the coordination failure, even when the government carrying it out is imperfect. Second, I show that in a two-sided market such as EVs and charging, demand subsidies can worsen the coordination failure and complement the equilibrium selection mechanism.

2 Institutional Background & Data

In 2012, the province of Québec, through state-owned Hydro-Québec, launched the Electric Circuit, the first government 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 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 government 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.[4]

The government 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 government chargers using a CCS adapter.[5] [6] [7]

[4] 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*.

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

[6] 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.

[7] Tesla vehicles can use CCS fast chargers via a CCS1 adapter; beginning in 2023-2025, Tesla has also

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 government network provision affects EV adoption and market structure.

2.1 Network Evolution

Table A.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 government deployment: government 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, government provision accounted for about two-thirds of all new chargers, versus about 20% from Tesla and 13% from private operators. As a result, the network became more government-owned over time: the government 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%.

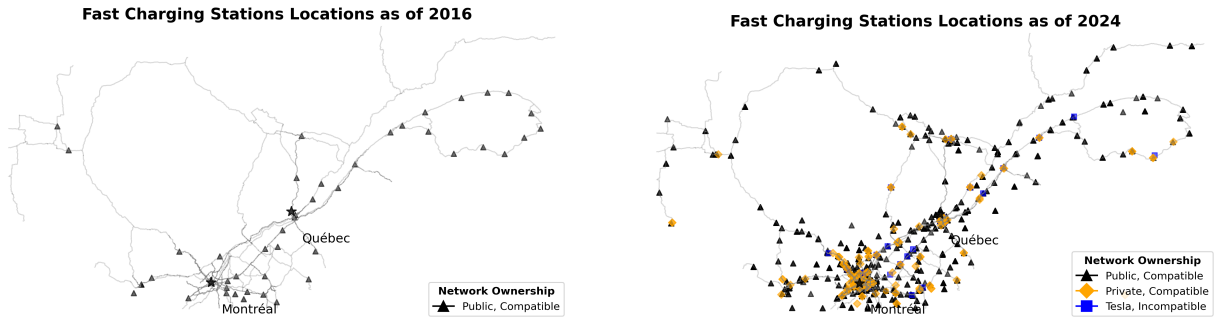
Table 1: Chargers by Network Type, 2016—2023

Year	All	government	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

Figure 1 shows the rapid expansion—and extensive coverage—of the government fast-

started opening select Superchargers to non-Tesla vehicles through NACS/Magic Dock, which post-dates the period studied here.

charging network as early as 2016, and how this network continued to grow, becoming much denser by 2023. Interestingly, there has been an increase in private fast-charging providers. However, these networks differ from the government ones, as they are denser near urban centers. From this figure alone, it remains unclear if this is a crowd-in or crowd-out story.



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

Figure 1: Government 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 Québec 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. [8]

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

[8] https://www.donneesQubec.ca/recherche/dataset/vehicules-en-circulation.

scraped trims with the registration data, conditional on model, brand, and weight.[9] [10]

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 government 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.[11] [12]

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 government network from 2017–2024. 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.

Charging Transaction

Table 2 summarizes the charging data. Panel A reports that the data contain 2,741,550 sessions generated by 134,804 unique consumers, or roughly twelve sessions per consumer per year. Grouping sessions into trips—defined as sequences of fast charges occurring within

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

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

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

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

Table 2: Summary Statistics: Charging Transaction Data, 2017–2023

Panel A. Sample	
Observations (charging sessions)	2,741,550
Unique consumers	134,804
Mean annual sessions per consumer	12
Mean annual trips per consumer	10
Panel B. Charging behavior	
Median state of charge at session start (%)	31
Median state of charge at session end (%)	74
Mean range added per session (km)	110
Mean routing distance from home (km)	127
Panel C. Fast-charging stops per trip	
One stop (%)	83
Two stops (%)	13
More than two stops (%)	3

Notes: Fast-charging transactions on Québec's Circuit électrique network, 2017–2023. Panel A: unit of observation is a charging session; consumers are identified by an anonymized `ConsumerID` stable over the sample period. Panel B: median state of charge (SOC) at session start and end, implied range added (computed from the change in SOC and the vehicle's rated battery capacity), and road-network distance from the station to the consumer's inferred home. Panel C: distribution of fast-charging sessions per trip, where a trip is a sequence of sessions within a 24-hour window.

a 24-hour window—yields about ten trips per consumer per year, implying an average of 1.2 fast charges per trip.

Panel B describes charging behavior within a session. The median session begins at 31% SOC and ends at 74%, adding approximately 110 km of range. The concentration of onset around one-quarter of battery capacity, and of termination just below 80%, is the pattern I document Figure 4. Sessions occur a mean routing distance of 127 km from the consumer's inferred home location, confirming that the network is used predominantly for intercity travel rather than for local use; the latter occur overwhelmingly on slow chargers, which I observe separately and is the focus of another paper (Cordeau et al., 2026).

Panel C reports the distribution of fast charges per trip. Eighty-three % of trips are completed with a single fast charge, thirteen % require two, and the remainder require more. This is the empirical basis for restricting the structural analysis to single-stop trips.

Car Registration

The second key input is the set of vehicles available to consumers. I merge SAAQ registration records with a model-level specification database covering price and range, the two characteristics that dominate the purchase decision in this market. Figure 2 plots the resulting choice set for 2022: each point is a trim-level BEV offered in Québec, with rated range on the horizontal axis and manufacturer's suggested retail price on the vertical axis. Five flagship models are labelled for reference.

The market comprises 21 distinct BEV models in 2022, accounting for 57,000 registered vehicles. The average listed price is $71,000, falling to $65,000 once federal and provincial purchase subsidies are applied, and the average range is 300 km.

The figure makes three features of the choice set apparent. First, the price–range relationship is convex and steep in the upper tail: moving from roughly 400 km to 600 km of range costs on the order of $70,000, whereas the same 200 km increment at the bottom of the range distribution costs closer to $20,000. Second, the models are not distributed continuously along the frontier but cluster into a low-range, low-price segment (Nissan LEAF, roughly 200 km) and a mid-range segment around 400 km (Tesla Model 3, Chevrolet Bolt), with a sparse and expensive long-range tail (Tesla Model S and X). Third, Tesla vehicles command a price premium conditional on range: the Model S and Model X sit well above the price–range locus traced by the remaining models, and the Model 3 is priced above comparable-range vehicles such as the Bolt.

The first two features bear on the demand side. The steep price of range in the upper tail means that a consumer who wants to complete long trips without charging must pay substantially for the privilege; whether a denser charging network is a substitute for that expenditure, or instead a complement to it, depends on preference parameters I estimate in Section 5.

The Tesla premium is not straightforward to interpret, and the ambiguity is central. A Tesla purchaser acquires both a vehicle and access to a proprietary, incompatible charging network. The observed premium therefore bundles the value of unobserved vehicle attributes—brand, performance, software—with the value of exclusive network access. Separately identifying the latter is part of this paper's contribution.

Subsidy

Policy variables encode Québec'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 Québec, the provincial rebate

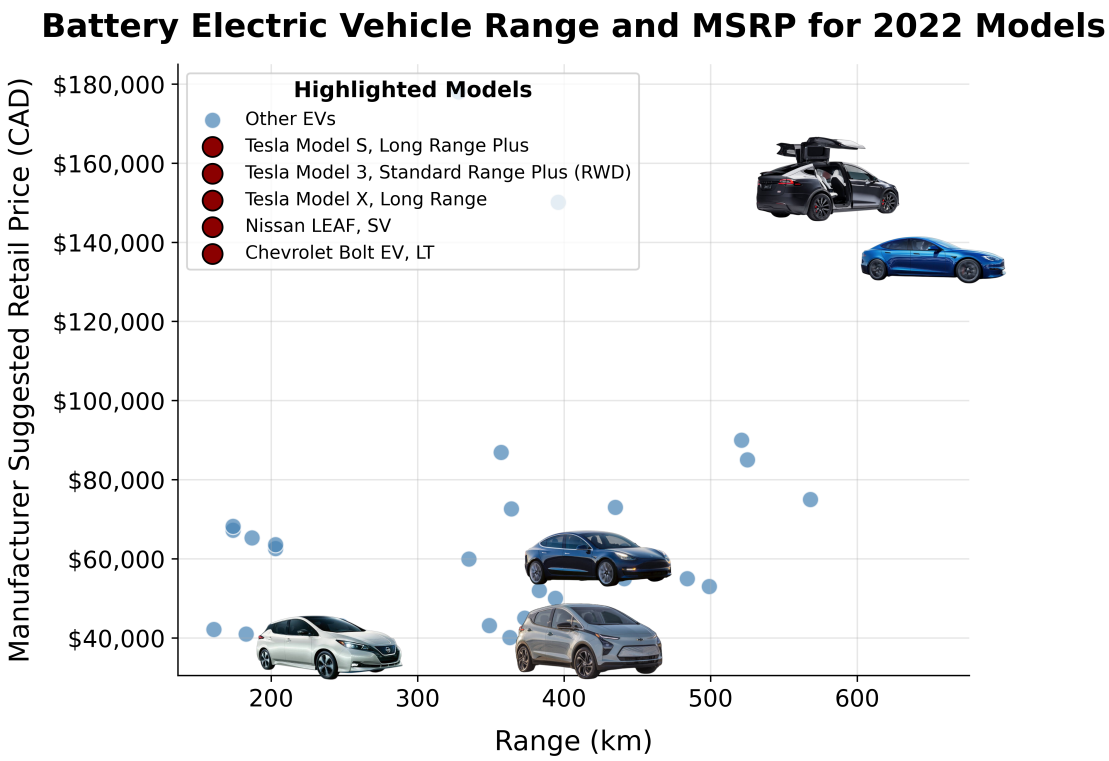


Figure 2: Battery Electric Vehicle Range and Manufacturer's Suggested Retail Price, 2022 Model Year

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 3 summarizes the structure of these programs.

Table 3: EV Purchase Incentives

Program	**Years**	**Ineligible above MSRP**	**BEV**	**PHEV**
Québec	2017—2019	BEV: > \$125k; PHEV: ≥\$75k	up to \$8k	up to \$8k
Québec	2020—2021	≥\$60k	\$8k	up to \$8k
Québec	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. Québec 2020—present PHEV amounts vary by kWh.

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.

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 government fast-charging network, focusing on Route 138 north of Québec City[13]. 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 Québec 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.

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 FEs α_c absorb time-invariant heterogeneity; year FEs 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 FEs, I estimate that connecting a city with gap no wider than 50 km raises log(EV share) by 0.788. In levels,

[13]See Figure A.1 in the appendix for the evolution of the network.

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 specification, I perform a sanity check in column (2) and find a nearly exact result, underscoring the robustness. Note that, as every city connected under the 50 km rule is also connected under the 100 km rule and that there are only 40 potential treated cities, having a similar result is expected.

Table 4: Average Treatment Effect, Network Connection

	(1) 50 km	(2) 100 km
log(EV Share)	0.788*	0.788*
	(0.3367)	(0.3285)
City FEs	Yes	Yes
Year FEs	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 3 reports these results. Estimation is doubly robust and uses not-yet-treated cities as controls at each event time.

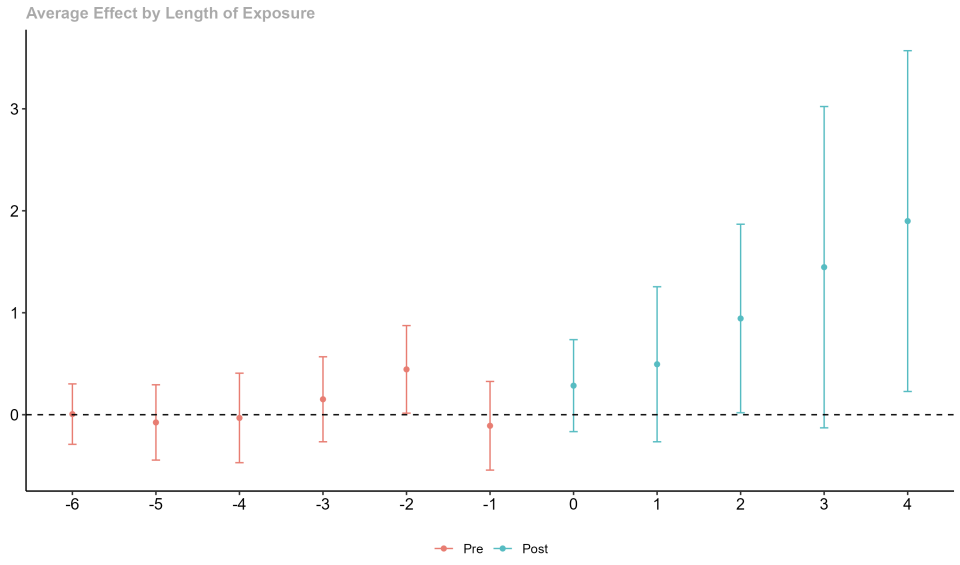


Figure 3: Dynamic Treatment Effects of Charging Network Connection on EV Adoption

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

3.2 Empirical Facts

This section documents five facts about fast-charging demand and supply that discipline the model. Facts 1, 2 and 3 show that fast charging is a remarkably standardized activity—one stop per trip, on the way home, initiated near a preferred battery level—which pins down the trip structure and the feasible choice sets on the demand side. Facts 4 shows that supply pairs a dominant government network with an atomistic compatible fringe alongside a VI firm, which pins down the market structure. Fact 5 contrasts Québec with British Columbia and motivates modeling Tesla as a vertically integrated firm that builds a proprietary network to sell vehicles.

Fact 1. Over 90% of fast-charging trips are completed with a single stop. Table 5 classifies trips involving at least one fast charge by the number of charging stops. A single stop completes 91.3% of trips by short-range models ($<$ 300 km) and 91.5% of trips by

long-range models ($\geq$ 300 km); two-stop trips account for about 7% in both classes, and longer sequences are rare. A single charging session is thus the overwhelming norm. Because fast charging is time-costly—and charging twice doubly so—this is unsurprising. What is striking is the near-invariance across range classes, which is consistent with consumers either selecting into the range class that matches their trip needs or, after purchase, undertaking only those trips that are worthwhile given the disutility of charging.

My model endogenizes both mechanisms by constructing, for each origin and vehicle model, the set of feasible destinations conditional on a single charging stop. Range-specific feasibility ties the vehicle purchase decision to the trips each model makes possible, capturing selection into range classes, while the outside option in the destination-choice model lets consumers forgo trips that are not worthwhile given the disutility of charging.

Table 5: Share of Trips by Number of Fast-Charging 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

Notes: Unique trips involving at least one fast-charging session on the Circuit électrique network, 2017–2023. Trips are constructed from transaction-level charging data as described in Section A.3.

Fact 2. Charging is concentrated around a 27% state of charge. I observe the battery state of charge (SOC) at the start and end of every session on the government network. Figure 4 plots the density of the SOC at the start of charging (left panel) and at the end of charging (right panel). Across range classes, sessions begin at an SOC centered around 27% and end around 80%.

I use these two regularities in the model. First, the 27% target generates heterogeneous station values: for each origin–destination pair and vehicle range, I project the SOC along the route, so a station's value depends on the battery level the consumer holds upon reaching it, with stations reached near the preferred 27% level being the most attractive. Second, the 80% endpoint disciplines the consideration set: since consumers rarely charge beyond this level, stations located at a distance requiring more than 80% of the battery to finalize the trip are deemed infeasible.

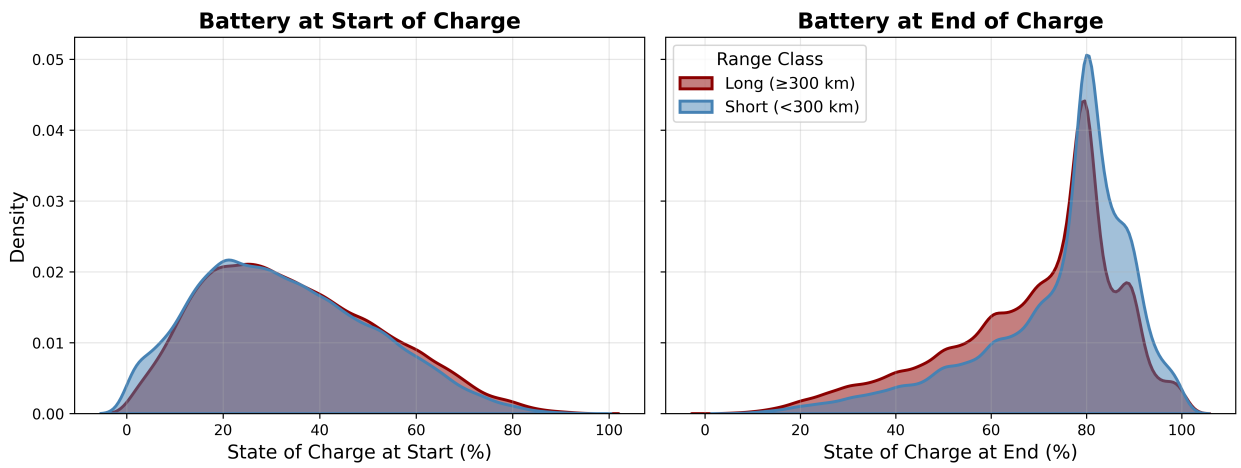


Figure 4: Distribution of the state of charge at the start (left panel) and end (right panel) of charging sessions.

