September 15, 2026

Eugene Perk Han Tan
PhD Candidate in Sustainable Development
Columbia University
420 W 118th St, New York, NY 10027
pt2535@columbia.edu | (631) 746-5004

Dear Professors Breza, Chandrasekhar, and Viviano,

I am writing to apply for the NBER Workshop on Development Economics and Networks for Early-Career Researchers. I am a doctoral candidate in Sustainable Development at Columbia University, shortly to defend next month, after which I will be heading to the University of Calgary as a postdoc. My research studies development, energy, and climate economics using network, industrial organization, and spatial tools.

I would like to present my paper, "Who Pays for the Last Mile? Connection Charges and the Private Provision of Utility Infrastructure." I believe it is an excellent fit for the conference as I believe it is at the forefront of empirical network research. The paper studies how passing through the costs of infrastructure network expansion creates strategic incentives to cooperation and collectively provision infrastructure or to free ride on neighbors who would provision it.

The reduced-form part of the paper tests if Rwandans behave strategically in electricity adoption. I address four identification strategies with four empirical methods. As utilities endogenously place infrastructure to meet the spatial configuration of anticipated demand, I construct anticipated externality matrices that represents customers' expectations over grid expansion. As cooperation and free-riding are mutually exclusive strategic responses to the same incentives, I use dynamics to separate them and jointly instrument them with a graph-based shift-share and censored peer-of-peer instruments, where the censoring is an methodological innovation that can apply when the full network is observed. I also control for a confounding social peer effects through a second, non-infrastructure network. The structural part of the paper, though admittedly incomplete, embeds the dynamic graph as a state space of a dynamic discrete choice game, embedding network formation within structural industrial organization.

I would be excited about learning from established researchers such as yourselves, other early-stage network researchers, and discussing the frontier of empirical network research.

Professors Scott Barrett, Eric Verhoogen, or Daniel Björkegren could provide a recommendation if I am shortlisted and can be contacted at sb3116@columbia.edu, ev2124@columbia.edu, or dan.bjork@columbia.edu.

Thank you for considering my application.

Sincerely,
Eugene PH Tan

I'm interested in attending the workshop because I'm interested in learning from other researchers at the forefront of empirical work and receiving feedback on my own work. The set of researchers who can use network tools in development are a relatively small community, and so it would be an ideal place to learn from others at this forefront.

In addition to the above paper, which studies the private provision of networked public infrastructure, I am also working on papers that study the impact of the European Union's Carbon Border Adjustment Mechanism on developing country supply chains and their incentives to adopt carbon policies, network externalities in electric vehicle charging stations in California, and doing activity analysis (as in Tjalling Koopman's Nobel) with recent graph neural network tools to estimate the internal production network of a firm.

Eugene PH Tan

Columbia University
School of International and Public Affairs
420 W 118th Street, New York NY
United States of America

Phone: +1 (631) 746-5004
Email: eugene.tan@columbia.edu
GitHub: www.github.com/taneugene

Completed PhD Courses

Networks in Economic Theory
Industrial Development, Public and Development Economics
Industrial Organization I, II
Collective Action, Environmental Economics, Energy Economics.
Micro, Metrics, and Macro First-Year Sequences
Materials Chemistry, Advanced Thermodynamics, Idealized Models of Climate Processes.

Work-in-Progress

Who Pays for the Last Mile? Connection Charges and the Private Provision of Utility Infrastructure

International Growth Center Travel Grant (3000USD)

Expanding utilities typically impose upfront connection charges either through flat fees that charge every new customer the same price or distance- and material-based charges that pass through the costs of extending infrastructure to each customer. This paper argues that the latter creates a dynamic public-good provision game resembling a war of attrition: first movers pay a high price to expand infrastructure that cross-subsidizes their en-route and downstream neighbors. I study how these strategic incentives affect grid electricity adoption in 176 rural Rwandan distribution networks. To overcome reflection, I construct an externality network to measure how provision externalities change as the grid evolves, jointly instrument current and anticipated neighbor provision with shift-share and censored peer-of-peer instruments, and use a second network to control for social peer effects. My theory predicts that strategic incentives delay adoption; I find that when prospective customers anticipate a 10% reduction in their connection charge, their weekly adoption hazard falls by 0.91 percentage points relative to a mean 4.0 percent hazard in weeks when the network expands. For adopters who benefit from a neighbor's provision, I find that 36 percent of them cooperate and connect jointly, while 64 percent free-ride and only connect after they receive the cross-subsidy. Realized connection charges and willingness to pay have a correlation of 0.20, while flat fees would not vary with willingness to pay. This is consistent with theory that high-value adopters connect earlier because they have more to lose from waiting, a price discrimination mechanism that directs the cross-subsidy from high to low-value neighbors. To study how strategic behavior affects aggregate outcomes, I develop a structural model that nests a participation game of strategic complements within a dynamic waiting game of strategic substitutes. The potential structure of the subgame links coalition choices to customers' relative values of connecting and waiting. I use the resulting coalition distribution to estimate an equilibrium in which waiting values are consistent with the network expansion generated by neighbors' choices. Counterfactuals compare flat, distance-based, and externality-correcting charges, quantifying trade-offs between spatial allocation, equity, dynamic efficiency, and fiscal sustainability.

The Impacts of Carbon Tariffs on Developing Countries

Balancing Climate and Industrial Objectives when Pricing Dirty Upstream Goods

Activity Analysis with Graph Neural Networks

Publications

Brandt, A. R., Y. Sun, S. Bharadwaj, D. Livingston, **E. Tan,** and D. Gordon (2015). Energy return on investment for forty global oilfields using a detailed engineering-based model of oil production. *PloS ONE*, 10(12).

Gordon, D., **E.P.H. Tan**, and J. Feldman (2016). Changing Oils, Changing Management. *Journal of Industrial Ecology*, 20(4), 673-675.

Livingston, D. and **E. Tan,** (2015) Shale's True Contribution to the Oil Market. *Carnegie Papers (not peer-reviewed).*

Gordon, D., **E. Tan,** and K. Garner (2015) The Abundance of Oils in the Water Stressed Rockies. *Carnegie Papers (not peer-reviewed).*

Education

Columbia University, School of International and Public Affairs

Doctoral Candidate

Sustainable Development PhD

References:

Scott Barrett sb3118@columbia.edu

Eric Verhoogen eric.verhoogen@columbia.edu

Daniel Björkegren dan.bjork@columbia.edu

Master of Philosophy

Master of Arts

Yale University, School of Forestry and Environmental Studies

Master of Environmental Science

Environmental and Natural Resource Economics

Colorado College

Bachelor of Arts

Environmental Science with Chemistry

Dean's List 2013

London School of Economics and Political Science

General Course (Study Abroad)

Teaching

Columbia University SIPA or Economics Department

Instructor

Economics of Energy (Spring 2023)

Extraordinary Teaching Award, SIPA Energy and Environment Concentration

Teaching Assistant

Climate Change Politics and Diplomacy (for Prof. Barrett, thrice between 2023-2026)

Advanced Economic Development (for Prof. Verhoogen, Spring 2024)

Economics of Energy (for Prof. Druckenmiller, Spring 2022)

Intermediate Microeconomics (for Prof. Musatti, Spring 2021)

Empirical Analysis of Energy Markets (for Mercadal, Fall 2020)

Yale University Math Department
Teaching Fellow
Ordinary and Partial Differential Equations with Applications (for Benett, Fall 2018).
Ordinary and Partial Differential Equations (for Rolf, Fall 2017).