Fact 3. Over 80% of single-stop charges occur on the return leg. Charging overwhelmingly takes place on the way home rather than en route to the destination. To show this, I exploit the `ConsumerID` panel: I retain days on which a consumer fast-charges exactly once—so each retained day is a unique single-stop event—and classify each event as outbound or return using feasibility logic anchored at the consumer's home (Appendix A.3). Conditional on the end-of-charge SOC, I project the SOC the driver would hold upon arriving home if returning immediately. The projected distribution is bimodal (Appendix Figure A.2): one mode implies an implausibly high arrival SOC, consistent with an outbound leg, while the other implies a moderate arrival buffer, consistent with return charging. Under this classification, more than 80% of single-stop events occur on the return leg. I embed this regularity in the model by assuming that consumers charge on the return leg: a trip takes the form $(o \to d,\ d \to s,\ s \to o)$, where the consumer travels from the origin city o to the destination d, stops at charging station s on the way back, and completes the leg home.

This structure pins down the value of each station along a route: given the vehicle's range, a station's position on the return leg determines the SOC the driver holds upon reaching it—and hence how close the stop falls to the preferred 27% level documented in Fact 2.

Fact 4. Québec's fast-charging market pairs a dominant government network with an atomistic compatible fringe. Table A.1 reports each network's share of stations (the extensive margin). The government owns 70.9% of sites and is by far the dominant player. ChargePoint's apparent 13.5% overstates its role: it is a hardware- and software-as-

a-service provider rather than a charging operator, and its stations are owned by dispersed local actors, so its share is better read as fringe ownership. Tesla owns 5.6% of sites; although its extensive-margin penetration is modest, its sites are larger and plausibly located at strategically chosen nodes to maximize vehicle sales through its vertically integrated proprietary network. FLO, the next-largest owner at 5.4%, both manufactures and operates stations, but its sites do not cluster spatially, so the network externality it internalizes is on the order of its share and I treat it as part of the fringe.

These observations motivate the market structure I impose: a dominant government network that moves first, a vertically integrated firm (Tesla) that moves second and invests to maximize EV sales, and a competitive fringe of atomistic single-station owners that moves last and captures none of the externality it generates. This is a key point of departure from Springel and Li (2025), whose analysis centers on the network externalities firms capture and the spillovers they impose on rival networks; I abstract from this inter-network channel—no private owner besides Tesla is large enough to internalize meaningful spillovers—and focus instead on how government provision shapes EV adoption and private build-out.

Table 6: 2023 Network Structure, Extensive Margin

Charging Network	**Stations (%)**
Government	70.9
ChargePoint	13.5
Tesla	5.6
FLO	5.4
Non-Networked	1.6
Petro-Canada	1.4
Electrify Canada	1.0
Couche-Tard	0.6

Fact 5. Tesla's vehicle share tracks its charging share: both are twice as large in British Columbia as in Québec. Table 7 compares Tesla's position in Canada's two largest EV markets in 2022. In British Columbia, where the government network holds under 10% of fast-charging capacity, Tesla operates 41% of the network and sells 56% of the battery-electric fleet. In Québec, where the government network holds 59%, Tesla's charging share falls to 21% and its vehicle share to 30%. The composition of the fleet shifts accordingly: the Model 3 and Model Y jointly account for half of British Columbia's registrations but only 29% of Québec's, while affordable compatible models—the Bolt, Kona, and Leaf, priced around $40,000 against $67,500 for a Model 3—make up 37% of Québec's fleet versus 20% in British Columbia (Appendix Tables A.3–A.5). The comparison is descriptive rather than

causal, but it is consistent with the two mechanisms at the heart of the model: a vertically integrated firm supplies charging in order to sell vehicles, and government provision of a compatible network is distributional—it erodes the proprietary network's advantage and shifts adoption toward cheaper compatible models.

Table 7: Tesla's Position in Québec and British Columbia (2022)

	Québec	British Columbia
Tesla share of BEV stock	0.30	0.56
Tesla share of fast-charging capacity	0.21	0.41
government network share of fast-charging capacity	0.59	0.09
Model 3 + Model Y share of BEV stock	0.29	0.50
Bolt + Kona + Leaf share of BEV stock	0.37	0.20

Notes: The government network is Circuit électrique in Québec and BC Hydro in British Columbia. Entries are computed from Appendix Tables A.3–A.5.

4 Model

The reduced-form findings of Section 3.1 indicate that consumers respond to the deployment of the charging network. To evaluate the network's impact under counterfactual policy scenarios and to account for the endogeneity of private station placement, I introduce a spatial equilibrium model. The model has two salient features. On the demand side, it captures how consumers use the charging network and aggregates this behavior into a network-quality measure that explains EV adoption. On the supply side, it models the charging-station investment decisions of both the vertically integrated firm and the competitive fringe. Fringe firms make free-entry decisions: taking rivals' entry as given, they do not compete on price or quantity; instead, realized charging demand at a node is split evenly across the stations located there. The VI firm, by contrast, builds a proprietary network to internalize part of the externality it generates, earning profits through vehicle sales rather than charging sessions.

In this section, I first present the EV demand model, then describe how the network-quality measure is constructed, introduce the firms' profit functions, and conclude with the equilibrium concept and its computation.

4.1 EV Demand

Consumers solve a three-stage decision problem. First, they choose a vehicle model; second, they choose which trips to undertake; and third, they choose where to charge along the way. I solve the problem by backward induction and condense the second and third stages into a network-quality index, Λ_{ojt}, which summarizes the expected utility of travel afforded by the charging network. The index varies across origin cities o and over time t, as the network expands, and across vehicle models j through driving range and, where applicable, compatibility.

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.

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 and charging network valuation. [14]

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 $\theta = (\alpha, \beta, \sigma_p, \sigma_\Lambda)$. I recover $\xi_{ojt}(\theta)$ as the

[14] Yet to be done.

residual term implied by observed market shares, following the inversion approach in Berry et al. (1993).

Network Quality and Trip Demand

This subsection maps the charging network $\mathcal{K}_{ojt}$ into a microfounded measure of network quality, Λ_{ojt}, the ex-ante expected utility that a consumer of vehicle model j residing in origin o at time t derives from access to the charging network. Constructing Λ_{ojt} requires (i) a model of the value the charging network delivers along each origin-destination route, and (ii) a model of origin-destination trip demand and the associated preferences for destination. The key point is that network quality is not captured by simple charger counts: 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.

The section proceeds in three layers. First, I microfound route-level charging utility with a discrete-choice model of charging-station selection along a given route (Section 4.1). Even when multiple stations are feasible, they deliver heterogeneous utility—some are systematically more valuable than others for a given origin-destination pair—up to a type-I extreme-value preference shock. Second, I summarize the implied charging value for each origin-destination pair and embed it in a destination-choice model (Section 4.1), in which consumers choose among destinations based on destination preferences and the charging path each destination entails. Third, I aggregate these objects into the ex-ante expected value of the network (Section 4.1).

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 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 for details on $\mathcal{S}_{odjt}$).

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

$$V^c_{iodsj} = \eta\, b_{odsj} + \lambda_{min}\text{Detour}_{ods} + \kappa \mathbb{1}\{\text{Alone within 10 km}\} + \zeta_s + \varepsilon^c_{iodsj}, \qquad s \in \mathcal{S}_{odjt} \quad (5)$$

The term η captures a preference over the state of charge at which charging is initiated, λ_{min} the minutes detour in the spirit of Houde (2012), κ captures the value of being the sole charging station within a 10 km radius, ζ_s captures unobserved station attributes and ε^c_{iodsj} is an EVT1 taste shock, implying a logit choice rule conditional on feasibility.

The specification b_s is a quadratic loss around a preferred charging-onset state of charge

b_s:

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

I set $\overline{\text{SOC}} = 0.27$ to match a salient empirical regularity: consumers disproportionately initiate charging when the battery is around 27% 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.

Charging Path Ex-ante Value. Despite its simplicity, the charging model generates substantial heterogeneity in the utility provided by charging stations (see Figure A.4). For a given trip (o, d) and vehicle type j, each station provides a different level of utility, with no added value 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 based on the quality of the charging network.

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 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 location 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.

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 for details on $\mathcal{D}_{ojt}$).

Utility from choosing destination d is

$$V^d_{iodj} = \lambda\, d_{od} \;+\; \zeta_d \;+\; \underbrace{\log\left(1 + \sum_{s \in \mathcal{S}_{odjt}} \exp\big(\bar{V}^c_{odjst}\big)\right)}_{OD_{odjt}} + \varepsilon^d_{iodj}, \qquad d \in \mathcal{D}_{ojt}. \tag{8}$$

The term d_{od} denotes origin-destination routing 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 $\mathcal{K}_t$ 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, 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.

Network Quality This subsection defines the microfoundation of network quality, Λ_{ojt}, which summarizes the ex-ante value of access to the charging network for a consumer residing in origin o and 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 trips they will need to undertake and where they will charge. I therefore construct Λ_{ojt} as the expected charging utility across potential destinations, using origin-specific trip-demand weights that reflect consumers' baseline destination preferences.

The charging network affects EV adoption primarily by relaxing charging frictions on origin-destination trips. Prior to purchasing a vehicle, a consumer in origin o faces uncertainty about which destination d will arise, while the distribution is known. Some destinations are highly valued, while others are not. A network that improves charging convenience on frequently traveled routes is more valuable than one that does so 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 Λ_{ojt} twice, 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. I model pre-purchase destination uncertainty by assuming that consumer i draws a finite set of potential trip destinations, $\mathcal{J}_{io} = \{d_{i1}, \ldots, d_{iM}\}$, whose elements are 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 ex-post value of the network depends on the realized set $\mathcal{J}_{io}$, while Λ_{ojt} is its ex-ante expectation under $P_0(\cdot \mid o)$. In particular, since each element of $\mathcal{J}_{io}$ is distributed according to $P_0(\cdot \mid o)$, 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 for origin-destination preferences that excludes the charging component, leaving (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 $\lambda_{\text{dist}} < 0$. The resulting origin-specific baseline destination distribution is

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

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 attractive and nearby destinations. Second, it is consistent with the discrete-choice structure and can be estimated in the destination-choice step. Third, by excluding the charging network, it yields weights that are invariant to the network, avoiding double counting when constructing network quality.[15]

Network quality index. For each origin-destination pair (o, d), vehicle model j, and time t, the charging network determines the route-level value of charging, OD_{odjt}, defined in (7). I aggregate these route-level objects into a single index of network quality by taking their expectation under the baseline destination distribution:

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

Intuitively, Λ_{ojt} is high when the network provides good charging opportunities for vehicle model j on the trips that residents of o are most likely to take. Because OD_{odjt} depends both on feasibility and on the distribution of charging utility across stations along the route, Λ_{ojt} captures extensive-margin improvements (more feasible routes) as well as intensive-margin improvements (better-located charging opportunities on feasible routes), in line with Springel and Li (2025), who show the role of charging connectivity in EV charging demand. More precisely, if no charging station is feasible for a given origin-destination pair, i.e., $\mathcal{S}_{odjt} = \emptyset$, then $OD_{odjt} = \log(1) = 0$. Moreover, adding any feasible charging option strictly increases route value: if a previously infeasible station $s^\star \notin \mathcal{S}_{odjt}$ becomes feasible, 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 one, and an infeasible station adds none.

Charging Demand

While the first stage of the demand estimation links the charging network $\mathcal{K}$ to EV adoption, $\mathcal{E}(\Lambda)$, through network quality, consumer trip and charging choices also matter on the supply side through the charging demand they imply. 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.[16] This aggregation collapses

[15] 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 $P_0(d \mid o)$.

[16] EV adoption is itself implied by the network: since Λ is a function of $\mathcal{K}$, I write $\mathcal{E}(\mathcal{K}) \equiv \mathcal{E}(\Lambda(\mathcal{K}))$.

complex trip and charging patterns into a scalar demand for each charging node.

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 o, and let M denote the expected number of trips per vehicle per year. The implied charging demand at node r is

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

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 a trip from o to d results in a charging session at station s. Summing over $s \in r$ aggregates the probability mass of charging at any station located in node r; collecting node-level demands yields the mapping introduced above, $\mathcal{C}(\mathcal{K}) \equiv \{Q_r(\mathcal{K}, \mathcal{E}(\mathcal{K}))\}_r$. Equation (11) is used only in the counterfactual analysis; because the counterfactual is a static game, I suppress time variation and drop the t subscript.

4.2 Charging Firm Supply

Three types of players form the supply side: the government, the VI firm, and the competitive fringe. Consumer behavior enters firms' problems through the demand objects derived above: network quality $\Lambda(\mathcal{K})$, EV adoption $\mathcal{E}(\Lambda)$, and charging demand $\mathcal{C}(\mathcal{K}, \mathcal{E})$. For simplicity, I consider consumer behavior as being embedded in firms' profits. The government moves first and (unless otherwise specified) does not behave strategically. The VI firm moves second: it observes the government's deployment and, when choosing its own network $\mathcal{K}^I$, anticipates the competitive fringe's best response. The fringe moves last: it consists of atomistic firms that move simultaneously and enter wherever expected profits are non-negative.

Following Li (2019), I treat supply as static and assume that static profits are proportional to dynamic values, so that the static entry condition is behaviorally consistent with a long-run dynamic game. Government stations are taken as given and are compatible with all EV models.

The Government

While the government does not act strategically, I specify the objective that rationalizes its deployment. The government chooses its network $\mathcal{K}^G = \{k_r^G\}_{r \in \mathcal{R}}$ to maximize equally weighted network quality across origins, subject to a capacity constraint:

$$\max_{\mathcal{K}^G} \; G(\mathcal{K}^G) = \sum_{o,j} \bar{\Lambda}_{oj}(\mathcal{K}) \qquad \text{s.t.} \qquad \sum_{r} k_r^G \leq \bar{K}^G, \tag{12}$$

where $\bar{\Lambda}_{oj}(\mathcal{K})$ is network quality in origin o averaged across the EV models available in 2022.

The Vertically Integrated Firm

The VI firm is the strategic player. It observes the government's deployment and invests in nodes to stimulate its own EV sales, internalizing the fringe's best response. By offering a proprietary charging network, the VI firm vertically differentiates its product and thereby raises demand for its vehicles. In this application, the VI firm is Tesla, and each vehicle sold earns a fixed per-vehicle margin (PCM). The VI firm is in the business of selling EVs, not charging sessions; its investment decisions therefore weigh only the effect on its vehicle sales, taking the fringe's response into account.[17] The VI profit function is

$$\Pi^{VI} = \sum_{o} \sum_{j \in \mathcal{T}} \text{Share}_{oj}(\Lambda)\, M_o\, \text{PCM} \;-\; \sum_{r} k_r^I\, F^I, \tag{13}$$

where $\text{Share}_{oj}(\Lambda)$ is the market share of Tesla model $j \in \mathcal{T}$ in origin o, which depends on network quality Λ_{oj}; M_o is the market size of origin o; PCM is the per-vehicle profit margin; k_r^I is the (endogenous) number of incompatible chargers in region r; and F^I is their per-charger fixed cost.

The Competitive Fringe

The entry model follows Bresnahan and Reiss (1991): the number of charging stations in a region 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, Q_r, defined in (11), and a per-charger demand $q_r(k_r) = Q_r/k_r$ that declines with the number of active chargers. Given the per-session price p, per-session variable cost c, and a fixed cost F, entry continues until the marginal entrant earns zero expected profit.

Each fringe firm owns and operates a single station; this single-owner benchmark represents a competitive economy without network externalities. The per-period profit of a single-owner station in region r is

$$\Pi_r^{so}(\mathcal{K}, \mathcal{E}) = \underbrace{(p - c)\, q_r(k_r)}_{\text{profit}} \;-\; \underbrace{F}_{\text{entry cost}}. \tag{14}$$

[17] This is supported by the fact that in the early deployment of the Supercharger network, charging sessions were free and restricted to Tesla vehicles. Following the National Electric Vehicle Infrastructure (NEVI) program, Tesla opened its NACS connector standard and later laid off a large part of its charging team, underscoring its willingness to let other firms build out the network. Only in 2024 did Tesla begin granting Supercharger access to other automakers such as Ford and GM.

Entry decision. Firms enter as long as their expected payoff covers the entry cost. In equilibrium, the number of single-owner stations in region r satisfies

$$(p-c)\,q_r(k_r) - F \;\geq\; 0 \;>\; (p-c)\,q_r(k_r+1) - F. \tag{15}$$

The Local-Manager Argument

The equilibrium is not computed by joint optimization over all nodes. A firm choosing its network jointly would compare 2^N configurations, so I instead let players evaluate nodes one at a time, in the spirit of Aguirregabiria and Ho (2012). The VI firm considers candidate nodes in order of proximity to Montréal, which can be read as a local manager assessing what a station at that node contributes to firm-wide profit. The device pins down a single network for the VI firm at given primitives. It does not remove the multiplicity at the fringe stage, which is the object of interest and which I bracket below.

The VI firm's rule is as follows. The nearest node is built if the additional margin on the EVs it induces exceeds its build-out cost. Because Tesla drivers can charge at both compatible and incompatible stations while other drivers can use only compatible ones, a VI build-out moves fringe entry in two directions: it steals charging demand from nearby compatible stations, and it raises EV adoption and hence charging demand everywhere. The local manager's calculation therefore incorporates the fringe's equilibrium response, which can amplify or attenuate the marginal value of an incompatible station at that node. The second-nearest manager then faces the analogous decision, and so on until every candidate node has been evaluated.

The fringe expands outward from the central node in the same way. The two forces just described operate here as well: strategic complementarity through EV adoption, and strategic substitutability through business stealing. The first initialization places no private station at any node, the second places a station at every node. Each iteration removes stations that are unprofitable given the current network, then considers the nearest unbuilt node and builds it if entry is profitable. The process continues outward until no unbuilt node is profitable and no standing station is unprofitable.

Local substitutability can produce cycling. Two neighbouring nodes are each profitable in isolation, but entry at either makes the other unprofitable, so the two alternate across iterations. When a cycle is detected I keep the more profitable of the two stations and treat the surrounding region as saturated, making it ineligible for further investment, so that the next candidate is the nearest unbuilt node outside a saturated region.[18] I refer to the two

[18] An earlier version of the algorithm dropped only the less profitable station and left the region open to entry. The resulting equilibria are close to those reported here.

resulting configurations as the Low and High equilibria.

4.3 Equilibrium

Let $\mathcal{R}$ denote the set of regions (nodes). A charging network is $\mathcal{K} = \{k_r, k_r^I\}_{r\in\mathcal{R}}$, where $k_r \in \mathbb{N}_0$ is the number of compatible chargers and $k_r^I \in \mathbb{N}_0$ the number of incompatible (VI) chargers in node r. Let $\Lambda(\mathcal{K})$ denote the vector of network-quality indices $\{\Lambda_{oj}(\mathcal{K})\}$. Let $\mathcal{E}(\Lambda)$ denote the resulting EV adoption and model shares (the EV stock) from the vehicle-demand problem given network quality Λ. Let $\mathcal{C}(\mathcal{K}, \mathcal{E}) = \{Q_r(\mathcal{K}, \mathcal{E})\}_{r\in\mathcal{R}}$ collect the region-level charging demand defined in (11), and assume that demand is split evenly across chargers within a region. Per-charger profits $\Pi_r^{so}(\mathcal{K}, \mathcal{E})$ are given by (14).

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 and charging behavior given the realized network.

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

3. **Fringe free entry:** for each region $r \in \mathcal{R}$, the number of compatible chargers k_r^* satisfies the integer free-entry conditions

$$\Pi_r^{so}(\mathcal{K}^*, \mathcal{E}^*) \geq 0 \;\text{ if } k_r^* \geq 1, \qquad \Pi_r^{so}(\mathcal{K}^* + \mathbf{e}_r, \mathcal{E}^*) < 0,$$

   where $\mathbf{e}_r$ adds one compatible charger in region r. 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.

4. **VI sequential optimality:** given the distance-to-Montréal ordering of candidate nodes, every VI station in $\mathcal{K}^*$ was incrementally profitable when evaluated, and no further build-out is profitable once fringe firms and consumers re-equilibrate:

$$\Pi^{VI}\big(\mathcal{K}^*(\mathbf{e}_r^I), \mathcal{E}^*(\mathbf{e}_r^I)\big) \;<\; \Pi^{VI}(\mathcal{K}^*, \mathcal{E}^*) \qquad \text{for all } r,$$

   where $\mathbf{e}_r^I$ denotes one VI station in region r, and $\big(\mathcal{K}^*(\mathbf{e}_r^I), \mathcal{E}^*(\mathbf{e}_r^I)\big)$ is the equilibrium reached after the fringe and consumers respond.[19]

[19] The VI firm deploys stations in groups of eight chargers, so $\mathbf{e}_r^I$ raises k_r^I by eight.

Conditions (1)–(4) 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** $\tilde{\mathcal{K}}$:
   (a) Low equilibrium: $\tilde{\mathcal{K}} \leftarrow \mathbf{0}$
   (b) High equilibrium: $\tilde{\mathcal{K}} \leftarrow \mathbf{0}$
2. **Government deployment:** $\tilde{\mathcal{K}} \leftarrow \tilde{\mathcal{K}} + \mathcal{K}^G$
3. **Consumers' beliefs:** set $\mathcal{K} \leftarrow \tilde{\mathcal{K}}$
4. **Trip and charging stage:** given $\mathcal{K}$, solve consumers' trip and charging problem and compute $\Lambda \leftarrow \Lambda(\mathcal{K})$
5. **Vehicle stage:** given Λ, solve vehicle demand and compute EV adoption $\mathcal{E} \leftarrow \mathcal{E}(\Lambda)$
6. **Fringe entry stage:**
   (a) Set firms' beliefs $\hat{\mathcal{K}} \leftarrow \tilde{\mathcal{K}}$
   (b) Given $(\hat{\mathcal{K}}, \mathcal{E})$, form the implied charging demand $\mathcal{C}(\hat{\mathcal{K}}, \mathcal{E})$ and per-charger profits $\Pi_r^{so}(\hat{\mathcal{K}}, \mathcal{E})$
   (c) Add one charger $\mathbf{e}_r$ at the node nearest to Montréal that (i) is not yet in $\tilde{\mathcal{K}}$ and (ii) is profitable given $(\hat{\mathcal{K}}, \mathcal{E})$, i.e., $\Pi_r^{so}(\hat{\mathcal{K}} + \mathbf{e}_r, \mathcal{E}) \geq 0$; update $\tilde{\mathcal{K}} \leftarrow \tilde{\mathcal{K}} + \mathbf{e}_r$
   (d) If a charger was added, return to step 3; otherwise set $\mathcal{K}^* = \tilde{\mathcal{K}}$, $\Lambda^* = \Lambda(\mathcal{K}^*)$, and $\mathcal{E}^* = \mathcal{E}(\Lambda^*)$
7. **VI entry stage:**
   (a) Set $i \leftarrow 1$ and order candidate nodes by distance to Montréal
   (b) Given $\mathcal{K}^*$, consider a build-out $\mathbf{e}^I_{r_i}$ at the i-th closest node
   (c) Perform steps 3–6 with $\tilde{\mathcal{K}} \leftarrow \mathcal{K}^* + \mathbf{e}^I_{r_i}$ to obtain the counterfactual equilibrium $(\mathcal{K}^*(\mathbf{e}^I_{r_i}), \Lambda^*(\mathbf{e}^I_{r_i}), \mathcal{E}^*(\mathbf{e}^I_{r_i}))$, in which fringe firms and consumers re-equilibrate in response to the VI station
   (d) If the build-out is incrementally profitable, i.e., $\Pi^{VI}\big(\mathcal{K}^*(\mathbf{e}^I_{r_i}), \mathcal{E}^*(\mathbf{e}^I_{r_i})\big) \geq \Pi^{VI}(\mathcal{K}^*, \mathcal{E}^*)$, set $\mathcal{K}^* \leftarrow \mathcal{K}^*(\mathbf{e}^I_{r_i})$, $\Lambda^* \leftarrow \Lambda^*(\mathbf{e}^I_{r_i})$, and $\mathcal{E}^* \leftarrow \mathcal{E}^*(\mathbf{e}^I_{r_i})$
   (e) Set $i \leftarrow i + 1$. If $i > \bar{\imath}$, the number of candidate nodes, stop; otherwise return to step 7(b)

Note that once $\mathcal{K}^G$ and the VI stations $\mathcal{K}^I = \sum_r \mathbf{e}_r^I$ are built, they remain in the network: the government and the VI firm are the first and second movers, respectively.

5 Estimation

This section presents the estimation strategy. The empirical analysis is separated into demand and cost estimation. The first block focuses on estimating charging parameters, destination preferences, and EV model preferences. The second estimates the free-entry cost of the fringe and the PCM of the vertically integrated firm, and calibrates the market size. I conclude with a validation exercise showing that the estimated parameters closely replicate the 2022 market.

5.1 Demand

Charging Parameters

This subsection estimates the parameters governing charging behavior. 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 observed station choices primarily reflect preferences rather than simply unreachable trips. Formally, when a consumer chooses station s over another feasible station s' in the feasible set $\mathcal{S}_{odjt}$, 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. 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 through availability, 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 and model types, implying that network quality is inherently route-model-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. Second, I limit detours by imposing a maximum detour of 5 km. Finally, feasibility is constrained by the battery technology: the remaining battery range at the station must be

sufficient to reach the final destination conditional on a max state of the battery of 80% at the end of the charge. This pattern is consistent with Figure 4, which shows that charging events rarely go beyond 80%. Together, these restrictions ensure that $\mathcal{S}_{odjt}$ captures the stations that are truly considered.

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 government 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 27%, 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.3.

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

$$V_{ns}(\theta_s) \;=\; \eta\, b_{ns} + \lambda_{min}\text{Detour}_{ods} + \kappa\mathbb{1}\{\text{Alone within 10 km}\} + \zeta_s + \varepsilon_{ns}, \qquad \theta_s = (\eta, \lambda_{min}, \kappa, \{\zeta_s\}) \tag{16}$$

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 27%. The penalty

is specified as a quadratic function, $(b_{ns} - 0.27)^2$, 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:

$$P_{ns}(s \mid o, d; \theta_s) = \frac{\exp(\eta\, b_{ns} \;+\; \lambda_{min}\text{Detour}_{ods} \;+\; \kappa \log(n_{50km}) \;+\; \zeta_s)}{\sum_{s' \in \mathcal{S}_{odjt}} \exp(\eta\, b_{ns'} \;+\; \lambda_{min}\text{Detour}_{ods'} \;+\; \kappa \log(n_{50km}) \;+\; \zeta_{s'})} \tag{17}$$

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 8 shows that, for a given trip, stations are perceived as heterogeneous (see Figure A.4). Namely, drivers prefer to initiate charging near a target SOC of 27%, which varies across origin, destination, model type, and network configuration.

Battery deviation is negative and statistically significant in all specifications where it is included, confirming a robust preference for initiating charging near the 27% target SOC. This consistency indicates that drivers incur meaningful utility losses when charging above or below their preferred battery level.

Detour costs are also negative and significant across specifications, confirming that time-consuming deviations relative to the shortest duration path reduce station choice probability. In column (4), κ shows that being the only charging station within a 10 km radius increases the likelihood of being chosen. This pattern is consistent with business-stealing effects and substitution across nearby stations.

Columns (1) and (2) include station fixed effects, while columns (3) and (4) do not. Specifications (3) and (4) exclude station fixed effects for two reasons. First, station fixed effects absorb all station-specific variation, leaving little identifying variation for coefficients that are correlated with persistent station attributes, and therefore mechanically attenuating economically relevant effects. Second, in the counterfactual analysis, I do not observe the unobserved preference components of private compatible charger or the Tesla stations; only observable characteristics are available. To maintain consistency between estimation and counterfactual simulations, I therefore rely on station heterogeneity from observable characteristics. Finally, on average, only 15 out of the roughly 400 existing government charging

stations are accessible for a given origin–destination pair.

Table 8: Charging Station Choice Estimation

	(1)	(2)	(3)	(4)
Battery Deviation at Start, η	-0.01^{***}	–	-0.07^{***}	-0.06^{**}
	(0.002)		(0.010)	(0.025)
Detour (minutes), λ_{min}	–	-0.05^{***}	-0.12^{***}	-0.11^{***}
		(0.001)	(0.008)	(0.030)
Alone within a 10 km Radius, κ	–	–	–	0.70^{**}
				(0.344)
Station FEs, ζ_s	Yes	Yes	No	No
Realized charging events	61,505	61,505	61,505	61,505
Mean choice set size	15	15	15	15
N Observations	916,044	916,044	916,044	916,044

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 5 km 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 at 80% of the battery. Standard errors are based on the inverse of the quasi-Newton Hessian approximation returned by the MLE routine.
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Destination Parameters

This subsection estimates long-trip destination demand. Using the inferred destinations trips observed in 2023, I estimate a maximum likelihood model of destination choice. 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, implemented as a reachable factor of 1.8 (i.e., up to 180% of the vehicle's range). This proxy corresponds to a one-stop long trip that starts charging at a low battery level (around $0 - 20\%$) and ends at a high battery level (around $80 - 100\%$).

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.

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

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

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:

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

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 network utility as an ex-ante object.

Estimation Results. Table 9 reports maximum likelihood estimates from the destination-choice model. Column (1) estimates a single continuous cost shifter—origin–destination

distance — while controlling for network quality. Column (2) additionally includes destination fixed effects, $\{\zeta_d\}$, which absorb baseline destination attractiveness (e.g., amenities, population, employment centers, tourism appeal, and other persistent attributes). These estimates discipline the key determinants of long-trip destination choice, which map into the origin-destination weight equation (10) and, in turn, deliver the ex-ante network value through (9).

The estimated distance coefficient is negative across specifications. In column (1), $\hat{\lambda} = -0.010$, consistent with standard spatial demand: distance is costly. In contrast, once destination unobserved preferences are included in columns (2), the point estimate remains similar, $\hat{\lambda} \approx -0.013$, but it becomes imprecise. Thus, conditional on feasibility and destination preferences, distance is not statistically different from zero.

Table 9: Results of Destination Estimation

	Alternative Specifications	
	(1)	(2)
Origin-Destination Distance, λ_{dist}	-0.010^{***}	-0.013
	(0.000)	(0.273)
Destinations FEs, ζ_d	–	Yes
Realized O-D trips	106,955	106,955
Mean choice set	53	53
N Observations	5,623,210	5,623,210

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 row corresponds to an origin–destination alternative within a trip-specific choice set. Column (1) includes only origin–destination distance. Column (2) additionally absorbs destination FEs ζ_d. Standard errors are based on the inverse of the quasi-Newton Hessian approximation returned by the MLE routine. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

On average, 53 out of 100 cities are accessible. Cities can be inaccessible for two reasons: (i) the vehicle cannot reach the destination even with a charging stop, or (ii) there is no feasible charging station along the route. Because I use the 2023 charging network, I attribute the remaining inaccessibility to destination preferences and to the intensive margin of network quality.

EV Demand

Using the 2014-2022 car registration data, I estimate the EV demand parameters. 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_{ojt} denotes the market share of model j in city o in year t, including only battery electric vehicles (BEVs). The outside option, s_{ot}^{O}, corresponds to the share of non–BEVs. In this specification, x_j collects vehicle characteristics and controls, p_{jt} is the price of model j, and Λ_{ojt} measures the charging–network quality relevant to consumers in city o considering model j. The consumer-vehicle-specific taste shock, μ_{iojt}, 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 $U_{io0} = 0$. 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 government provision being the primary driver. This time variation is crucial for identification, as government 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.[20]

[20]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

Because vehicle prices are endogenously set by manufacturers, I instrument them with cost shifters. First, I use the exchange rate between the country where each vehicle is produced and Canada, which shifts the cost of imported vehicles and components. Second, I use lithium prices and their interaction with battery size, which capture fluctuations in global commodity markets that move marginal costs differentially across models. The interaction with battery size, a key cost component of EVs, is crucial: capacity varies substantially across models with the same range, so lithium price shocks move marginal cost without moving the attribute consumers value. Lithium price shocks therefore shift marginal costs through capacity while, conditional on range, leaving consumer utility unaffected. These instruments generate supply-side variation in prices that is plausibly orthogonal to unobserved local demand shocks, delivering consistent estimates of the price coefficient.

Estimation Results. Table 10 reports demand estimates using city-year-model observations and a rich set of unobserved preference controls. Across specifications, the price coefficient is negative and tightly estimated (around -0.028 to -0.044), implying meaningful price sensitivity.

The positive and statistically significant coefficient on Λ_{ojt}, the network-quality variable, implies that consumers prefer a denser charging network, consistent with previous studies on EV demand (Li, 2019; Springel and Li, 2025; Springel, 2021; Remmy, 2025). The estimate is stable across specifications, ranging from approximately 0.73 to 0.86. The mean own-price elasticity ranges from 1.98 to 3.15. These estimates lie within the range commonly found in the EV demand literature, where elasticities typically fall between -2 and -4 (Springel, 2021; Li, 2019; Muehlegger and Rapson, 2022; Remmy et al., 2022), and are on the lower end of the automobile demand literature, where they range from -4 to -7 (Berry et al., 1993; Goldberg, 1995; Petrin, 2002; Berry et al., 2004).

I also compute the EV category elasticity of Λ and find it to be between 1.25 and 1.44. This implies that a 1% increase in network quality increases the overall EV adoption rate by roughly 1.3 to 1.4%, thereby reducing the share of the non-EV outside option.

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 I study.