Colorado College Environmental Program
Teaching Assistant
Air: Atmospheric Chemistry and Physics (for Drossman and Whitten, Spring 2014).
Energy: Energetics and Thermodynamics (for Whitten, Fall 2013).

Invited Conference Presentations

2025:
Climate Pipeline Project at Harvard University

2024:
SusDever at Columbia SIPA

2023:
Northeast Workshop on Energy Policy and Environmental Economics at Cornell University
AERE at Eastern Economic Association

2022:
Heartland Environmental and Resource Economics Workshop at University of Illinois Urbana-Champaign

Employment outside Academia

World Bank, Global Facility for Disaster Reduction and Recovery (Jan-August 2019)
Data Science and Economics Short-Term Consultant

- Building tools to assess climate change resilience using Living Standards Measurement Study (LSMS) data, contributing to the Unbreakable report.

Center for Applied Data Science (April 2016 – May 2018)
Lead Instructor, Enterprise Data Science

- Led and taught three 42-day courses for Business Intelligence teams within enterprises across Asia. Each course covered about 2 semesters of Statistics and Machine Learning content, using and adapting content from The Data Incubator (now Pragmatic Institute).

Carnegie Endowment for International Peace (August 2014 – August 2015)
Junior Fellow, Energy and Climate Program.

- Provided research assistance for David Livingston and Deborah Gordon
- First research assistant compiling data and managing public release of the Oil-Climate Index.

Non-Research Honors

Fellowships

Junior Fellowship, Carnegie Endowment for International Peace (2014-2015)
Perdana Fellowship, Government of Malaysia (2013)

Scholarships

Davis UWC Scholarship, Davis UWC Scholars Program (2010-2014)

Funding

President's Special Projects Fund, Colorado College, 2012
Selig International Study Fund, Colorado College, 2012

Service and Leadership

Lead Organizer, SusDever at Columbia SIPA, 2024.
Organization Committee, Interdisciplinary PhD Workshop in Sustainable Development, Columbia SIPA, 2023.
PhD Student Committee, Columbia University 2021-2022.
Producer, Carnegie International Nuclear Policy Conference, Carnegie Endowment for International Peace 2015.

Computational Skills

Programming Languages: Python, JavaScript, Spark
Computational and Statistical Languages: GAMS, SQL, Stata, R.

Who Pays for the Last Mile? Connection Charges and The Private Provision of Utility Infrastructure *

Eugene Tan

September 2026

Expanding utilities typically impose upfront connection charges either through flat fees that charge every new customer the same price or distance- and material-based charges that pass through the costs of extending infrastructure to each customer. This paper argues that the latter creates a dynamic public-good provision game resembling a war of attrition: first movers pay a high price to expand infrastructure that cross-subsidizes their en-route and downstream neighbors. I study how these strategic incentives affect grid electricity adoption in 176 rural Rwandan distribution networks. To overcome reflection, I construct an externality network to measure how provision externalities change as the grid evolves, jointly instrument current and anticipated neighbor provision with shift-share and censored peer-of-peer instruments, and use a second network to control for social peer effects. My theory predicts that strategic incentives delay adoption; I find that when prospective customers anticipate a 10% reduction in their connection charge, their weekly adoption hazard falls by 0.91 percentage points relative to a mean 4.0 percent hazard in weeks when the network expands. For adopters who benefit from a neighbor's provision, I find that 36 percent of them cooperate and connect jointly, while 64 percent free-ride and only connect after they receive the cross-subsidy. Realized connection charges and willingness to pay have a correlation of 0.20, while flat fees would not vary with willingness to pay. This is consistent with theory that high-value adopters connect earlier because they have more to lose from waiting, a price discrimination mechanism that directs the cross-subsidy from high to low-value neighbors. To study how strategic behavior affects aggregate outcomes, I develop a structural model that nests a participation game of strategic complements within a dynamic waiting game of strategic substitutes. The potential structure of the subgame links coalition choices to customers' relative values of connecting and waiting. I use the resulting coalition distribution to estimate an equilibrium in which waiting values are consistent with the network expansion generated by neighbors' choices. Counterfactuals compare flat, distance-based, and externality-correcting charges, quantifying trade-offs between spatial allocation, equity, dynamic efficiency, and fiscal sustainability.

*This paper benefited from discussions and preliminary analyses with Gabriel Gonzalez Sutil, who I thank for insight and support. I thank Jin Hyun Park for excellent research assistance, and Rwanda's Ministry of Infrastructure for information about Rwanda's electricity policy.

1 — Introduction

There are extremely high capital costs of expanding networked utility infrastructure. Utilities must choose whether to recover such costs via public finance, through "fixed cost recovery" in consumption or network tariffs, or in upfront *connection charges* on new users. While from the perspective of usage markets, network infrastructure is a fixed cost, network components such as lines, pipes, and poles are all marginal costs of connecting new infrastructure users. So connection charges create a pricing problem because network components may eventually be shared: to what extent should utilities pass through infrastructure costs to users who require a marginal network extension?

In practice, distribution utilities use one of two approaches to pricing entry into distribution networks. The first approach is a flat fee, under which the utility pools planned, projected, and potentially subsidized capital costs, then divides the remainder by its intended user base to set the price of connection. Its apparent advantage is ex ante equity: connection charges do not depend on how remote new users are when they connect. Its disadvantage is that prices do not ration costly network extensions, so utilities need to plan coverage and decide who is eligible to connect in advance. *Distance-based charges*, the other approach, pass through marginal supply-side costs of network expansion to new adopters. Remote, isolated users decide for themselves whether the network extension they need to connect is worth the passed-through connection charge, so by comparison, distance-based charges are a market-based approach to connection pricing. But it has a catch. When first movers provision infrastructure extensions that prospective en-route and downstream users would share, they create a positive pecuniary dynamic externality that cross-subsidizes their neighbors.

This paper asks how marginal-cost pricing strategies within a market for utility connections affect who privately provisions infrastructure, when they adopt, and how their local decisions aggregate to affect overall outcomes. The net effect, even on adoption quantities, depends on an efficiency-equity tradeoff. On one hand, positive externalities are inefficient, because users who anticipate a neighbor's provision of shared infrastructure have a strategic incentive to delay uptake even when they can afford it. On the other hand, high-value users have more to lose from waiting and therefore have stronger incentives to provision infrastructure first. They may self-select into expensive early connections that equitably cross-subsidize their lower-value neighbors, allowing them to adopt at lower rates. To make sense of this mechanism, I start by developing a theoretical model that shows how distance- or materials-based connection charges create a dynamic public-good provision game. In addition to predicting inefficient delay and directed cross-subsidies, the model shows that strategic complements that predict joint adoption and dynamic strategic substitutes that manifest as free-riding. Next, I develop empirical tests of whether adoption responds to strategic incentives. I estimate the extent to which prospective customers anticipate, cooperate with, and free-ride on their neighbors' infrastructure provision using a case study of Rwandan electricity distribution networks, finding that prospective utility customers do indeed behave strategically. Finally, I estimate a structural model that nests a coalition-formation subgame within a dynamic waiting game on an evolving infrastructure network. I use it to quantify how, under counterfactual connection policies, local strategic responses aggregate to form spatial allocation, cost incidence, and utility financial outcomes.