Table 10: Result of EV Demand Estimation

	Alternative Specifications		
	(1)	(2)	(3)
Price	-0.044^{***}	-0.028^{***}	-0.029^{***}
	(0.003)	(0.003)	(0.003)
Λ_{ojt}	0.742^{***}	0.856^{***}	0.732^{***}
	(0.067)	(0.067)	(0.153)
Range (per 10 km)		-0.058^{***}	-0.058^{***}
		(0.008)	(0.008)
$\Lambda_{ojt}\times$ Range (per 10 km)			0.002
			(0.003)
Mean own-price elasticity	−3.15	−1.98	−2.05
EV category elasticity of Λ	1.25	1.44	1.38
Product, ξ_j	Yes	Yes	Yes
City×Year, μ_{ct}	Yes	Yes	Yes
Supply Instruments	Yes	Yes	Yes
N (Observations)	19,400	19,400	19,400

Notes: Each observation is a city-year-model-trim (2016–2022). The regression is restricted to BEVs; the outside option includes hybrids and gasoline vehicles. Column (1) includes Price and Λ_{ojt}. Column (2) adds vehicle Range (per 10 km). Column (3) further includes the interaction $\Lambda_{ojt}\times$Range. All specifications absorb multiple layers of unobserved demand heterogeneity. Product×year unobserved demand components ξ_{jt} capture model-specific shocks common across markets in a given year. City×year unobserved demand absorbs local market conditions and policy environments. Standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Surprisingly, range has a negative coefficient in columns (2) and (3). This is due to an identification issue: range is used to construct Λ, the measure of network quality. Although the two variables are not perfectly collinear, they are highly correlated, making it difficult to separately identify their effects on EV adoption. I examine this issue further in Table A.10. Specifically, I examine whether increasing EV range by 100 km while keeping the charging network constant raises consumers' WTP. A negative amount would point to a problem with the estimation. On the contrary, I find that adding 100 km of range increase the WTP by \$6,600, which is roughly the cost of providing an additional 100 km of range. By comparison, increasing network quality by one standard deviation raises WTP by about \$15,000. This confirms that column (1) captures both range and network improvements.

Readers interested in the mechanisms through which charging stations affect consumer decisions can find additional discussion in Section A.3.1. There, I estimate the marginal value of adding a charging station and show that its value varies substantially across locations. In particular, Figure A.4 shows that the distribution of average WTP across charging stations is approximately exponential.

For the cost estimation and counterfactual simulations, I restrict attention to the specification in column (1) for several reasons. First, it implies a price elasticity of -3.14, which lies squarely within the range typically reported in the EV demand literature (between approximately -2 and -4). Second, both the marginal value of range and the marginal value of network quality are strictly positive, as economic intuition would suggest and are of reasonable magnitude. By contrast, columns (2) and (3) imply that increasing vehicle range by 100 km reduces consumers' WTP, which is inconsistent with both intuition and the evidence presented in Table A.10. Finally, I estimate the column (1) specification using vehicle-model fixed effects (ξ_j) in Table 10 and vehicle-model-year fixed effects (ξ_{jt}) in Table A.8, and find the results to be robust across specifications. Although using ξ_{jt} would be the more conventional approach, I prefer ξ_j because it allows greater identification of the role of charging network on EV adoption. With ξ_{jt}, part of the variation associated with Tesla's proprietary charging network would instead be absorbed by the fixed effects, making the role of network quality more difficult to identify.

5.2 Cost Estimation

In this section, I exploit realized private entry decisions together with the model-predicted profits at each node to estimate the fringe firms' fixed entry cost and the VI PCM per Tesla sold required to justify network expansion.

Following Bresnahan and Reiss (1991), entry is characterized by a static equilibrium condition in which the observed number of firms N implies an inequality bracket on per-firm profits and fixed costs. I extend this static framework to a panel setting by exploiting entry timing, comparing pre-entry and entry-year profitability conditions to recover analogous cost bounds from observed $0 \to N$ transitions. Rather than specify a parametric distribution for entry-cost shocks, I estimate a common cost by minimizing a moment-inequality criterion that penalizes the squared magnitude of revealed-preference violations. I then use the estimated entry cost as the charger cost faced by the VI and infer the PCM that rationalizes the observed charging station build-out, given the model-predicted increase in EV sales generated by each additional charging station.

While fringe firms are assumed to operate in a competitive environment given the market

conditions, the VI is modeled as a Stackelberg leader. The VI observes the government's network deployment and chooses whether to build a charging station, taking into account the fringe firms' best-response entry decisions. Solving the model by backward induction allows me to infer the effect of a VI charging station on the firm's EV sales and the PCM required to rationalize the observed deployment decisions.

To ensure the model reproduces the observed market outcome, I calibrate the market size to $M = 15\%$ of the registered passenger-vehicle stock in each region, the value at which simulated EV adoption matches the 2022 EV stock in the data. The same calibration applies everywhere EV counts enter the estimation. As a first validation, I show that this calibration and estimated entry cost yields charger costs consistent with engineering estimates. I also find that it implies PCMs that are consistent with those reported by (Remmy, 2025).[21]

Finally, I validate the baseline model using the estimated demand and cost parameters against the 2022 data, finding that it provides a good fit.

Fringe Entry Cost Estimation

Chargers are sunk once installed, so I identify the common per-charger flow cost from observed entry decisions. Consider a charging node that first transitions from zero chargers in year $t-1$ to $N_r \geq 1$ chargers in year t. Revealed preference implies two inequalities.

Because the firm did not enter in year $t-1$, installing the first charger would not have covered its annual cost,

$$F \;\geq\; R^e_{r,t-1}, \tag{23}$$

where $R^e_{r,t-1}$ denotes the predicted variable profit from entering with one charger in year $t-1$, evaluated holding the rest of the charging network fixed.

Once the node enters in year t with N_r chargers, the observed charger configuration must cover its cost,

$$F \;\leq\; R_{rt}(N_r) = \frac{(p-c)Q_{rt}}{N_r}, \tag{24}$$

where Q_{rt} is the demand predicted by the charging-demand model at the observed network configuration.

Each entry event therefore identifies an interval

$$R^e_{r,t-1} \;\leq\; F \;\leq\; R_{rt}(N_r).$$

Because permitting delays, local construction costs, and other unobserved factors prevent a single common cost from satisfying every interval exactly, I estimate the common per-

[21] (Remmy, 2025) find an average markup of about 6,770 euros, which represents about 8,000 $.

charger cost by minimizing the squared magnitude of the violated inequalities,

$$\widehat{F} = \arg\min_{F \geq 0} \sum_{r \in \mathcal{E}} \left[\max\left(0, R^{e}_{r,t-1} - F\right)\right]^2 + \left[\max(0, F - R_{rt}(N_r))\right]^2, \tag{25}$$

where $\mathcal{E}$ denotes the set of observed $0 \to N$ entry events.

Equation (25) is a moment-inequality analogue of the entry model of Bresnahan and Reiss (1991). Rather than specifying a parametric distribution for entry-cost shocks and estimating an ordered-probit likelihood, it estimates the common per-charger cost that minimizes the total dollar value of revealed-preference violations, leaving the cost-shock distribution unspecified. Standard errors are obtained by bootstrapping entry events.

The estimated entry cost is consistent with the engineering cost of a fast-charging site. A \$150,000 site amortized against $\widehat{F} = \$30,236$, $(2,092)$ per year breaks even in roughly 7.2 years at a 15% discount rate or, over a 5-year horizon, implies an internal rate of return of about 15.3%. The implied annuity factor of 4.96 is close to those used in industry benchmarks, including NREL E-FAST (8 years, 10%), Argonne CHECT (15 years, 12%), and RMI (10 years, 8%).

Vertically Integrated Firm Cost Estimation

The vertically integrated firm (Tesla) is not in the charging business but in the EV business: it deploys chargers to raise demand for its cars through a vertically differentiated product characteristic, namely the charging network. Its network choice therefore trades the charging cost off against the vehicle profit the network generates, and the object identified by its entry behavior is the per-vehicle PCM m, not a charging margin. I take Tesla's per-charger cost technology to be the fringe's, $\widehat{F}$, and evaluate all counterfactuals in the 2022 environment.

For each observed Tesla station i I measure the Tesla adoption it supports in general equilibrium,

$$\Delta x_i \;=\; \mathrm{EV}^{(1)} - \mathrm{EV}_i^{(2)}, \tag{26}$$

where $\mathrm{EV}^{(1)}$ is equilibrium Tesla adoption under the full network and $\mathrm{EV}_i^{(2)}$ re-solves the Big-Push fringe equilibrium with station i removed from the VI network deployment. Δx_i thus nets out fringe re-optimization: it is the Tesla demand that would be lost, after competitors adjust, if station i were not provided by the VI.

Optimality of the observed network requires that each Tesla station's vehicle profit cover its charging cost, $PCM^{VI}\,\Delta x_i \geq N_i\,\widehat{F}$, with the marginal station roughly breaking even. Imposing a common margin and pooling across the network gives the aggregate break-even

estimator

$$\widehat{PCM}^{VI} = \frac{\sum_i N_i \widehat{F}}{\sum_i \Delta x_i} = \frac{\overline{N_i \widehat{F}}}{\overline{\Delta x_i}}, \tag{27}$$

that is, the margin at which Tesla's total charging deployment equals the vehicle profit on the total adoption its network supports. The estimate is $\widehat{PCM}^{VI} = \$10,385$ per induced vehicle, a plausible per-vehicle contribution margin comparable in magnitude to Tesla's reported automotive gross profit per vehicle near the end of the sample period and supported by (Remmy, 2025) estimated average markup of $8,000 using the German EV market.

Table 11: Cost-Side Parameter Estimates

Parameter	Symbol	Estimate	Method
Contribution margin per charging event	$p - c$	\$10	data
Market size normalization	M	15%	calibrated
Per fast-charger flow cost	F	\$30,236	moment inequalities
Tesla per-vehicle margin (PCM)	PCM^{VI}	\$10,385	revealed preference

5.3 Estimation Validation

In order to validate the model performance at matching the data, Table 12 summarizes the 2022 market outcome, which I take to be a steady state, as well as the model predicted outcome.

Table 12: Equilibrium Outcomes: Data versus Big-Push Counterfactuals (2022)

	Data	Low	Laissez-faire High	Big Push
EV adoption				
Total EV stock	57,930	37,963	50,785	56,666
Chargers	1,220	331	348	984
government (Big Push)	741	0	0	741
Private fringe	183	147	228	99
Tesla	296	184	120	144

Low and High are counterfactuals *without* the government Big-Push investment (government chargers = 0); BP adds the government floor. Data columns are observed 2022 values.

I can find that the $M = 15\%$ normalization allows the total EV stock on the data and the Big Push to be in line. One thing to note is the under-deployment of private compatible

charging stations in the Big Push case. On the other hand, in the absence of the government network deployment, the private response is much more robust. I find Tesla to under-deploy the charging network, but that is by construction as I restrict the deployment space to 0 or 8 chargers.

6 Counterfactual

The EV charging market suffers from two distinct market failures. The first is a coordination problem: because a station's profitability depends on EV adoption, which in turn depends on the network other stations provide, private investment is a game of strategic complements with multiple equilibria, and nothing guarantees the industry lands on a good one. The second is a classical externality: even the best private equilibrium underprovides, because a station on a remote corridor generates adoption that accrues to other markets, other firms, or vehicle manufacturers—none of which the local entrant can capture. These failures can be addressed in many ways; this paper focuses on government intervention through direct network deployment. Because charging is a two-sided market, I further ask whether purchase subsidies for EVs complement the deployment—raising its returns by thickening the demand side—or substitute for it, so that the two instruments together deliver less than the sum of their parts. Finally, I study the role of VI by Tesla, which internalizes the network externality its proprietary charging network generates.

The counterfactual of interest is the causal impact of the government network deployment on (i) EV adoption and (ii) crowding out of private investment. Because the entry game is one of weak strategic complementarities, many equilibria may be sustained. To facilitate interpretation, I reduce this dimensionality to the Low and High equilibria attainable without government intervention. Each policy environment therefore delivers three market outcomes: Low, High, and the Big Push.

Moving from Low to the Big Push measures the total effect of the Big Push. This decomposes into two parts: the move from Low to High is the coordination component, capturing what better coordination alone could achieve, while the move from High to the Big Push is the externality component, capturing the deployment that could not be sustained privately even under full coordination.

Each counterfactual is defined by three binary levers: whether an EV purchase subsidy is in place, whether the government deploys the network (Big Push), and whether Tesla invests in charging infrastructure. Table 13 summarizes the resulting design.

The three blocks map directly onto the questions asked in Section 6. CF 1.1 identifies the effect of the deployment itself. CF 1.2 repeats the exercise with a purchase subsidy in place,

Table 13: Counterfactual design

	Subsidy	Big Push	Tesla VI
CF 1.1: The causal impact of the Big Push			
Benchmark	0	0	1
Big Push	0	1	1
CF 1.2: Do subsidies complement or substitute for the Big Push?			
Benchmark	1	0	1
Big Push	1	1	1
CF 2: The value of VI			
Benchmark	0	0	0
Big Push	0	1	0

Notes: Each row is a market environment; entries indicate whether the lever is active (1) or not (0). Within each environment, outcomes are computed at the Low and High equilibria of the private entry game. CF 1.1 and CF 1.2 are compared to isolate the interaction between demand-side subsidies and supply-side deployment. CF 2 is contrasted with CF 1.1, which shares the same subsidy and Big Push settings but allows Tesla to invest, so that the difference isolates the value of VI.

so that comparing the two reveals whether the instruments are complements or substitutes. CF 2 shuts down Tesla's investment and, read against CF 1.1, measures how much of the network a vertically integrated firm internalizes on its own.

6.1 Big Push & Subsidy

The Causal Impact of the Big Push

I compare the Big Push equilibrium against the outcomes that obtain under laissez-faire. Because the entry game has multiple equilibria, laissez-faire is not a single state but two: a low- and a high-adoption equilibrium. I therefore evaluate the Big Push against both, which decomposes its total effect into two parts. The move from the Low to the High equilibrium is the coordination component—what better equilibrium selection alone would achieve—and the move from the High equilibrium to the Big Push is the externality component—the additional network that no private configuration sustains, even under full coordination. To recover the Low and High equilibria I follow Vives (2005), whose algorithm returns the Low and High equilibria of the game with strategic complementarities.[22]

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

[22]The algorithm iterates the best-response map from an extremal starting point: from no charging stations at any node for the Low equilibrium, and from a station at every node for the High.

High- rather than the Low-adoption Nash equilibrium. EV adoption rises by 31.3% (from 23,785 to 31,233 vehicles) as the network shifts to the High-adoption state, improving network quality by 0.24 (worth about $5,333 per vehicle) and raising consumer surplus by $166.5M. Climate benefits add $41.3M at an SCC of $185, and manufacturer surplus increases by $68.8M. Notably, the total charger count barely moves (249 to 236), but the number of active stations more than doubles (35 to 77): coordination reallocates capacity from a few deep sites to a broader set of locations. 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 government first-mover investment once the market is already in the High-adoption equilibrium. Relative to the High equilibrium, the Big Push raises EV adoption by a further 27.1% (to 39,691 vehicles) and improves network quality by 0.36 (about $8,000 per vehicle) and raising consumer surplus by $317.5M. Climate benefits rise by $46.9M (SCC = $185; $12.7M at $50) and manufacturer surplus by $75.1M. These gains come from government provision at high-connectivity nodes that are socially valuable but privately unprofitable: the 741-charger government network raises the number of chargers from 236 to 921 while crowding out part of the private fringe, whose chargers fall from 132 to 52.

Together, the two counterfactual experiments separate equilibrium selection from network externalities. The coordination component captures extensive-margin gains from moving from Low to High adoption (+31.3%), while the externality component captures gains from improving network topology within the High state by internalizing unpriced spillovers at connector nodes (+27.1%). The manufacturer-surplus and climate components are additive across the two channels, while the adoption effects compound multiplicatively ($1.31 \times 1.27 = 1.67$), yielding the total Big Push effect relative to Low: a 67% increase in EV adoption, $88.2M in climate benefits (SCC = $185), and $143.9M in manufacturer surplus. Unlike laissez-faire, the Big Push requires government outlays ($22.4M for 741 government chargers at $30,236 each) but delivers substantially larger consumer, producer, and climate benefits.

Table 14: CF 1.1: The causal impact of the Big Push

Panel A. Decomposition of the Big Push effect

	Coordination	**Externality**	**Total**
EV adoption (Δ)	+7,448	+8,458	+15,906
relative (×)	1.31	1.27	1.67
Network quality (Δ)	+0.24	+0.36	+0.60
WTP per vehicle ($)	+5,333	+2,667	+8,000
Consumer surplus ($M)	166.6	150.9	317.5
Manufacturer surplus ($M)	68.8	75.1	143.9
Tesla	33.5	35.2	68.7
non-Tesla	35.3	39.9	75.2
Climate benefit ($M)	41.3	46.9	88.2

Panel B. Equilibrium states (levels)

	Low	**High**	**Big Push**
EV adoption	23,785	31,233	39,691
Total chargers	249	236	921
Government	–	–	741
Private fringe	81	132	52
Tesla	168	104	128
Active charging stations	35	77	359

Government cost: $22.4M (741 chargers at $30,236)

Note: Panel A reports levels at each equilibrium of the entry game; under laissez-faire two stable equilibria emerge (Low and High), and the Big Push adds the 741-charger government network. Panel B decomposes the total Big Push effect (relative to the Low equilibrium) into a coordination component (Low to High) and an externality component (High to Big Push); absolute effects sum to the total and relative effects multiply to it. Chargers and stations are distinct: a station is an active site that may host several chargers. Consumer surplus, network quality, and per-vehicle willingness to pay are measured relative to the Low equilibrium. Manufacturer surplus is the price–cost margin times EVs sold, at $8k for non-Tesla makes (Remmy, 2025) and a $10,400 Tesla margin net of charger costs. Climate benefits value each EV at 30 tCO_2 lifetime abatement.

While Table 14 quantifies the overall effectiveness of the Big Push, it does not explain what makes the policy effective. 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 5 summarizes this spatial placement.