The first contribution of this paper is to show that marginal-cost pricing creates a *dynamic* public-good

provision game in utility infrastructure network expansion.[1] Existing literature shows that income distributions affect private public-good provision by changing who voluntarily contributes (Warr 1983; Bergstrom, Blume, and Varian 1986); that Shapley values provide a standard solution concept for cooperatively sharing infrastructure costs in networked public-good games (Shapley 1953; Hart and Mas-Colell 1989; Littlechild and Owen 1973); and that users' locations and directed externalities determine efficiency and equilibrium provision (Allouch 2015; Elliott and Golub 2019). Adding dynamics means that the physical network changes endogenously and consequently affects the shape of externalities, the set of remaining candidate adopters, and connection charges.[2] As the grid expands, monotonically declining connection charges over time create a war of attrition similar to Bliss and Nalebuff (1984), where users with higher values connect first because they have more to lose from waiting.[3] In equilibrium, they select into cross-subsidizing the connection charges of lower-value en-route and downstream neighbors. The theory predicts that the utility can price discriminate an otherwise homogeneous connection when it uses distance-based charges and rations the rollout of the grid.

I test for strategic incentives within Rwandan electricity distribution networks. Sub-Saharan African electricity distribution networks are growing utilities where the stakes are particularly high. There are harsh fundamentals in low willingness to pay and high capital costs (Lee, Miguel, and Wolfram 2020b), but also incredibly ambitious electricity access targets spurred by Sustainable Development Goal 7 and Mission 300. Rwanda's institutional setting is one such context where the game occurs. The public utility uses distance-based connection charges. Rather than expanding the grid first and then inducing uptake, last-mile distribution network expansion is mostly triggered by customer connection requests. I combine unusually detailed, timestamped geospatial customer connections with distribution infrastructure network data to reconstruct the expansion of 177 rural distribution networks over time. The resulting data make the game observable. When network expansions occur, I can observe both sides of the externality: who provisions infrastructure and who subsequently benefits.

My theory proposes that adoption hazards around an anticipated charge-reducing network expansion should follow the shape of a heartbeat. Hazards decrease from a baseline equilibrium into the anticipated event, spike with joint adoption, and then regress into a higher baseline after connection charges have declined. The second contribution is to identify and decompose the strategic incentives that pump the heartbeat. There are four substantial identification challenges. Utilities endogenously place infrastructure to meet the spatial configuration of anticipated demand. Cooperation and free-riding are mutually exclusive strategic responses to the same infrastructure externality, making them difficult to identify separately.

1. Utility infrastructure is non-rival up to a capacity constraint. It is not strictly non-excludable because the utility can exclude users. Even so, utilities typically have legal duties and responsibilities to serve conditional on upfront charges. What matters for the game is that prospective users cannot exclude each other. While shared network components such as electricity poles theoretically are a congestible club good rather than a pure public good, I use the public-good label as shorthand for the private-provision game over them.

2. For this reason, I found that using a network formation approach such as Haller (2016), or more generally Bala and Goyal (2000) or Jackson and Wolinsky (1996) intractable. Utility customers do not choose who to link to, instead they have a binary choice of whether to adopt and the utility builds a predetermined least-cost extension. This avoids a dynamic problem over every feasible tree and sequence of adoption decisions.

3. By analogy to Bliss and Nalebuff (1984), waiting for a neighbor to connect first resembles awkwardly waiting for someone to start the party while encircling an empty ballroom dance floor. The one who can least stand it enters first, slaying the dragon of requesting that the utility provision network infrastructure.

Strategic incentives depend on neighbors' endogenous adoption decisions, which creates simultaneity and a reflection problem (Manski 1993). And infrastructure network externalities may be correlated with social network peer effects.

I match these four empirical challenges to four identification strategies. First, I construct and measure an anticipated *externality matrix* that represents customers' expectations over how the grid will evolve due to a neighbor's connection (Elliott and Golub 2019). Its construction removes the demand pull of other users on the infrastructure network's topography (Lipscomb, Mobarak, and Barham 2013; Faber 2014; Banerjee, Duflo, and Qian 2020). Second, anticipation, joint adoption, and free-riding incentives manifest in adoption behavior before, together with, and after a neighbor's externality, so I use dynamic measures of the externality matrix to separate them. Third, I jointly instrument cooperation and anticipation with respective supply-side shift-share and demand-side peer-to-peer instruments (Goldsmith-Pinkham, Sorkin, and Swift 2020; Bramoullé, Djebbari, and Fortin 2009). In the peer-of-peer instrument, I develop a method of censoring linkages that provide direct externalities that would break the exclusion restriction. Both instruments identify only those future network expansions that are anticipable from decision-time variation, which is likely a subset of what prospective customers can observe in their local neighborhoods. Finally, I control for confounding social peer effects through a second, non-infrastructure network constructed by interacting electricity connection status with a bilateral inverse-distance network (Blume et al. 2015).[4]

My reduced-form empirical results show that strategic incentives are quantitatively important. Using the above empirical strategies, I find that an anticipated 10 percent reduction in distance to the grid over the next eight weeks lowers the weekly adoption hazard by 0.910 percentage points. This corresponds to a 23 percent reduction relative to a 3.909-percentage-point hazard in weeks where another customer adopts within the network, and 2.4 times the unconditional mean weekly hazard of 0.379 percentage points. I test that this is not driven by supply-side activity in a placebo, and find that nearby network extensions that do not provide externalities do not result in the same strategic delay. Using a dyadic regression that compares joint to sequential adoption rates, I find that where strategic complements are theoretically weakest, 32 percent of customer pairs jointly adopt, while where they are strongest, 66 to 81 percent of them do. Then, using cross-sections of both utility and separate survey data, I show that customers with larger buildings or higher ultimate consumption pay higher connection charges. Conditional on having grid electricity, connection charges tend to have about a 0.2 correlation with indicators that reflect willingness to pay and about a 0.05 correlation with indicators that reflect income. A flat fee would result in a zero correlation by construction, so these results indicate selection into private provision and evidence that the mechanism price discriminates. Overall, all reduced-form results are consistent with theory.

The third contribution of this paper is to develop a structural model that allows me to nest a within-period participation subgame of strategic complements within an outer dynamic game of strategic substitutes. I relax the standard assumption in two-step conditional-choice-probability (CCP) estimators of discrete dynamic games that actions are conditionally independent and chosen simultaneously (Aguirregabiria, Collard-Wexler, and Ryan 2021). Instead, I allow for interdependent actions within a period and show how to construct model-consistent joint action probabilities, state transitions, and continuation values. In the

4. Social network peer effects in electricity primarily operate through consumption rather than connection decisions (Bernard and Torero 2015).

inner subgame, branch-relevant neighbors have a meeting where they receive revision opportunities to join a tentative coalition that jointly provisions their shared infrastructure. I show it is an exact potential game (Monderer and Shapley 1996). The critical part of the model is that the terminal coalition of the meeting must be described by a stationary Gibbs distribution (Blume 1993). The Gibbs distribution represents the multiple equilibria that result from strategic complements as separate peaks, fully describing interdependent participation decisions.

In the outer game, prospective customers construct continuation values that depend on neighbor coalition formation. Unlike the standard approach, customers observe a tentative coalition as they make decisions, so information sets differ and actions are correlated in equilibrium. Nevertheless, I show that Hotz–Miller inversion can still be applied to equate the coalition-conditional choice-value differences and exact potential differences to the log-odds of Coalition-Conditional Choice Probabilities (CoCCPs). Then I invert CoCCPs using Bayes' rule to form the coalition distribution conditional on waiting, which allows me to construct one-step-ahead continuation values as a function of Gibbs states. The structural part of the paper joins a recent industrial organization literature in modeling infrastructure adoption (Beatty 2022; dos Santos and Guidetti 2024; Johnston, Liu, and Yang 2023), more broadly, network externalities and economies of scale (Ackerberg and Gowrisankaran 2006; Björkegren 2019; Jia 2008), and a small literature at the intersection of industrial organization and network formation (Mele 2017). I simulate several counterfactual scenarios motivated by real-world policies, theory, and the ongoing structural transformation of the Sub-Saharan electricity sector, and assess how local strategic incentives aggregate to affect endogenous network expansion and electricity grid adoption.

Beyond the more fundamental economic questions above, the paper contributes to the literature on electricity access in developing countries. This literature has focused primarily on the impacts of providing electricity infrastructure on development outcomes, with substantially mixed evidence.[5] Very little literature studies connection pricing institutions. Lee, Miguel, and Wolfram (2020b) estimates the demand for grid connections in a randomized control trial and finds willingness-to-pay is low. Blimpo, McRae, and Steinbuks (2018) argues that low electricity consumption and limited tariff revenue mean Sub-Saharan utilities struggle to generate revenue from consumption, forcing them to charge high upfront connection fees to build out the grid. The current paper shifts attention from the impacts of infrastructure coverage and connection prices to economic mechanisms and market design. The natural prescription for expanding electricity access when focusing on infrastructure coverage or price levels is to increase public or concessional subsidies. By contrast, a market-design focus argues instead that economic mechanisms can harness strategic incentives to provide utilities with more bang for their buck. The present paper may also explain empirical irregularities observed in previous studies. For example, Walter and Moneke (2024) shows that electricity access tends to be bimodal across rural Zambian villages, with industrial users present in high access regions. Their mechanisms section argues that cost externalities may explain it, but do not quite point to the Zambian utility's distance-based connection charges.

The dynamic public-good provision game studied here extends beyond Rwandan electricity distribution

5. For example, see Lee, Miguel, and Wolfram (2020a), Khandker, Barnes, and Samad (2009, 2013), Lipscomb, Mobarak, and Barham (2013), Dinkelman (2011), Chakravorty, Emerick, and Ravago (2016), Grogan and Sadanand (2013), Rud (2012), Masselus et al. (2024), van de Walle et al. (2013), Van de Walle et al. (2017), Burlig and Preonas (2022), and Jeuland et al. (2021).

networks. To close, I surveyed utilities elsewhere and found that some have adopted complementary policies alongside distance-based charges to address or soften strategic incentives. For example, some condition connection charges on local population density ex ante, while others refund first movers ex post when second-movers connect to the infrastructure they provision. Other network utilities, from electricity transmission to piped water and gas to fiber-optic networks, face the same choice in pooling expansion costs or passing them through to new users. Overall, this paper argues that the dynamic externality created by marginal-cost connection pricing creates strategic incentives that shape the cost incidence, spatial extent, and the dynamic allocation of utility network infrastructure.

2 — A Network Expansion Game

This section presents a toy model of the strategic game created by distance-based utility connection policies. Despite its simplicity, the toy model incorporates network cost, geometry, timing, a bargaining subgame, and policy pricing. This section presents the game, its equilibria, and propositions relevant to the empirical exercises. I do not consider welfare consequences here because the structural model will compare the efficiency, equity, and financial consequences of alternative connection policies directly. The key insights of this section are:

1. Distance-based policies create positive externalities and cross-subsidies between neighbors.
2. Willingness to pay (WTP) is a necessary condition for adoption, but not sufficient. Players with strategic incentives may choose to delay adoption despite WTP.
3. Distance-based charges may create coordination failures, incentives to free-ride, or price discrimination. The first decreases connections, the second temporarily delays adoption, and the third may increase connections.
4. Strategic complements and strategic substitutes both occur, with the former within periods and the latter across periods. Strategic complements are stronger with the proportion of shared infrastructure.
5. The unique Nash equilibrium in a repeated game is one in which the higher-value player adopts first.
6. Known neighbor externalities produce a heartbeat-like shape in adoption hazards. Adoption hazards decrease in anticipation of the event, joint contemporaneous adoption results from cooperation, and free-riders adopt sequentially after charges have declined.

   I will test insights 4, 5, and 6 in the reduced form.

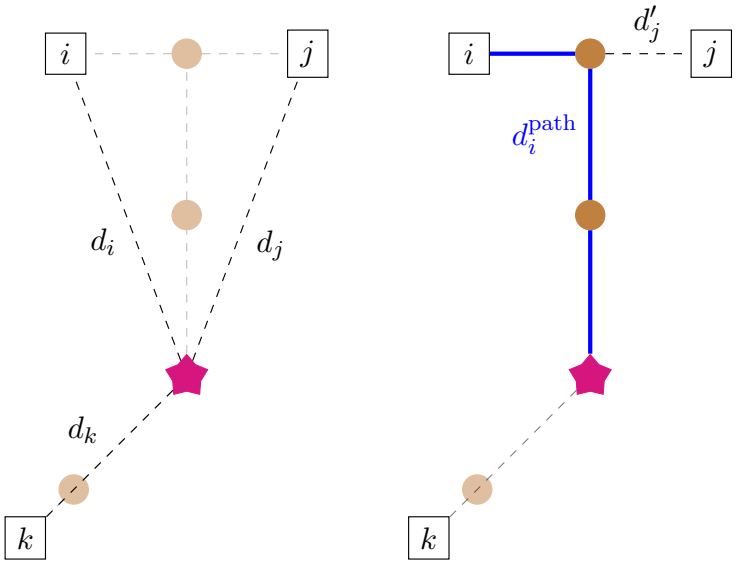
2.1. Positive Externalities and Cross-Subsidies

Figure 1 visualizes the externality in an infrastructure network graph.[6] Three potential players, i, j, and $k \in \mathcal{I}$, are depicted as squares. I refer to them as *players* or *adopters* throughout and as *buildings* or

6. Network nodes may be players, junction points, or sources, and edges represent the relevant linear infrastructure, such as pipes or wires. For my electricity-distribution application, nodes are buildings, poles, and transformers; edges are low-voltage lines.

prospective customers in the empirical analysis. The left panel is the initial state of infrastructure. The magenta star is its closest *access point* on the existing network to the three players, a point from which the utility can extend the network further. The utility has a least-cost plan minimizing the total cost of connecting all players, whose *junction nodes* are muted brown. These represent poles or pumps for wire- and pipe-based networks, respectively. Players do not observe the plan but may have beliefs over it. Customers observe their ex-ante straight-line distances to their closest points on the existing infrastructure, denoted d_i, d_j, and d_k, the closest point, and the spatial configuration of neighbors around them.