The left panel plots cumulative EV adoption against distance to Montréal, for the Low (blue), High (orange), and Big Push (green) equilibria. Within the first 150 km, adoption differs little across the three equilibria; beyond 150 km, the three curves diverge monotonically. The right panel provides the explanation: it plots the cumulative number of fast chargers against the same distance axis. The Big Push nearly quadruples the size of the charging network relative to the High equilibrium (from 236 to 921 chargers, and from 77 to 359 active stations), and aggregate EV adoption rises by 27.1%. This aggregate figure masks substantial heterogeneity: the relative increase in EV adoption is far larger in remote areas, where the government network fills gaps the private sector leaves unserved. The range composition of adopters shifts accordingly, with the share of short-range (150–200 km) vehicles falling from 10.1% to 7.0% and the share of long-range vehicles rising, consistent with the Big Push expanding coverage into corridors where range anxiety previously bound.

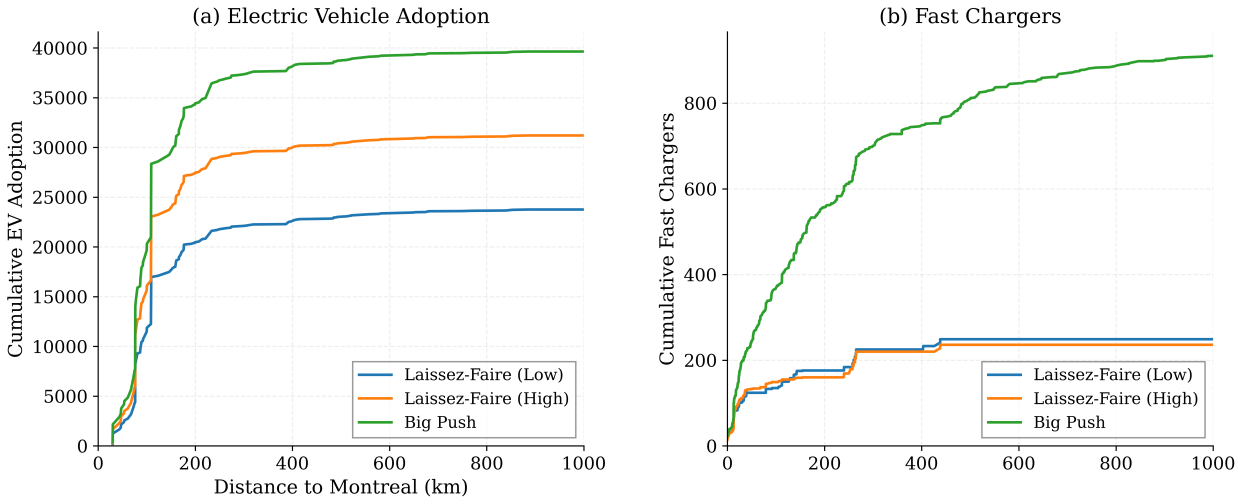
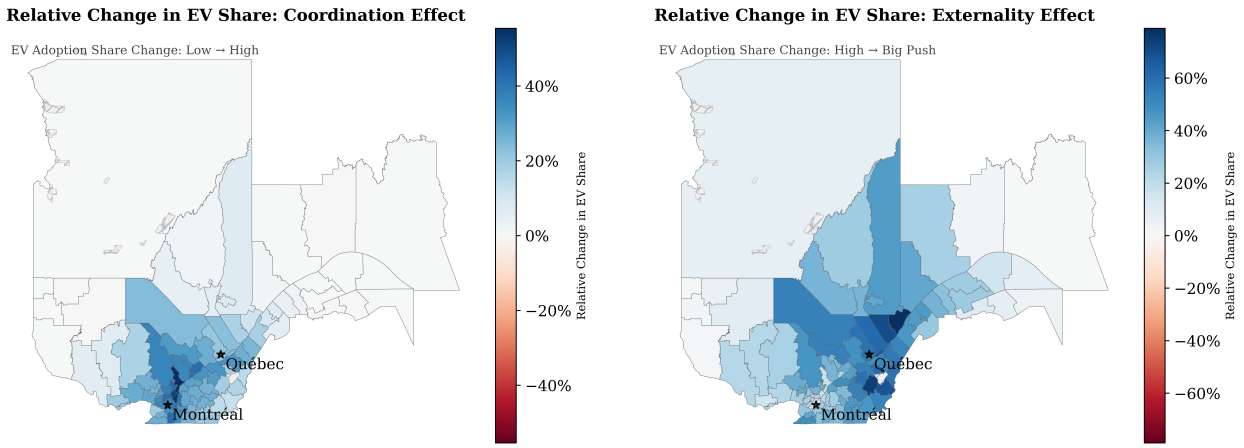


Figure 5: Cumulative EV Adoption and Charger Deployment by Distance to Montréal

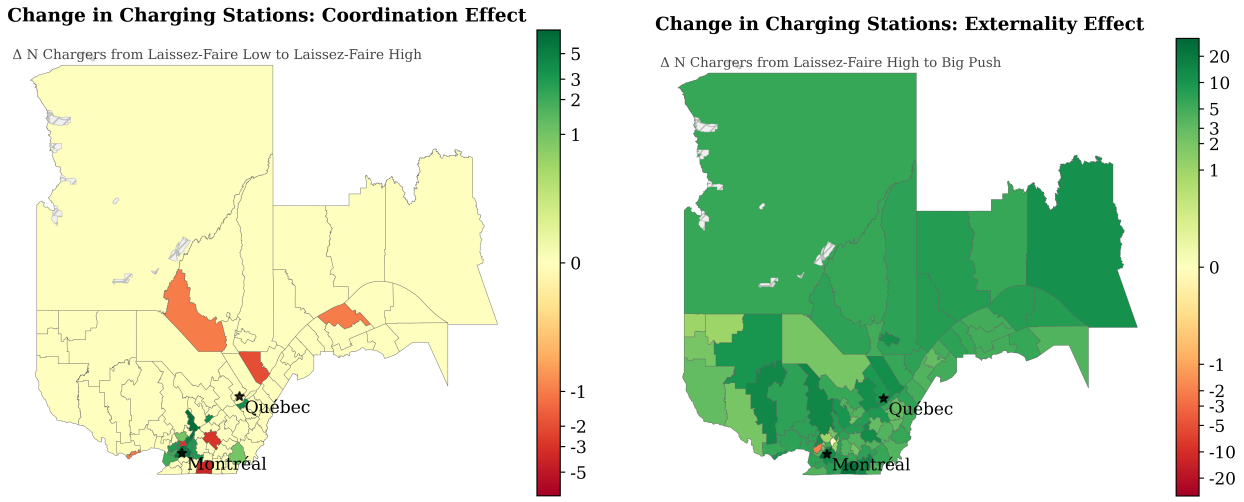
Figure 6 makes this spatial heterogeneity explicit by mapping the relative change in EV adoption share for each component of the decomposition. The coordination effect (left panel) is concentrated in the Montréal–Québec corridor: moving from the Low to the High equilibrium raises adoption most where the private network densifies. The externality effect (right panel) displays the opposite geography: the largest relative gains occur in peripheral regions — along the St. Lawrence east of Québec and in the mid-northern regions — precisely where the government network adds connector nodes that no private player would build. In the Montréal core, by contrast, the externality effect is close to zero, since the private network already serves these markets in the High equilibrium.



(a) Coordination effect (Low → High) (b) Externality effect (High → Big Push)

Figure 6: Relative Change in EV Adoption Share by Region

To understand the network deployment underlying these changes in EV market share, Figure 7 maps the corresponding spatial changes in the charging network for each component of the decomposition.



(a) Coordination effect (Low → High) (b) Externality effect (High → Big Push)

Figure 7: Change in the Number of Fast Chargers by Region

The two panels reveal that the coordination and externality components operate through very different deployment patterns. The coordination effect (left panel) is a reallocation rather than an expansion: the total network grows only modestly (from 331 to 348 chargers), but its composition shifts markedly, with the private fringe adding stations in the dense Montréal–Centre-du-Québec corridor while Tesla retreats (its chargers fall from 184 to 120) and a handful of peripheral locations lose coverage. Consistent with the countercyclical

investment behavior documented above, Tesla's proprietary network partially substitutes for the missing compatible network in the Low equilibrium and contracts once the fringe coordinates on the High state. The externality effect (right panel) is, by contrast, an almost uniform expansion: the 741 government chargers extend coverage across virtually every region of the province, including the remote north and the Lower St. Lawrence, with crowd-out of the private fringe confined to a few cells around Montréal. This mirrors the adoption geography in Figure 6: the coordination component densifies the core, while the externality component extends the network's reach to regions the private sector leaves unserved under any equilibrium selection.

Such network deployment also affects firms' entry decisions. Figure 8 shows how the private fringe responds to both (i) government investment and (ii) the induced increase in EV adoption. Green indicates crowd-in (more private entry) and red indicates crowd-out (less private entry).

While I find an aggregate increase in the number of charging stations (see Figure 5), the government investment induces some crowding in and out. The pattern is spatially differentiated: private entry is crowded out in the dense Montréal–Québec corridor, where government provision substitutes directly for fringe stations and depresses the returns to private entry, while it is crowded in at select remote nodes, where the larger induced EV fleet raises utilization and makes previously marginal locations profitable. Overall, government investment displaces private entry in the network core, but complements it at the periphery through the demand channel of the complementary market (the EV market).

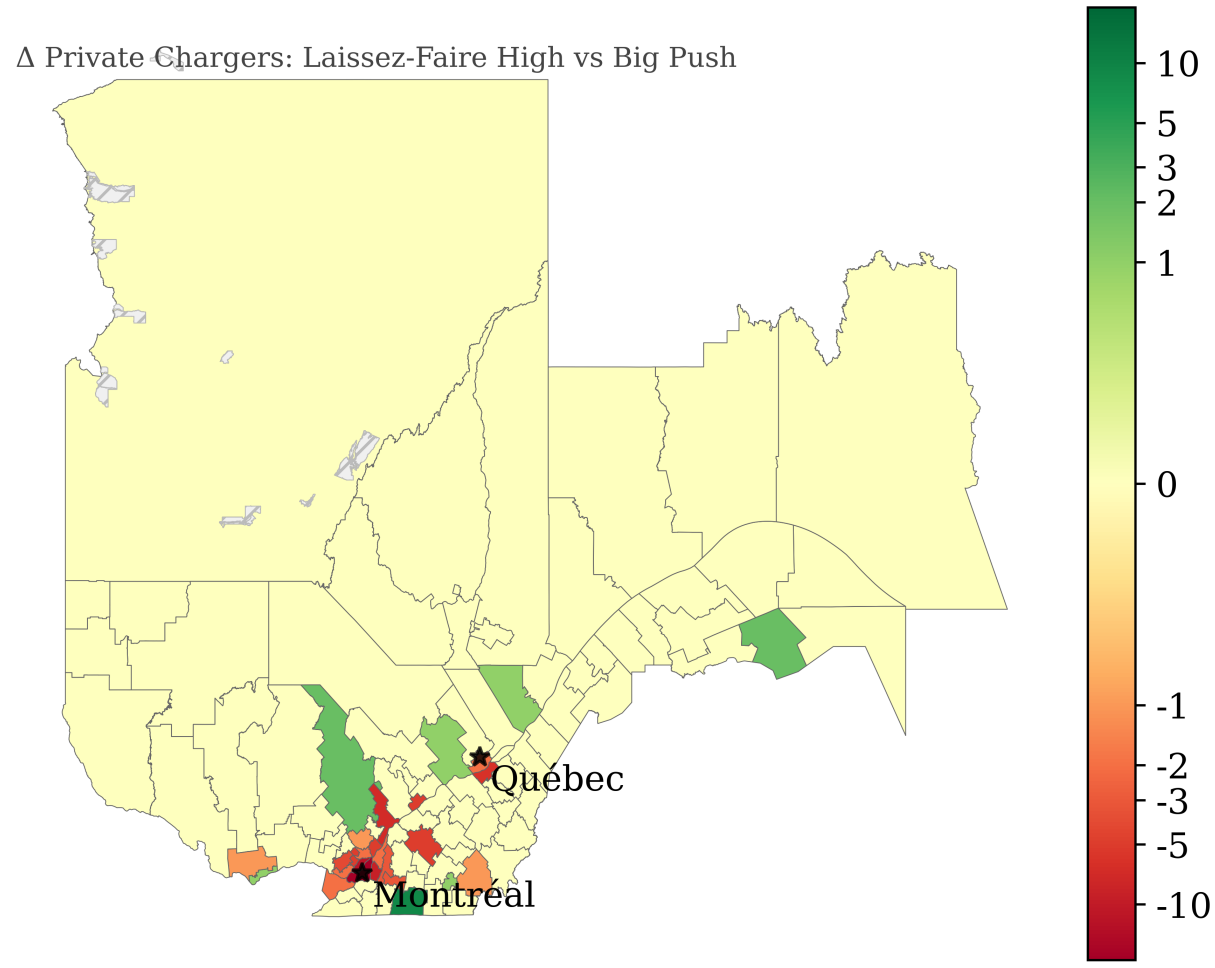


Figure 8: Crowd-In vs. Crowd-Out of Private Chargers

Do subsidies complement or substitute for the Big Push?

Comparing the realized deployment under the observed subsidy regime with a world in which the subsidy is in place but the Big Push is not, I ask whether the two instruments complement or substitute one another. The answer depends on the margin.

In absolute terms, they are complements. Holding the deployment fixed, the Big Push converts 20,386 additional EVs when the subsidy is active, against 15,906 without it: the levels is positive, and the same deployment network add roughly 4,500 more vehicles once the demand side has been thickened by the EV subsidies. In relative terms, the ranking reverses. The Big Push multiplies adoption by 1.67 absent of the subsidy but only by 1.55 with it. The two are not in tension: the subsidy raises the baseline enough that a larger absolute gain is a smaller proportional one. Read together, the two margins deliver the

two-sided-market message. Each additional station converts proportionally fewer marginal drivers once the subsidy has already pulled in the most responsive consumers, even as total conversions rise.

This aligns with Springel (2021) on two margins. First, it confirms the non-neutrality of subsidies: instruments targeting the two sides of the market are not interchangeable, and their relative effectiveness depends on which side is already thick. Second, it reproduces the diminishing marginal value of charging stations, visible here in the externality component, which falls from 8,458 to 6,940 additional vehicles once the subsidy is in place.

I bring a novel result: subsidies do not resolve the coordination failure but widen it. Demand-side support raises the Low–High gap from 7,448 to 13,446 vehicles, so the cost of landing on the wrong equilibrium—and equivalently the value of a mechanism that selects the right one—rises with the subsidy rather than falling.

Table 15: Big Push × Subsidy Interaction

	No subsidy	**With subsidy**	**Δ**
Panel A. EV adoption, levels			
Low	23,785	36,879	+13,094
High	31,233	50,325	+19,092
Big Push	39,691	57,265	+17,574
Panel B. Decomposition, absolute (EVs)			
Coordination	7,448	13,446	+5,998
Externality	8,458	6,940	−1,518
Total	15,906	20,386	+4,480
Panel C. Decomposition, ratio (×)			
Coordination	1.313	1.365	+0.052
Externality	1.271	1.138	−0.133
Total	1.669	1.553	−0.116
Panel D. Network quality, total effect			
Network quality	0.600	0.535	−0.065
WTP per vehicle ($)	13,479	12,019	−1,460
Aggregate WTP ($M)	535.0	688.3	+153.3

Note: Columns differ only in whether an EV purchase subsidy is in place; the subsidy is assumed to pass through fully to consumers. The government network is identical in both Big Push columns (741 chargers, $22.4M). Panels B and C decompose the total Big Push effect into a coordination component (Low to High) and an externality component (High to Big Push); absolute components sum to the total, ratio components multiply to it. The total cross-partial in levels is positive (+4,480 EVs; +$153.3M in willingness to pay), so the instruments are complements in levels. It is negative in ratios (−0.116), since the subsidy raises the no-deployment baseline by enough that the larger absolute gain is a smaller proportional one. The decomposition locates the source of both patterns: the subsidy raises the coordination component on both margins (+5,998 EVs; +0.052) while lowering the externality component on both (−1,518 EVs; −0.133).

The comparison with Table 15 is informative because the two tables share an identical government network (741 chargers, $22.4M) and differ only in whether the demand side is subsidized. To further understand how the subsidy reshapes the entry game itself—rather than only the size of the Big Push effect—Table 16 reruns the full decomposition of CF 1.1 with the purchase subsidy in place.

Three differences stand out. First, the subsidy shifts weight from the externality compo-

nent toward the coordination component. Without the subsidy the two channels are roughly balanced (1.31×1.27); with it, coordination does most of the work and the externality margin shrinks by half (1.36×1.14). A thicker demand side makes private entry viable at nodes that were marginal before, so more of the socially valuable network is sustainable in a well-coordinated laissez-faire equilibrium and less of it must be supplied by the government. The corollary is that the subsidy does not resolve the coordination failure—it enlarges it. Even with generous purchase incentives, the gap between the Low and High equilibria widens, so the return to whatever mechanism selects the good equilibrium is higher, not lower, once demand-side support is in place. Demand-side policy and coordination are therefore not substitutes: a subsidy raises the stakes of the selection problem it leaves unsolved.

Second, the private sector is far more active in every laissez-faire column. The Low equilibrium already supports 76 active stations against 35 without the subsidy, and the High equilibrium reaches 141 against 77: the subsidy is, in effect, an indirect station subsidy, since a thicker demand side makes entry profitable at sites that were marginal before. The Big Push reverses this. Private stations fall from 126 under High to 30, a sharper contraction than the corresponding fall in CF 1.1, because the subsidy had drawn private entry into precisely the high-connectivity connector nodes the government network then occupies.