Figure 1 — Distance-based connection charges create externalities.



(a) Before adoption (b) After player i adopts

Notes. Panel (a) shows the initial state, in which players i, j, and k observe their Euclidean distances d_i, d_j, and d_k to their closest points on the existing infrastructure (magenta star). Ex-ante connection charges are based on each player's hypothetical unilateral extension path. The utility has an expansion plan based on spatial configuration (muted colors), with the planned node and edge extensions shown in light brown. Panel (b) shows the state of the network after player i connects. The utility installs nodes (dark brown circles) and connects them with edges (blue lines) to reach i. This is the network-path distance d_i^{path}. This extension reduces player j's distance to the network d'_j, creating a positive pecuniary externality and cross-subsidy equal to $c(d_j) - c(d'_j)$. Player k remains unaffected and receives no externality because they are not en route or downstream from the expansion that player i induces.

Customers understand the *connection charge policy* $c(\cdot)$ chosen by utilities, regulators, or governments.[7] Under a *distance-based policy*, player i pays $c(d_i)$; under a *materials-based policy*, they pay $c(d_i^{\text{path}})$. The former charges customers for the *minimum technical solution*, even if the utility chooses to extend the grid in a way amenable to connecting neighbors. The latter charges customers for whatever the utility actually builds; therefore, the realized network-path distance d_i^{path} may be longer than the straight-line distance d_i because of obstacles or the utility's least-cost plan. The charge may also depend on other variables such as player characteristics.

7. I do not assume differentiability because the policies are typically step functions.

Without loss of generality, the right panel of Figure 1 assumes that i connects first. In response to i making a connection request and paying for it, the utility installs the two nodes and three edges to connect player i to the existing infrastructure. Player i's connection request creates a positive *externality* by provisioning infrastructure that makes j closer to the network. The extension shortens j's distance to the closest access point from d_j to d'_j and lowers their connection charge from $c(d_j)$ to $c(d'_j)$. The size of the externality or *cross-subsidy* is $c(d_j) - c(d'_j)$. Who receives the cross-subsidy depends on the order of connections. Under a per-node material-based charge, costs are allocated to i, j, and k in a $(2, 0, 1)$ ratio if i connects first and $(0, 2, 1)$ if j connects first. Under a flat-fee policy, costs are allocated $(1, 1, 1)$. To achieve an equal allocation $(1, 1, 1)$, a distance-based policy would require i and j to cooperatively adopt and split the cost.

Realized connection-charge-induced externalities and cross-subsidies occur only between players who ultimately share infrastructure. Consider player k. As they have no neighbors en route or downstream from the planned network extension, they are isolated and would not induce any positive externalities or cross-subsidies upon adoption. Ex ante to adoption, however, the anticipated prospect of receiving externalities is uncertain. Players do not observe the utility's least-cost plan; if the utility had counterfactually expanded the grid towards i westwards instead of northwards, k might have received an externality from i. With incomplete information on the least-cost plan, strategic incentives may arise between neighbors who anticipate they are either en route to, or downstream of, a neighbor's network expansion.

The externality is *pecuniary, positive, local, directed, dynamic, state-dependent, and policy-induced.* I justify these properties as follows:

- **Pecuniary:** The externality affects prices.
- **Positive:** Extensions only decrease neighbors' connection charges.
- **Local:** It affects only en-route or downstream players relative to the network extension.
- **Directed:** Suppose that i were connected to the first node instead of the second in Figure 1. In this case, their adoption would reduce half of j's cost, but j's adoption would reduce all of i's cost in relative terms.
- **Dynamic and state-dependent:** Each infrastructure extension changes the state of the utility's infrastructure network. The externality from i to j would be smaller if the first junction node already existed. As infrastructure expands, possible externalities decrease in both magnitude and number.
- **Policy-dependent:** The externality sizes depend on the shape of the policy $c(\cdot)$. Positive externalities arise only if $c(\cdot)$ is weakly increasing with distance.

The anticipated externality matrix I develop for the reduced-form analysis inherits these properties.

I refer to the utility's supply-side capital costs as *connection costs* and the upfront fees faced by prospective players as *connection charges.*[8] Distance-based connection charges pass through marginal utility network-expansion costs to players, but flat fees do not. Let C^m denote the utility's *total cost of shared infrastructure*

8. These terms follow electricity industry parlance (Golumbeanu and Barnes 2013), but I use them generically.

for a bilateral pair of neighbors i and j who have strategic interactions.[9] Suppose the utility expects to recoup *total player charges for shared infrastructure* $C = \alpha C^m$ in connection charge revenue, where the remaining $(1-\alpha)C^m$ is subsidized with government subsidies or usage charges. Then the total amount C recovered through connection charges satisfies the bounds

$$\max\{c(d_i), c(d_j)\} \leq C \leq c(d_i) + c(d_j). \tag{1}$$

In the lower bound, C must cover at least the subsidized cost of connecting the more distant player. In the upper bound, there are weakly positive savings from sharing infrastructure instead of building two unilateral extensions. I let $M_{ij} = M_{ji} = c_i + c_j - C$ represent the symmetric *shared infrastructure savings*.

2.2. Strategic Network Game with Distance-Based Charges

Now we focus on the strategic incentives between i and j in a two-player game, assuming they share some infrastructure. The players choose whether to adopt or not in each period t. A *pure strategy* selects $s_{it} \in \{\text{Adopt}, \text{Wait}\}$, while a *mixed strategy* adopts with probability $\sigma_{it} \in [0,1]$. I abstract from distances and let players observe a *unilateral* or *first-mover connection charge* $c_i = c(d_i)$ and $c_j = c(d_j)$ and *second-mover connection charges* $c'_i = c(d'_i)$ and $c'_j = c(d'_j)$ at the beginning of the game. Players' private *connection values* V_i, V_j, *outside options* W_i, W_j, and all connection charges are assumed to be nonnegative: $V_i, V_j, W_i, W_j, c_i, c_j, c'_i, c'_j \geq 0$. For exposition, I initially assume that, conditional on the first mover adopting, the other player subsequently connects: $s_{it+1} \mid (s_{-it} = \text{Adopt}) = \text{Adopt}$. This assumption holds if and only if each player is willing to pay their residual connection charge in the following period: $V_i - c'_i > W_i$ and $V_j - c'_j > W_j$, as there are no remaining strategic incentives. Table 1 shows the normal-form payoffs of the network game under this setup.