The distributional reading is therefore mixed. Consumers and manufacturers do best under the subsidized Big Push, which delivers the largest fleet (57,265 vehicles), the highest consumer surplus ($688.3M), and the highest manufacturer surplus ($181.8M) of any column considered. Private ChargePoint operators do best under the subsidized High equilibrium, where a subsidized demand side supports 126 stations with no competitor, and worst under the subsidized Big Push, where they retain only 30.

Third, welfare gains are uniformly larger in levels while the multiplicative effects are smaller. Consumer surplus, manufacturer surplus, and climate benefits all exceed their CF 1.1 counterparts for the same $22.4M of government investment, even as adoption rises by 55.3% rather than 66.9%. This is the level–ratio distinction of the previous subsection expressed across the whole decomposition: the subsidy raises the baseline enough that a larger absolute gain registers as a smaller proportional one.

Table 16: CF 1.2: The causal impact of the Big Push (with subsidy)

Panel A. Decomposition of the Big Push effect

	Coordination	**Externality**	**Total**
EV adoption (Δ)	+13,446	+6,940	+20,386
relative (×)	1.36	1.14	1.55
Network quality (Δ)	+0.31	+0.23	+0.54
WTP per vehicle ($)	+6,871	+5,149	+12,020
Consumer surplus ($M)	345.8	294.9	640.6
Manufacturer surplus ($M)	121.4	60.4	181.8
Tesla	53.8	24.5	78.3
non-Tesla	67.6	35.9	103.5
Climate benefit ($M)	74.6	38.5	113.1

Panel B. Equilibrium states (levels)

	Low	**High**	**Big Push**
EV adoption	36,879	50,325	57,265
Total chargers	325	344	996
Government	–	–	741
Private fringe	141	224	103
Tesla	184	120	152
Active charging stations	53	126	374

Government cost: $22.4M (741 chargers at $30,236)

Note: Panel B reports levels at each equilibrium of the entry game; under laissez-faire two stable equilibria emerge (Low and High), and the Big Push adds the 741-charger government network. Panel A decomposes the total Big Push effect (relative to the Low equilibrium) into a coordination component (Low to High) and an externality component (High to Big Push); absolute effects sum to the total and relative effects multiply to it. Chargers and stations are distinct: a station is an active site that may host several chargers. Consumer surplus, network quality, and per-vehicle willingness to pay are measured relative to the Low equilibrium, with each stage's quality gain valued at the fleet present after that stage. Manufacturer surplus is the price–cost margin times EVs sold, at $8k for non-Tesla makes (Remmy, 2025) and a $10,385 Tesla margin net of charger costs. Climate benefits value each EV at 30 tCO_2 lifetime abatement at a $185 social cost of carbon.

6.2 The Value of VI

To isolate the role of Tesla's proprietary network, I re-solve the model with Tesla charger investment switched off: Tesla sells vehicles but cannot build its own network. The comparison with the baseline of CF 1.1 measures how much the vertically integrated firm's network shapes the adoption equilibrium.

Two patterns stand out. First, the vertically integrated firm's network matters most precisely when the market is stuck in the Low equilibrium: banning Tesla charger investment lowers Low-equilibrium adoption by 2,455 vehicles (−10.3%) and cuts Tesla's share by 5.0 %age points. The effect shrinks once coordination improves—adoption falls by only 1,122 vehicles (−3.6%) in the High equilibrium and by 1,015 (−2.6%) under the Big Push, with Tesla's share down 1.7 and 1.3 %age points respectively—because the compatible network is by then dense enough that Tesla's proprietary network is partly redundant.

Table 17: Effect of Tesla's Vertically Integrated Network on Adoption and Market Share

	Total EV Adoption			**Tesla Market Share**		
	Tesla	No Tesla net.	Δ	Tesla	No Tesla net.	Δ (pp)
Low (L)	23,785	21,330	−2,455	39.5%	34.5%	−5.0
High (H)	31,233	30,111	−1,122	39.8%	38.0%	−1.7
Big Push	39,691	38,676	−1,015	40.0%	38.7%	−1.3

Note: "Tesla" lets Tesla operate its proprietary, incompatible network; "No Tesla net." bans Tesla charger investment (mandated compatibility). Δ is the change from baseline to the no-network counterfactual; market-share changes are in %age points.

This also shows up in the decomposition: without Tesla's network to prop up the laissez-faire Low state, the coordination component is larger (+41.2% vs. +31.3% with Tesla present), because the market is worse off when it fails to coordinate. The vertically integrated network thus matters mainly as insurance against the worst case: it cushions the Low equilibrium, while its effect in the High equilibrium and under the Big Push is roughly a third as large. Notably, the network raises adoption under every equilibrium selection: incompatibility is privately advantageous to Tesla, yet adoption-enhancing for the market as a whole.

Table 18: CF 2: The value of VI

Panel A. Decomposition of the Big Push effect

	Coordination	**Externality**	**Total**
EV adoption (Δ)	+8,781	+8,565	+17,346
relative (×)	1.41	1.28	1.81
Network quality (Δ)	+0.28	+0.37	+0.65
WTP per vehicle ($)	+6,253	+8,410	+14,663
Consumer surplus ($M)	188.3	325.3	513.5
Manufacturer surplus ($M)	80.0	76.9	156.9
Tesla	42.5	36.5	79.0
non-Tesla	37.5	40.4	77.9
Climate benefit ($M)	48.7	47.5	96.3

Panel B. Equilibrium states (levels)

	Low	**High**	**Big Push**
EV adoption	21,330	30,111	38,676
Total chargers	80	131	790
Government	–	–	741
Private fringe	80	131	49
Tesla	–	–	–
Active charging stations	29	76	356

Government cost: $22.4M (741 chargers at $30,236)

Note: Tesla cannot invest in chargers (mandated compatibility), so it builds no network in any column; its manufacturer surplus is therefore gross of charger costs. Panel B reports equilibrium levels; under laissez-faire two stable equilibria emerge (Low and High), and the Big Push adds the same 741-charger government network. Panel A decomposes the total effect (relative to Low) into a coordination component (Low to High) and an externality component (High to Big Push); absolute effects sum to the total and relative effects multiply to it. Consumer surplus, network quality, and per-vehicle willingness to pay are measured relative to Low, each stage's quality gain valued at the fleet present after that stage. Manufacturer surplus is the margin times EVs sold ($8k non-Tesla, $10,385 Tesla). Climate benefits value each EV at 30 tCO_2 lifetime abatement at a $185 social cost of carbon. Government investment is 741 chargers at $30,236 each.

Figure 9 makes the mechanism visible by overlaying the no-Tesla counterfactual (dashed) on the baseline equilibria (solid). Panel (b) shows that removing Tesla's network is most consequential in the Low equilibrium: the network shrinks from 249 to 80 chargers, and the

remaining fringe capacity is concentrated near Montréal — the dashed curve is essentially flat beyond 150 km. In the Low state, Tesla's proprietary chargers provide nearly all long-distance coverage. In the High equilibrium and under the Big Push, by contrast, the dashed curves run parallel to the solid ones: the network is smaller without Tesla, but its spatial reach is intact, because the compatible network (private fringe in High, government chargers in the Big Push) already spans the province. Panel (a) shows the adoption consequence. The Low equilibrium displays a wide gap that persists across the entire distance range, while the High and Big Push curves are nearly coincident with their no-Tesla counterparts. Tesla's network is thus pivotal only where the compatible network fails to reach — which, outside the Low equilibrium, is almost nowhere.

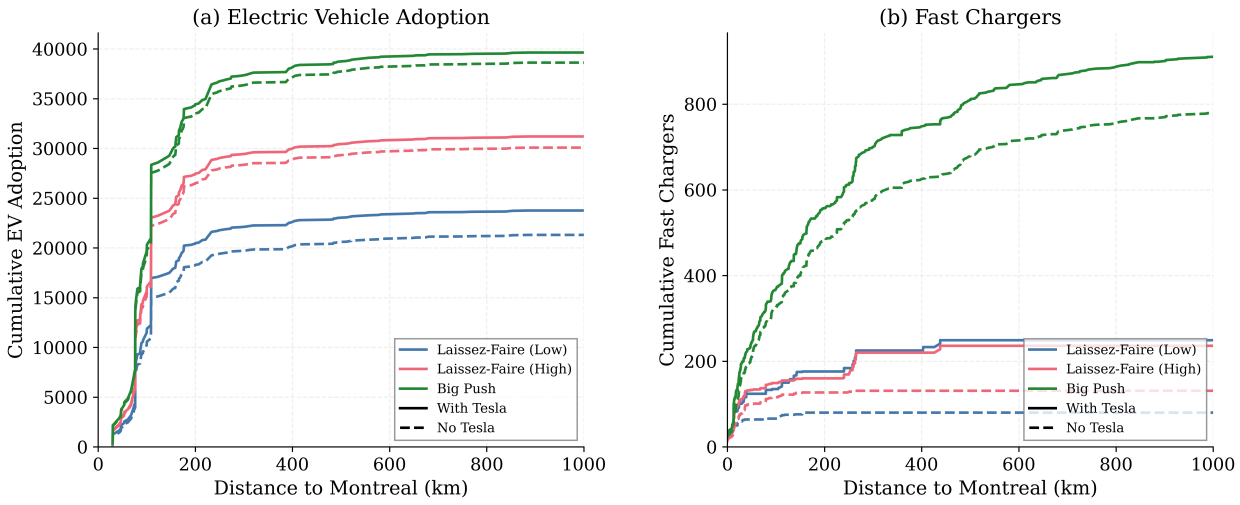


Figure 9: Cumulative EV Adoption and Charger Deployment by Distance to Montréal, With and Without Tesla's Proprietary Network

While I find a result in the opposite direction from Li (2019), the two are not a one-to-one contradiction, because the counterfactual and modelling assumption differ. I shut down the VI firm altogether, whereas Li studies how mandated compatibility would affect investment by competing oligopolistic VI firms. That is, she asks whether business stealing or strategic complementarity dominates. Put differently, Li asks whether mandated compatibility is a good policy, while I show that prohibiting the VI firm from investing in the network would lower EV sales.

The modelling and counterfactual differences reside in the research question and in the market studied. Competing VI firms are not a feature of the Québec market, nor of the rest of Canada, so I focus instead on how the VI firm responds to the Big Push, and on how the VI firm internalizes part of the uninternalized externality, which I then contrast with the government's Big Push.

Summary Table

Table 19 collects the three counterfactuals against a common benchmark, in which Tesla invests, no purchase subsidy is in place, and the Big Push adds the realized 741-charger government network. Three results carry the paper's message.

The Big Push is effective in every environment, raising EV adoption by between 55 and 81 % over the Low equilibrium, and its two channels contribute in comparable measure throughout: coordination—the move from the Low to the High equilibrium—and externalities—the move from the High equilibrium to the Big Push—each account for roughly half of the total effect. The deployment works not by building a larger network than a coordinated market would, but by supplying capacity at nodes that no private firm would serve.

Demand-side subsidies (CF 1.2) raise adoption at every equilibrium but lower the proportional return to the Big Push, from 1.67 to 1.55. The two instruments are complements in levels—a fixed government network converts more vehicles once demand is subsidized—and substitutes in ratios, because the subsidy raises the no-deployment baseline enough that a larger absolute gain registers as a smaller share of it. The same force widens the gap between the Low and High equilibria, so the value of coordinating the market rises with the subsidy rather than falling: demand-side support raises the stakes of the selection problem it leaves unsolved.

Finally, Tesla's proprietary network (shut down in CF 2) matters most precisely where the market fares worst. Removing it costs the most adoption in the Low equilibrium and far less once coordination improves, because the compatible network is by then dense enough to make Tesla's chargers partly redundant. VI thus functions as insurance against the coordination failure rather than a remedy for it: by internalizing the externality its own chargers generate, the integrated firm sustains investment that props up the worst-case equilibrium, but incompatibility confines the benefit to its own customers.

Table 19: Summary of counterfactuals

	CF 1.1 No subsidy	**CF 1.2** Subsidy	**CF 2** No Tesla VI
EV adoption (vehicles)			
Low	23,785	36,879	21,330
Big Push	39,691	57,265	38,676
Big Push effectiveness (×, relative to Low)			
Coordination (H/L)	1.31	1.36	1.41
Externality (BP/H)	1.27	1.14	1.28
Total (BP/L)	1.67	1.55	1.81
Chargers, Low → Big Push			
Government	0 → 741	0 → 741	0 → 741
Private fringe	81 → 52	141 → 103	80 → 49
Tesla	168 → 128	184 → 152	–
Total	249 → 921	325 → 996	80 → 790
Consumer surplus ($M, BP vs. Low)			
Total	317.5	688.3	567.1

Note: Each column summarizes one counterfactual; full results appear in Tables 14, 16, and 18. All three share the same 741-charger government network under the Big Push. CF 1.1 is the benchmark (Tesla invests, no subsidy); CF 1.2 adds the observed EV purchase subsidy; CF 2 bans Tesla charger investment (mandated compatibility). The Big Push effect decomposes multiplicatively into a coordination component (Low to High) and an externality component (High to Big Push). Charger entries show the count in the Low equilibrium and under the Big Push; a station may host several chargers.

7 Conclusion

This paper asks whether direct government provision of network infrastructure can resolve the coordination failure and network externality. Using vehicle registration and charging transaction data from Québec, I estimate a structural model of EV demand in which consumers value the charging network through the origin–destination trips it supports, combined with a Stackelberg model of station investment in which the government moves first, the vertically integrated firm second, and the private fringe last.

Under laissez-faire, two stable equilibria emerge. Québec's realized Big Push moves the market past both, raising EV adoption by 67% for $22.4M of government investment—an

implied cost of roughly $47 per ton of CO_2 abated, well below the $185 social cost of carbon. The decomposition shows why the policy works: coordination and externalities contribute in roughly equal measure. Coordination barely changes the size of the network but reallocates it; the externality channel expands it fourfold at high-connectivity nodes that no private entrant would build. Private investment is reshaped rather than displaced—crowded out in the dense core, crowded in at the periphery through the demand channel.

Two further results speak to policy design. Demand-side subsidies raise the absolute return to the deployment but widen the gap between equilibria, so the value of coordinating the market rises with the subsidy rather than falling. And Tesla's proprietary network, while privately motivated, raises adoption under every equilibrium selection. Its value is concentrated where the market fares worst: by internalizing the externality its own chargers generate, the integrated firm sustains investment that no fringe entrant would undertake, cushioning the Low equilibrium in particular. VI thus acts as insurance against the coordination failure rather than a remedy for it—incompatibility confines the benefit to the integrated firm's own customers, leaving the network available to everyone else unchanged.

The policy implication is that industrial policy in network markets should be spatially and topologically targeted: what matters is not how much capacity is built but where. Jurisdictions with government ownership can provide directly; decentralized settings can approximate the same incentives through geographically differentiated subsidies or temporary exclusivity rights. More generally, industrial policy can be climate policy. Many decarbonization technologies—hydrogen refueling, district heating, the grid investments behind electrification—are network goods whose value hinges on complementary investments that no private actor can fully coordinate or capture. Where adoption is stalled by such failures, carbon pricing and purchase subsidies are incomplete instruments: they raise the payoff to the high-adoption equilibrium but need not select it. Building the network does both—it substitutes for the coordination the market cannot achieve and supplies the nodes that even a coordinated market would leave unbuilt.

Four extensions suggest themselves. The dynamic incentives to invest strategically so as to influence the equilibrium number of standards remain unexplored. The model abstracts from charging-station pricing, as price variation is weak over the study period. The framework could be used to characterize the minimal network sufficient to overcome the coordination failure—a question with direct bearing on corridor programs such as the U.S. NEVI and NRCan's Trans-Canada initiatives, which impose fixed spacing requirements without reference to network topology. Finally, this paper assumes home charging is available and focuses on the intercity network. Households without a driveway or garage must instead rely on local municipal infrastructure—on-street chargers and neighbourhood hubs—for routine

overnight charging, so their binding constraint is local rather than corridor coverage.

References

Aguirregabiria, V. and C.-Y. Ho (2012). A dynamic oligopoly game of the us airline industry: Estimation and policy experiments. *Journal of Econometrics 168*(1), 156–173.

Barwick, P. J., H.-S. Kwon, S. Li, and N. B. Zahur (2025). Drive down the cost: Learning by doing and government policies in the global ev battery industry. Technical report, National Bureau of Economic Research.

Berry, S., J. Levinsohn, and A. Pakes (2004). Differentiated products demand systems from a combination of micro and macro data: The new car market. *Journal of political Economy 112*(1), 68–105.

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 cashback : 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.

Cole, C., M. Droste, C. Knittel, S. Li, and J. H. Stock (2023). Policies for electrifying the light-duty vehicle fleet in the united states. In *AEA Papers and Proceedings*, Volume 113, pp. 316–322. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

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.

Goldberg, P. K. (1995). Product differentiation and oligopoly in international markets: The case of the us automobile industry. *Econometrica: Journal of the Econometric Society*, 891–951.

Houde, J.-F. (2012). Spatial differentiation and vertical mergers in retail markets for gasoline. *American Economic Review 102*(5), 2147–2182.