Table 1—Payoff matrix for the two-player dynamic networked public-good provision game

		Player j	
		Adopt	**Wait to Adopt**
Player i	**Adopt**	$V_i - \kappa_i,\ V_j - \kappa_j$	$V_i - c_i,\ \beta(V_j - c'_j)$
	Wait to Adopt	$\beta(V_i - c'_i),\ V_j - c_j$	$\beta W_i,\ \beta W_j$

Notes. Player values are V_i and V_j. Unilateral connection charges are $c_i = c(d_i)$ and $c_j = c(d_j)$. Cooperative connection charges κ_i and κ_j are the result of bargaining over the shared extension. If player i waits while player j adopts first, player i's subsequent connection charge is $c'_i = c(d'_i)$, and similarly for j. The parameter $\beta \in [0,1]$ weights second-period payoffs relative to first-period payoffs. The outside options are W_i and W_j. These payoffs assume that if one player adopts in the first period, the other player is willing to pay the connection charge: $V_i - c'_i > 0$ and $V_j - c'_j > 0$.

If a *leader* unilaterally adopts, they incur charges c_i or c_j to receive their connection values V_i or V_j. The *follower* connects in the second period at the *cross-subsidized connection charge* c'_i or c'_j and receives payoffs $\beta(V_i - c'_i)$ or $\beta(V_j - c'_j)$. The value of second-period outcomes is discounted at rate $\beta \in [0,1]$. Greater pure time preferences, a longer wait between periods in calendar time or slower utility responses, or risk-averse

9. In the figure, it would represent the total cost of the two junction nodes and the edges that connect them.

preferences all make β smaller.

A participation subgame is nested within the *(Adopt, Adopt)* cell. If the two players cooperatively adopt, they incur a *bargaining cost* ρC, where $\rho \geq 0$, and divide the total shared-infrastructure cost through *cost shares* that satisfy $\kappa_i + \kappa_j = (1+\rho)C$. I model κ using Shapley cost shares, which are a standard solution concept for cooperative games and also arise as an equilibrium outcome in non-cooperative bargaining games with transferable utility (Shapley 1953; Hart and Mas-Colell 1989, 1996). The Shapley cost share of player i conditional on reaching *(Adopt, Adopt)* is

$$\kappa_i = \frac{1}{2}\left[(1+\rho)C + c_i - c_j\right] = c_i - \frac{1}{2}\left[M_{ij} - \rho C\right]. \tag{2}$$

I derive first equality of Equation 2 in Appendix subsection B.1.[10] The second equality of Equation 2 substitutes in M_{ij}. Cooperation saves each player half the shared infrastructure savings, $\frac{1}{2}M_{ij}$, at a bargaining cost of $\frac{1}{2}\rho C$. Conditional on reaching *(Adopt, Adopt)*, Shapley cost shares, unlike Shapley values, do not depend on willingness to pay. Cooperation may not be optimal if the bargaining cost exceeds the infrastructure savings, $\rho C > M_{ij}$. Therefore, utilities can facilitate coordinated adoption by reducing ρ. Joint adoption requires both players to adopt simultaneously, so it is contemporaneous. By contrast, the benefits from free-riding accrue in the future, are anticipated at t, and are discounted to the present in $\beta(V_j - c'_j)$.

It will be helpful to define special cases for distance-based charges and outside options.

Definition 1 (Mechanical distance-based policy)**.** Infrastructure costs are passed through at a flat pass-through rate α. Let c_i^m and c_j^m denote the utility's initial infrastructure costs of connecting players i and j so that the corresponding connection charges are $c_i = \alpha c_i^m$ and $c_j = \alpha c_j^m$. Second movers pay the residual charge, $c'_i = \alpha c_i^{m\prime} = C - c_j = \alpha(C^m - c_j^m)$. Then for all $k \in \{i, j\}$, the externality, total connection charges, and Shapley cost shares satisfy

$$M_{ij} = c_k - c'_k \qquad C = c_k + c'_{-k}, \qquad \kappa_k = \frac{1}{2}[(c_k + c'_k) + \rho(c_k + c'_{-k})].$$

Savings are symmetric in absolute terms but asymmetric in relative terms because neighbors may be upstream or downstream of each other.[11] In the mechanical baseline, infrastructure cost spillovers are fully passed through at rate α to connecting players. Therefore, α never changes whether an externality exists, the direction of cross-subsidies, and whether infrastructure savings cover the bargaining cost.

Consider two options for the game's temporal structure:

10. The reader may ask what would happen if I reversed the design, centering the participation game and assigning the free-riding value as the opportunity cost of the participation game. This would imply that the participation game concerns how to split the consumption value of utility service rather than the connection cost. It would also make the repeated game less tractable.

11. Three example graphs may help illustrate. If the two players' paths do not overlap, then $M_{ij} = 0$ and there is no externality. If one path is fully nested within the other, so player i uses no infrastructure that would not also be used by player j, then $M_{ij} = c_i$ and $C = c_j$: if j adopts first, player i's connection is free, $c'_i = 0$, and the size of j's cross-subsidy to i is i's initial cost $c_i - c'_i = c_i$; if i adopts first instead, player j's second-mover charge is the residual cost $c'_j = c_j - c_i = C - M_{ij}$, the total less the shared infrastructure. In a Y-shaped graph, both players have some private infrastructure and a shared trunk, so $0 < M_{ij} < \min\{c_i, c_j\}$: the first mover pays $c_i > M_{ij}$ and the second pays the residual cost of the total less the trunk $c'_j = c_j - M_{ij}$.

Definition 2 (One-shot game)**.** The game ends if both players wait, so the outside option value is $W_i = W_j = 0$. Second movers still receive $\beta(V_k - c'_k)$.

Definition 3 (Repeated game)**.** In a stationary *repeated game*, if both players wait, the same game is played again next period. The outside option W_k is the undiscounted one-period-ahead equilibrium payoff with present value βW_k.

In the following section, I analyze when each cell of Table 1 is a pure-strategy Nash equilibrium (PSNE) in the one-shot game. Next, I present mixed-strategy Nash equilibria (MSNE) with a generic outside option, resembling a chicken game. Finally, I analyze MSNE in a repeated game with private values, which resembles a war of attrition.

2.3. Pure Strategy Nash Equilibria in a One-Shot Game

In this subsection, I assume the game is one-shot. Adopting is a dominant strategy for k if cooperation beats the second-mover payoff, $V_k - \kappa_k \geq \beta(V_k - c'_k)$, and unilateral adoption beats the outside option, $V_k - c_k \geq 0$. Waiting is a dominant strategy if players are unwilling to pay the unilateral upfront connection charge, $0 \geq V_k - c_k$, and the second-mover payoff is preferred to cooperative adoption, $\beta(V_k - c'_k) \geq V_k - \kappa_k$. The game at its core consists of two comparisons: the static incentive of whether connection values exceed the unilateral connection charge, and the strategic incentive of whether cooperation exceeds second-mover surplus.