Jia, P. (2008). What happens when wal-mart comes to town: An empirical analysis of the discount retailing industry. *Econometrica 76*(6), 1263–1316.

Juhász, R. (2018). Temporary protection and technology adoption: Evidence from the napoleonic blockade. *American Economic Review 108*(11), 3339–3376.

Kline, P. and E. Moretti (2014). Local economic development, agglomeration economies, and the big push: 100 years of evidence from the tennessee valley authority. *The Quarterly journal of economics 129*(1), 275–331.

Lane, N. (2025). Manufacturing revolutions: Industrial policy and industrialization in south korea. *The Quarterly Journal of Economics 140*(3), 1683–1741.

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

Liu, E. (2019). Industrial policies in production networks. *The Quarterly journal of economics 134*(4), 1883–1948.

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.

Murphy, K. M., A. Shleifer, and R. W. Vishny (1989). Industrialization and the big push. *Journal of political economy 97*(5), 1003–1026.

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

Petrin, A. (2002). Quantifying the benefits of new products: The case of the minivan. *Journal of political Economy 110*(4), 705–729.

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

Remmy, K. et al. (2022). Adjustable product attributes, indirect network effects, and subsidy design: The case of electric vehicles. *University of Bonn and University of Mannheim, Germany*.

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.

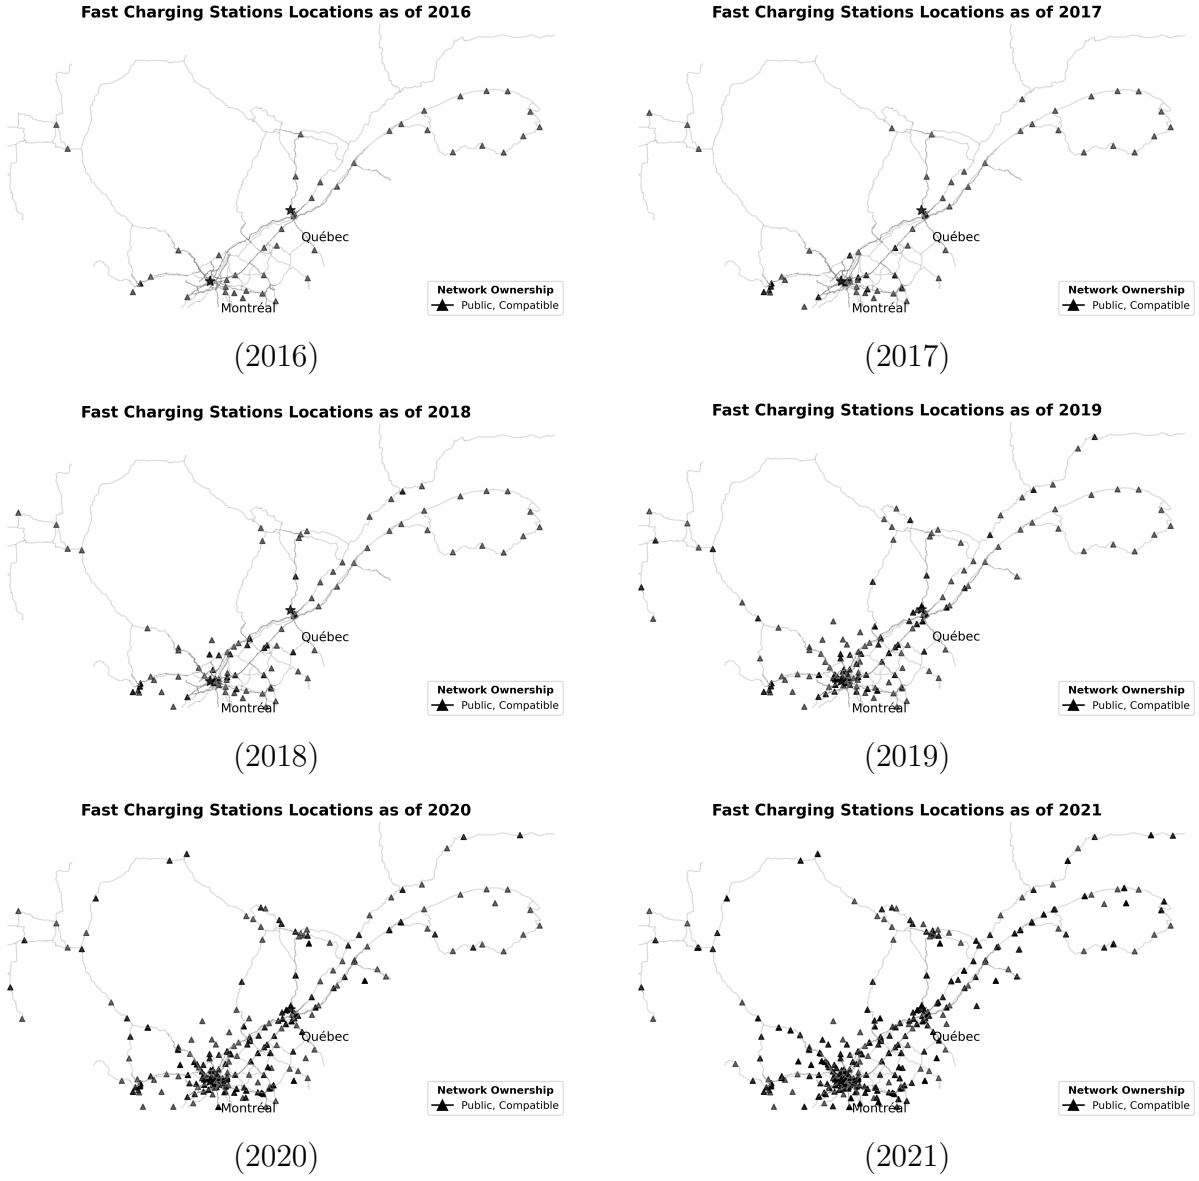
Tarski, A. (1955). A lattice-theoretical fixpoint theorem and its applications.

Vives, X. (1990). Nash equilibrium with strategic complementarities. *Journal of Mathematical Economics 19*(3), 305–321.

Vives, X. (2005). Complementarities and games: New developments. *Journal of Economic Literature 43*(2), 437–479.

A.1 Network Evolution

Figure A.1: Network Evolution



A.2 Additional Empirical Evidence

A.2.1 SOC projections

The left panel of Figure A.2 plots the distribution of the 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), the other to charges implying a moderate arrival buffer (return events). The separation is consistent with drivers targeting a charging onset around 25% SOC. The right panel reports the corresponding distribution conditional on classifying the event as outbound. Section A.3 details the classification algorithm and further validation.

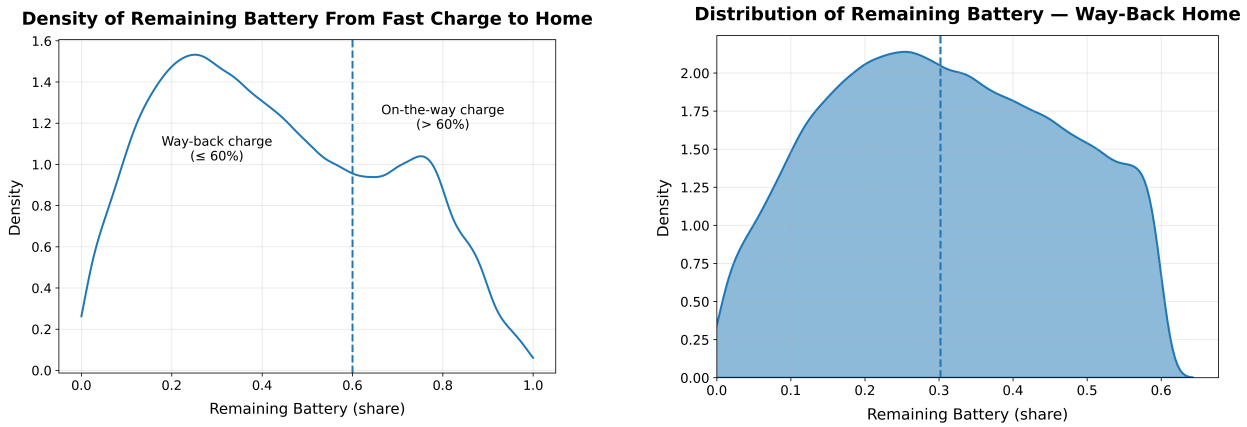


Figure A.2: Distribution of projected arrival SOC at home, computed conditional on the end-of-charge SOC

A.2.2 Station Ownership and Station Size.

Table A.1 reports each network's share of stations (the extensive margin). The government owns 70.9% of sites. ChargePoint's apparent 13.5% overstates its role as an operator: it is a hardware- and software-as-a-service provider, and its stations are owned by dispersed local actors. Tesla owns 5.6% of sites; although its extensive-margin penetration is modest, its sites are larger and plausibly located at strategically chosen nodes to maximize vehicle sales through its vertically integrated proprietary network. FLO, the next-largest owner at 5.4%, both manufactures and operates stations, but its sites do not cluster spatially, so the network externality it captures is on the order of its share.

Table A.1: 2023 Network Structure, Extensive Margin

Charging Network	**Stations (%)**
Government	70.9
ChargePoint	13.5
Tesla	5.6
FLO	5.4
Non-Networked	1.6
Petro-Canada	1.4
Electrify Canada	1.0
Couche-Tard	0.6

Table A.2 reports the distribution of chargers per station in 2022. Two-charger sites are the mode (57.5%), and 82.7% of stations carry at most two chargers.

Table A.2: Distribution of Chargers per Station (2022)

Chargers per station	Stations	Share (%)	Cumulative stations	Cumulative share (%)
1	89	25.2	89	25.2
2	**203**	**57.5**	**292**	**82.7**
3	20	5.7	312	88.4
4	26	7.4	338	95.8
5+	15	4.3	353	100.0

A.2.3 Québec and British Columbia in 2022.

Tables A.3–A.5 underlie Fact 5. Table A.3 reports BEV adoption by make: Tesla accounts for 30% of BEVs in Québec versus 56% in British Columbia. Table A.4 reports the top models: Tesla models dominate the top of the distribution in British Columbia, whereas Québec's top models include a broader set of affordable compatible options (Bolt, Kona, Leaf). Table A.5 reports the composition of each province's fast-charging network.

Table A.3: BEV Adoption by Make in Québec and British Columbia (2022)

	Québec (QC)		British Columbia (BC)	
Make	Count	MS	Count	MS
Tesla	31,708	0.30	45,559	0.56
Hyundai	19,124	0.18	6,219	0.08
Chevrolet	16,488	0.16	5,875	0.07
Nissan	10,519	0.10	8,422	0.10
Ford	6,446	0.06	2,500	0.03

Table A.4: BEV Adoption by Model in Québec and British Columbia (2022)

	Model characteristics		Market share (MS)	
Model	Range (km)	Price (kCAD)	Québec (QC)	British Columbia (BC)
Model 3	495.5	67.5	0.22	0.35
Bolt	417.0	38.2	0.15	0.05
Kona	415.0	43.9	0.12	0.05
Leaf	317.3	40.2	0.10	0.10
Model Y	525.0	85.0	0.07	0.15

Notes: Prices are mean list prices in thousand CAD (pre-subsidy), averaged across trims within model-year. Range is the mean (km) by model in 2022.

Table A.5: Charging-Network Composition in Québec and British Columbia (2022)

Québec (QC)

Charging Network	MS
Circuit électrique	0.59
Tesla	0.21
ChargePoint Network	0.07
FLO	0.06
Couche-Tard	0.03

British Columbia (BC)

Charging Network	MS
Tesla	0.41
BC Hydro	0.09
FLO	0.08
Non-Networked	0.07
Electrify Canada	0.03

Notes: Top networks shown; shares need not sum to one. The government network is Circuit électrique in Québec and BC Hydro in British Columbia; Tesla is private and incompatible; all other networks are private and compatible with the public standard.

A.3 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 government 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 return leg (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

$$\widehat{L}_n = \text{km}_j \left(1 - b_n^{\text{start}}\right). \tag{28}$$

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

$$r_n = \frac{\widehat{L}_n - \text{dist}(o, s)}{2}. \tag{29}$$

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:

$$\mathcal{D}_n = \{d : \ \text{dist}(d, s) \leq r_n, \ \text{dist}(o, d) \geq \text{dist}(o, s)\}. \tag{30}$$

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:

$$\hat{d}_n = \arg\min_{d \in \mathcal{D}_n} |r_n - \text{dist}(s, d)|. \tag{31}$$

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.3 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 trips. Importantly, this validation relies on a minimal subset of the data and is not used in estimation.

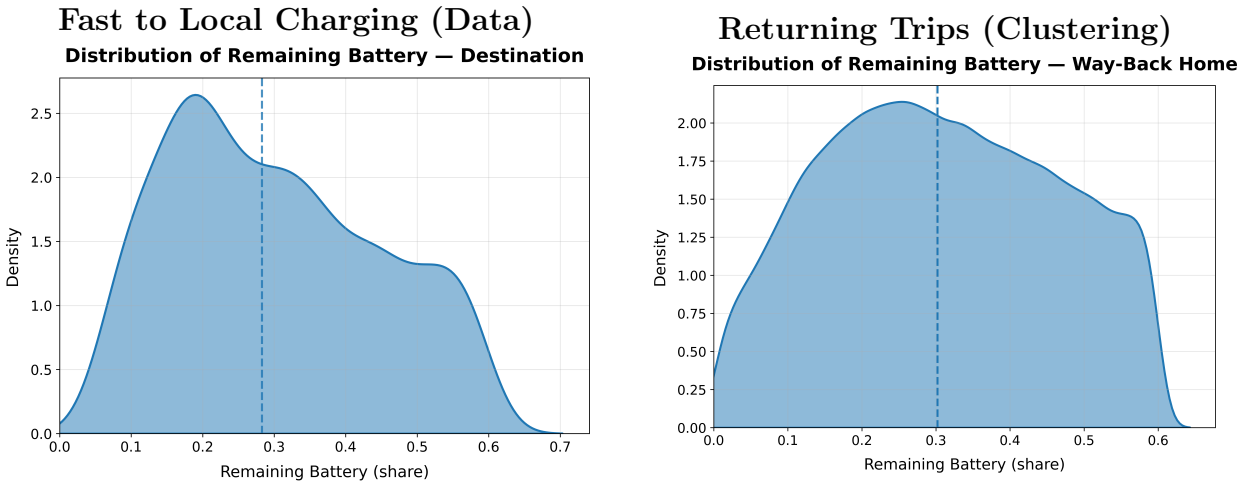


Figure A.3: Validation of returning-trip classification using fast-to-local charging sequences

Robustness

This section evaluates the robustness of the destination imputation rule introduced in Steps 3–5. I consider four specifications. In all cases, the inferred destination must lie in the feasible set $\mathcal{D}_n$ defined in Step 4, i.e., it cannot be farther than the implied destination radius r_n, and it must satisfy $\text{dist}(o,d) \geq \text{dist}(o,s)$. The hard constraint $\text{dist}(s,d) \leq r_n$ is mechanical: given the observed start-of-charge SOC, the destination cannot be beyond the remaining range implied by the data.

Model A: Distance-only best match (baseline)

The baseline imputation assigns the destination that best matches the implied one-way distance r_n:

$$\hat{d}_n^A = \arg\min_{d \in \mathcal{D}_n} \left| r_n - \text{dist}(s,d) \right|.$$

Model B: Empirical Bayes within the feasible set

Model B combines (i) a prior capturing destination attractiveness and (ii) a likelihood that favors destinations that make greater use of the implied remaining distance, conditional on feasibility.

Let π_d denote an empirical prior proportional to the number of hotel rooms in destination d:

$$\pi_d = \frac{\text{HotelRooms}_d}{\sum_{d' \in \mathcal{D}} \text{HotelRooms}_{d'}}.$$

The logic is that destinations differ in inherent attractiveness, and hotel capacity is a transparent proxy for such differences.

Conditional on $d \in \mathcal{D}_n$, I specify a monotone likelihood in the station-to-destination distance:

$$\ell_n(d;\alpha) = \exp\left(\alpha \cdot \frac{\text{dist}(s,d)}{r_n} \right), \qquad \alpha > 0,$$

so that, among feasible destinations, those closer to the maximum implied distance r_n receive greater likelihood weight. The resulting (unnormalized) posterior is

$$p_n(d;\alpha) \propto \pi_d \cdot \ell_n(d;\alpha), \qquad d \in \mathcal{D}_n.$$

Four specifications

I implement the following four destination rules:

1. Distance-only best match (Model A):

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

2. Empirical Bayes MAX (Model B, MAX):

$$\hat{d}_n^{B,\text{MAP}} = \arg\max_{d \in \mathcal{D}_n} p_n(d; \alpha).$$

3. Empirical Bayes RAND (Model B, RAND):

$$\hat{d}_n^{B,\text{RAND}} \sim \text{Uniform}(\mathcal{D}_n).$$

4. Empirical Bayes 90th %ile (Model B, 90th):

$$\hat{d}_n^{B,90} = \inf \{d \in \mathcal{D}_n : F_n(d; \alpha) \geq 0.90\} ,$$

The MAX rule selects the most likely destination under the posterior, while the RAND rule provide extreme and non-informative alternatives within the same feasible set. I compare results across the four specifications to show that the main findings are not driven by the particular destination imputation rule.