From the point of view of utilities that target the connection quantities as a key performance indicator, such as electricity access measures in electricity distribution, *(Adopt, Adopt)* is ideal because both players connect immediately, achieving player surplus $V_i + V_j - (1 + \rho)C$. It is a PSNE if one period of connection value exceeds the difference between cooperative and second-mover charges $(1 - \beta)V_i \geq \kappa_i - \beta c'_i$, a demanding constraint when the second period may be as little as a week away. The right-hand side is $\kappa_i - \beta c'_i = c_i - \beta c'_i - \frac{1}{2}(M_{ij} - \rho C)$, so for joint adoption to occur, the discounted externality $c_i - \beta c'_i$ must be small, and cooperation costs must not be too high.[12]

A *(Wait, Wait)* PSNE may often be relevant when both players are unwilling to pay the upfront costs, $0 \geq V_k - c_k$. This PSNE is inefficient when there are gains from cooperation, $V_k - \kappa_k \geq 0 \geq V_k - c_k \forall k \in \{i, j\}$, but no player deviates to adopt. There can be a prisoner's dilemma if waiting is weakly dominant and free-riding beats cooperation in $\kappa_k - \beta c'_k \geq (1 - \beta)V_k \quad \forall k \in \{i, j\}$. There can also be a stag hunt if *(Wait, Wait)* and *(Adopt, Adopt)* are both equilibria; low connection rates can be the result of a coordination failure. Under distance-based charges, low connection rates could reflect strategic incentives rather than static WTP.

The most interesting behavior is in the *(Adopt, Wait)* or *(Wait, Adopt)* cells. It is a PSNE if at least one player is willing to pay the full initial cost of connection, $V_i - c_i \geq 0$, but the other player prefers to wait:

$$c_j - \beta c'_j - \frac{1}{2}(M_{ij} - \rho C) = \kappa_j - \beta c'_j \geq (1 - \beta)V_j. \tag{3}$$

12. Under the mechanical distance-based policy, this expression is $\frac{1}{2}c_i + (\frac{1}{2} - \beta)c'_i + \frac{1}{2}\rho C$, because the size of the externality is the shared infrastructure.

As they compare upfront connection charges or cooperation costs to just one period of flow benefits, the condition is not demanding. This PSNE nests three special cases. First, when Equation 3 holds, and the follower is willing to pay the upfront charge, there is a *chicken game* in which the leader provisions the public infrastructure and the follower free-rides. I will return to analyzing the chicken game's MSNE in the following subsection. Second, when Equation 3 holds, and the follower is unwilling to pay the upfront charge unilaterally, it produces *price discrimination* that expands connections in equilibrium. Finally, free-riding and price discrimination can occur simultaneously, in which case the charge can redistribute without expanding equilibrium quantities. I call $\beta(V_i - c'_i)$ the *second-mover payoff* when the second mover cannot afford the initial cost and the *free-riding payoff* when they can.

Free-riding occurs if both players are willing to pay the initial connection charge, $V_i - c_i \geq 0$ and $V_j - c_j \geq 0$, but one chooses to *Wait* because $\kappa_j - \beta c'_j > (1-\beta)V_j$. The leader provisions the shared infrastructure unilaterally, while the follower saves the externality cost, $c_k - c'_k$, even though it can pay the initial charge. Free-riding delays adoption but does not prevent it. Players wait for their neighbor to connect, which lowers their connection charge, and then take advantage of the cross-subsidy immediately afterward.

Price discrimination can expand the extensive margin of who can connect to utility infrastructure. Suppose high type j is willing to pay $V_j - c_j \geq 0$, a unilateral connection charge higher than a flat fee, but the low type is willing to pay only a second-mover connection charge

$$V_i - c'_i \geq 0 \geq V_i - c_i. \tag{4}$$

High type j *selects* into early adoption and connects at a high price, which creates a cross-subsidy of size $c_i - c'_i$ to i, allowing them to connect at a low price as a second mover. Low type i would not connect unilaterally. This price discrimination happens without coercion, as it is not from j's benevolence that i receives a cheap connection, but instead from j's own self-interest. Price discrimination can increase the quantity of connections.

Price discrimination can also redistribute on the intensive margin. Without loss of generality, suppose player i and player j's net surplus from unilateral adoption is ordered $V_i - c_i > V_j - c_j$. Lemma 1 shows that the parameter region in which *(Wait, Adopt)* is a Nash equilibrium is contained within the parameter region in which *(Adopt, Wait)* is a Nash equilibrium. If one of these is a unique first-mover equilibrium, the higher-surplus neighbor adopts first in it.

Lemma 1 (Higher net-surplus players lead if there is a unique asymmetric equilibrium.)**.** Without loss of generality, suppose that i has a higher unilateral adoption surplus $V_i - c_i \geq V_j - c_j$. If the first-mover relative surplus is larger than the second-mover relative surplus, $(V_i - c_i) - (V_j - c_j) \geq \beta\left[(V_i - c'_i) - (V_j - c'_j)\right]$, the support of *(Wait, Adopt)* being a PSNE is contained within the support of *(Adopt, Wait)* being a PSNE. Therefore, if there is a unique equilibrium, it is *(Adopt, Wait)* and player i is the first mover. If there are multiple equilibria, joint player surplus is weakly greater in the equilibrium where the higher-value player adopts first.

Proof. See Appendix subsection B.2.

The conditions for Lemma 1 are weak. Some ordering of net surplus will generically hold: $V_i - c_i \geq V_j - c_j$. The second condition $(V_i - c_i) - (V_j - c_j) \geq \beta \left[(V_i - c_i') - (V_j - c_j')\right]$ holds as long as i is not much further from the existing infrastructure in connection-charge terms than their difference in values. While support containment does not necessarily mean that the higher-surplus player will cross-subsidize the lower-surplus player in general, there are certain parameter supports where they will. If one further assumes positive income elasticities and symmetric costs, Lemma 1 shows that the cross-subsidy is directed, redistributing from wealthier to poorer neighbors. Later, Proposition 2 strengthens this lemma to show that when the game is repeated, there is a unique symmetric cutoff equilibrium in which price discrimination occurs.

2.4. Mixed Strategy Equilibria

In this subsection, I no longer assume that the game is one-shot and assume arbitrary continuation values. A chicken game exists when both players' unilateral surplus exceeds their continuation values, $V_k - c_k > \beta W_k$, but both players have incentives to free-ride via Equation 3. A MSNE can then occur where players randomize between adopting and waiting. I limit the following discussion to MSNE predictions that differ from those of the one-shot PSNE. These include strategic complements and substitutes, a refinement of Lemma 1 that makes price discrimination the unique equilibrium, and how players should react when they anticipate a future neighbor externality.

A mixed-strategy Nash equilibrium exists if both players are unilaterally willing to pay and free-riding benefits exceed savings from cooperative adoption. As mixed strategy equilibria imply indifference between strategies conditional on the neighbor's mixed strategy, i's equilibrium strategy is a function of j's primitives

$$\sigma_i = \frac{\overbrace{V_j - c_j - \beta W_j}^{\text{WTP given outside option}}}{\underbrace{\beta(V_j - c_j')}_{\text{free-riding benefit}} - \underbrace{\tfrac{1}{2}(M_{ij} - \rho C)}_{\text{cooperative savings}} - \underbrace{\beta W_j}_{\text{outside option}}}. \tag{5}$$

Comparative statics of Equation 5 show that adoption increases in neighbor value V_j, decreases in the neighbor's outside option W_j and standalone charge c_j, if c_j', M_{ij}, and C are held fixed. If cooperation is useful, $M_{ij} \geq \rho C$, then adoption also decreases in the discount factor.

My first proposition is that the MSNE can result in both *strategic substitutes and strategic complements*.