Table A.6: Imputation Robustness Check

	Baseline	EB MAX	EB RAND	EB 90th
Battery Deviation at Start, η	-0.10^{***}	-0.15^{**}	-0.10	-0.12
	(0.028)	(0.065)	(0.118)	(0.096)
Detour (minutes), λ_2	-0.11^{***}	-0.09^{***}	-0.04^{***}	-0.01
	(0.017)	(0.030)	(0.007)	(0.011)
Local Competition, κ	-0.39^{***}	-0.37^{*}	-0.37	-0.44
	(0.047)	(0.211)	(0.402)	(0.279)

Notes: This table reports maximum likelihood estimates of the charging-station choice model under alternative destination imputation rules. The baseline corresponds to the distance-only best match. EB MAX selects the posterior mode within the feasible set. EB RAND draws uniformly within the feasible set (Python seed =123). EB 90th selects the 90th %ile of the posterior distribution. $\alpha = 700$ is chosen as the value that maximizes the log-likelihood of the EB MAX specification. Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

The Baseline and EB MAX specifications — which are most consistent with plausible consumer destination choice and the structure of the data — yield coefficients with the expected negative signs and strong statistical significance. Battery deviation, detour time, and local competition all reduce the probability of selecting a station, as predicted by the model.

When the imputation rule is deliberately perturbed (EB RAND and the 90th-%ile rule), statistical precision declines and some coefficients lose significance. However, the estimated signs remain predominantly negative, and the magnitudes remain of similar order.

Overall, specifications that more closely approximate underlying consumer behavior produce sharper and more precise estimates, while alternative rules reduce precision without reversing the economic interpretation. The main qualitative conclusions are therefore robust to the destination imputation rule.

Estimation

A.3.0.1 Charging

Table A.7: Charging Station Choice Estimation

	(1)	(2)	(3)	(4)	(5)	(6)
Battery Deviation at Start, η	-0.11^{***}	-0.10^{***}	–	–	-0.11^{***}	-0.10^{***}
	(0.002)	(0.002)			(0.008)	(0.030)
Detour (minutes), λ_{min}	–	–	-0.12^{***}	-0.03^{***}	-0.12^{***}	-0.03
			(0.001)	(0.001)	(0.022)	(0.08)
Local Competition, κ	–	–	–	–	–	–
Station FEs (ζ_s)	No	Yes	No	Yes	No	Yes
Realized charging events	61,505	61,505	61,505	61,505	61,505	61,505
Mean choice set size	15	15	15	15	15	15
15						
N Observations	916,044	916,044	916,044	916,044	916,044	916,044

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 5% 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%. Columns (2), (4), and (6) include station fixed effects. Column (7) introduces a measure of local competition defined as $\log(n_{50km})$. Standard errors are based on the inverse of the quasi-Newton Hessian approximation returned by the MLE routine.
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

A.3.1 EV demand

Table A.8: Robustness: Result of EV Demand Estimation

	Alternative Specifications		
	(1)	(2)	(3)
Price	-0.045^{***}	-0.019^{***}	-0.020^{***}
	(0.002)	(0.003)	(0.003)
Λ_{ojt}	0.617^{***}	0.824^{***}	1.423^{***}
	(0.070)	(0.072)	(0.167)
Range (per 10 km)		-0.078^{***}	-0.066^{***}
		(0.008)	(0.009)
$\Lambda_{ojt}\times$ Range (per 10 km)			-0.012^{***}
			(0.003)
Mean own-price elasticity	-3.18	-1.37	-1.39
EV category elasticity of Λ	1.05	1.40	1.55
Product×Year, ξ_{jt}	Yes	Yes	Yes
City×Year, μ_{ct}	Yes	Yes	Yes
Supply Instruments	Yes	Yes	Yes
N (Observations)	19,400	19,400	19,400

Notes: Each observation is a city-year-model-trim (2016–2022). The regression is restricted to BEVs; the outside option includes hybrids and gasoline vehicles. Column (1) includes Price and Λ_{ojt}. Column (2) adds vehicle Range (per 10 km). Column (3) further includes the interaction $\Lambda_{ojt}\times$Range. All specifications absorb multiple layers of unobserved demand heterogeneity. Product×year unobserved demand components ξ_{jt} capture model-specific shocks common across markets in a given year. City×year unobserved demand absorbs local market conditions and policy environments. Standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Value of a Marginal Station

Value of a Marginal Station. To provide economic intuition for the magnitude of the estimates, I conduct a set of back-of-the-envelope counterfactuals on the charging network. For each candidate location, I add a single charger while holding all other primitives fixed and recompute the model-implied network value. The difference relative to the baseline defines the marginal value of that location. All numbers are computed for a steady state using 2022 technology and the 2016 charging network.

Table A.9 unpacks the underlying mechanism by providing comparative statics for adding a single charging station, decomposed into three effects. First, the *network effect*: how a given charging station affects charging demand. Second, the *infra-marginal effect*: how it affects the utility derived from owning an EV. This is downstream of the network effect—as the new station enables further trips and higher-quality charges, the network quality perceived by existing owners increases. Third, this improvement in network quality also induces marginal adopters, captured by the *marginal effect*.

Starting with the network effects, the table reports two variables: (i) Demand and (ii) Net Addition. Demand reflects the charging events captured by the new station. Because the 2016 network is sparse, most of these charges are additional. Some business stealing does occur, however—the difference between Demand and Net Addition measures its extent. Figure A.6 further illustrates this mechanism by mapping the static substitution across stations, holding the EV stock fixed.

Turning to the infra-marginal effect, the improvement in network quality directly raises the value of owning an EV. The distribution is highly heterogeneous: the median WTP is \$45 per BEV (\$340k in aggregate), the mean is \$118 (\$900k), and the 99th %ile reaches \$585 (\$4.4M). Figure A.4 illustrates this heterogeneity by plotting ranked WTP across locations—the distribution resembles an exponential one, with a small number of locations generating very large gains and most delivering modest value. The figure is shown for a given vehicle range; the table reports summary statistics averaged across the range distribution. While central nodes tend to offer the highest average value, more remote locations can also generate substantial gains for nearby markets, as shown in Figure A.5.

Finally, the addition of a charging station also affects the marginal decision to purchase an EV. The best locations induce up to 337 additional EVs from a single charger, but this decays rapidly—the median charger induces only 5. Central nodes and gap-bridging locations contribute the most. Notably, stations located near existing chargers can have a negative effect on neighboring stations: excessive density in an area reduces the marginal value of each additional charger.

Table A.9: Back-of-the-Envelope Value of Adding One Charger to the 2016 Charging Network

	Network	*Infra-marginal*		*Marginal*
Rank	Net Add.	WTP/BEV	Agg. WTP	New EVs
99th pct.	11,331	3,306	25,330,572	306
Mean	2,357	66	505,692	59
Median	120	31	237,522	6

Notes: Each row reports the effect of adding one charger at a candidate location, given the 2016 network. *Network* effects capture network improvements and business stealing: "Net Add." is the net additional charging events induced, after business stealing. *Infra-marginal* effects hold the EV stock fixed: "WTP/BEV" is willingness-to-pay per BEV; "Agg. WTP" multiplies this by 7,662 BEVs in Québec. *Marginal* effects capture EV adoption responses: "New EVs" is the implied increase in the EV fleet.

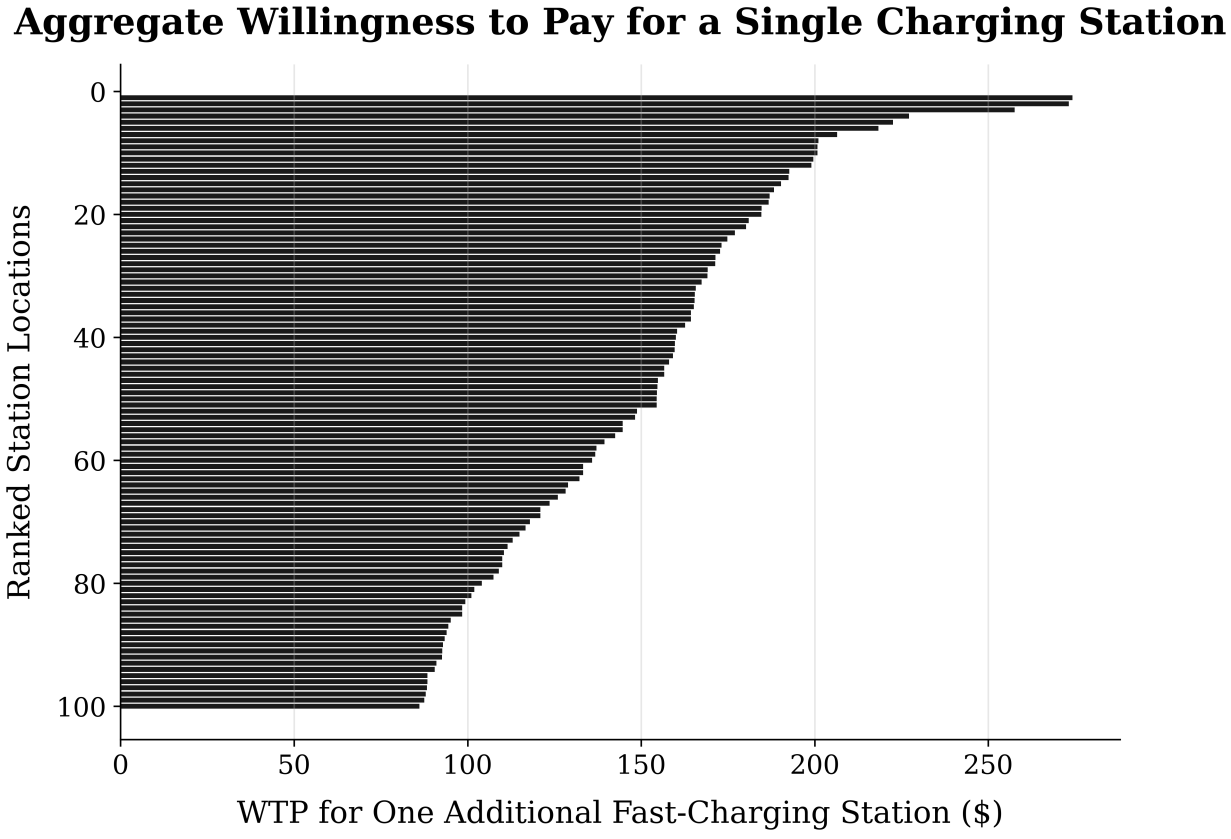


Figure A.4: Ranked WTP Across Candidate Locations

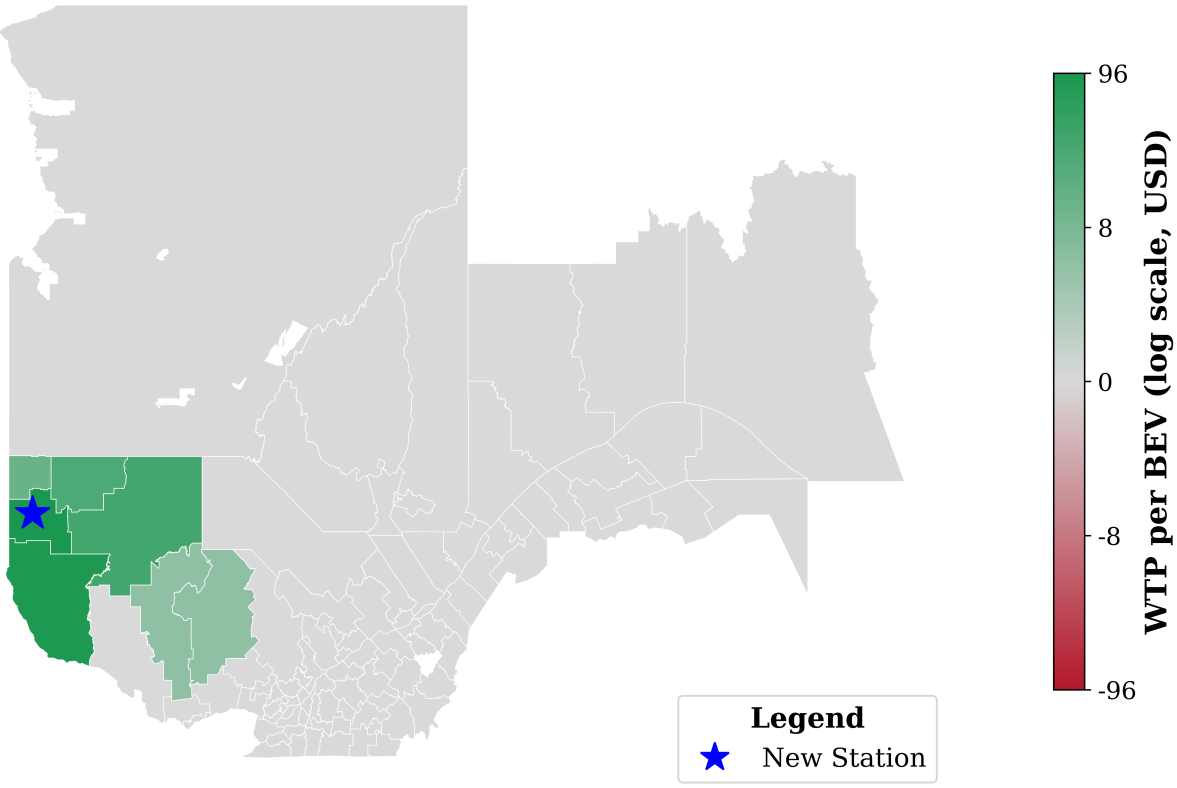


Figure A.5: Local WTP from Adding a Charging Station (2016)

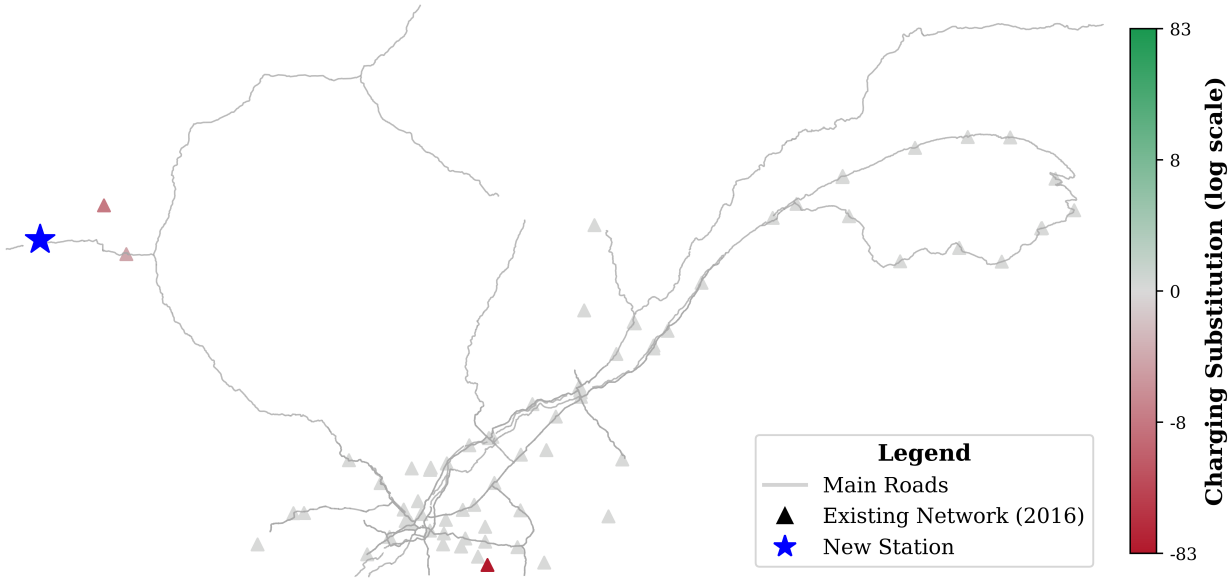


Figure A.6: Business Stealing from Adding a Charging Station (2016)

Willingness to Pay for Network Quality and Range

Table A.10: Willingness to Pay for Network Quality and Range

	Column (1)	Column (2)	Column (3)
WTP for a 1 s.d. increase in network quality (σ_Λ)			
Range = 344 km (mean)	\$15,036	\$27,650	\$25,150
WTP for a 10 km increase in range, Λ held fixed at $\bar{\Lambda}$			
	–	\$-2,072	\$-1,921
WTP for a 100 km increase in range, charging network held fixed			
Total	\$6,651	\$-8,484	\$-7,218
through network access ($\Delta\Lambda$)	\$6,651	\$12,231	\$11,461
direct range effect	–	\$-20,715	\$-18,679
Interaction effect			
100 km increase in range × 1 unit increase in Λ	–	–	\$762

Notes: WTP = $\Delta u/(-\alpha) \times 1,000$ (prices enter in \$1,000). Estimated price coefficients: $\alpha_1 = -0.044$, $\alpha_2 = -0.028$, $\alpha_3 = -0.029$. Column (1) includes only network quality Λ; Column (2) adds range; Column (3) adds the range × Λ interaction. The 100 km experiment moves each car in the estimation sample up two 50 km range groups and re-evaluates $\Lambda = E[K_{ojt}]$ for the same city, year, and charging network; entries are sales-weighted means across the estimation sample. In Column (1) the entire effect runs through Λ. The mean vehicle range is 344 km; $\sigma_\Lambda = 0.900$; mean $\Delta\Lambda$ from +100 km is 0.352.