Proposition 1 (Strategic complement and substitute regimes)**.** Assume that continuation values W_i are fixed. There are strategic complements if and only if

$$\frac{1}{2} \underbrace{(M_{ij} - \rho C)}_{\text{cooperation benefits}} + \underbrace{\beta W_i}_{\text{outside option}} > \underbrace{\beta(V_i - c_i')}_{\text{free-riding payoff}}, \tag{6}$$

and strategic substitutes if the opposite holds.

Proof. Strategic complements occur if and only if the cross-partial of payoffs to own and neighbor strategies is positive; strategic substitutes occur if it is negative. Player i's payoff is $\pi_i(\sigma_i, \sigma_j) = \sigma_i \sigma_j (V_i - \kappa_i) +$

$\sigma_i(1-\sigma_j)(V_i-c_i)+(1-\sigma_i)\sigma_j\beta(V_i-c_i')+(1-\sigma_i)(1-\sigma_j)\beta W_i$. Its derivative is the marginal value of adoption,

$$\frac{\partial \pi_i}{\partial \sigma_i} = \sigma_j(V_i-\kappa_i)+(1-\sigma_j)(V_i-c_i)-\sigma_j\beta(V_i-c_i')-(1-\sigma_j)\beta W_i.$$

Taking the second derivative in neighbor strategies gives the cross-partial

$$\frac{\partial^2 \pi_i}{\partial \sigma_i \partial \sigma_j} = \underbrace{\left[(V_i-\kappa_i)-\beta(V_i-c_i')\right]}_{\text{cooperate or free-ride?}} - \underbrace{\left[(V_i-c_i)-\beta W_i\right]}_{\text{unilaterally adopt or continue?}} = \tfrac{1}{2}(M_{ij}-\rho C)+\beta W_i-\beta(V_i-c_i').$$

The last equality substitutes $\kappa_i = c_i - \frac{1}{2}(M_{ij}-\rho C)$. □

Proposition 1 shows how primitives interact to define a *strategic regime*. When the gains from sharing infrastructure M_{ij} are large relative to bargaining costs ρC, players have an incentive to coordinate and adopt jointly. When the outside option W_i is large, players would delay to coordinate adoption rather than adopt unilaterally. By contrast, second-mover payoffs $\beta(V_i-c_i')$ create an incentive to free-ride and favor strategic substitutes. Recall that β represents how the second-period payoffs are discounted relative to the first, and so increases with a responsive utility, short time between periods, low discount rates, and low risk aversion. All these would create strategic substitutes, as would lower pass-through.[13]

All primitives enter Proposition 1, so strategic regimes will vary a lot across neighborhoods and even within bilateral dyadic pairs of neighbors, and across utility types, jurisdictions, contextual, spatial, and social network settings. Given the Adopt and Wait strategies, the strategic regime can be observed in relative adoption timing. Strategic complements predict *synchronous* adoption. Coalitions bargain internally over how to share infrastructure-extension costs and collectively negotiate the extension with the utility. The utility then connects all coalition members to utility infrastructure at once. Strategic substitutes predict *sequential* adoption. The first-mover provisions the network extension, and then their neighbor observes a lower cost and adopts only then. These patterns will allow me to distinguish free-riding from cooperation, strategic substitutes from complements when we turn to empirics.

Neither strategic substitutes nor strategic complements are necessarily preferable to the other. Although strategic substitutes create a dynamic efficiency cost through free-riding, the equity benefits resulting from Lemma 1 may mean that more players can connect in equilibrium. Furthermore, strategic substitutes do not necessarily mean that connections are slow; the sequence may result in a "snowball" effect once the first mover induces the initial infrastructure (Dahl, Løken, and Mogstad 2014). Conversely, strategic complements are not necessarily good. They can result in multiple equilibria with high and low adoption, and players may not connect because they are stuck in a coordination failure.[14] However, both strategic incentives delay network adoption relative to what static incentives alone would predict. Under strategic

13. A higher pass-through rate (smaller subsidies) expands the strategic-complements regime. Under a mechanical baseline policy, for example, each infrastructure-related charge in Equation 6 is proportional to α, so the cross-partial equals $\alpha\left[\frac{1}{2}(M_{ij}^m-\rho C^m)+\beta c_i^{m\prime}\right]-\beta(V_i-W_i)$. With the spatial structure of shared infrastructure M_{ij}^m and the second-mover cost $c_i^{m\prime}$ held fixed, the cross-partial is zero when $\alpha_i^* = \beta(V_i-W_i)/\left[\frac{1}{2}(M_{ij}^m-\rho C^m)+\beta c_i^{m\prime}\right]$. The network switches from strategic substitutes to strategic complements when α moves above this threshold.

14. More precisely, anticipating that a neighbor will soon be ready to adopt raises W_k and decreases the payoff from adopting now: $\frac{\partial^2 \pi_i}{\partial \sigma_i \partial W} = -(1-\sigma_j)\beta < 0$. If bargaining were instantaneous, there would be no delay.

complements, a player delays to coordinate with neighbors who may not yet be ready. Under strategic substitutes, they delay in the hope that the neighbor would move first.

2.5. Price Discrimination and Anticipation in the War of Attrition

Assume now a repeated game in Definition 3, and that connection values are private. This results in a war of attrition. I show two results. First, Proposition 2 refines Lemma 1 to show that with private values, there exists a symmetric cutoff equilibrium in which the high-type neighbor selects into adopting first. Under strategic substitutes, this cutoff equilibrium is unique. Second, I describe the shape of mixed-strategy equilibria around an anticipated neighbor adoption in Proposition 3, a result that provides testable hypotheses for my empirical exercises.

For simplicity, I assume symmetric charges: $c_i = c_j = c$, $c'_i = c'_j = c'$, and $\kappa_i = \kappa_j = \kappa$. This reflects Figure 1 and does not change the qualitative results.[15] Connection values V_k are known only to player k and are now random values drawn independently from a common distribution F that has positive density f on $[\underline{V}, \overline{V}]$. The charges c, c', and κ, the discount factor $\beta < 1$, and the distribution F are all common knowledge, but individual connection values V_i, V_j are private. Let $W(V_k)$ replace the continuation value W_k for player k, so that I can assume it is differentiable in the connection value and further that it satisfies the condition $\beta W'(V_k) < 1$: higher connection values raise adoption payoffs more than they raise the discounted continuation value. Then let $s_j(V_j) \in \{\text{Adopt}, \text{Wait}\}$ denote a neighbor's decision rule, now a function of their connection value draw, and let j's mixed strategy be $\sigma_j = \Pr\big[s_j(V_j) = \text{Adopt}\big]$. Player i's net benefit from adopting is

$$\Delta_i(V_i; \sigma_j) \equiv \frac{\partial \pi}{\partial \sigma_i} = \sigma_j \underbrace{\big[(V_i - \kappa) - \beta(V_i - c')\big]}_{\text{Cooperate or free-ride?}} + (1 - \sigma_j) \underbrace{\big[(V_i - c) - \beta W(V_i)\big]}_{\text{Unilaterally adopt or continue?}}. \tag{7}$$

Equation 7 shows that i performs a cost-benefit analysis when deciding whether to adopt. They compare cooperation to free-riding and unilateral adoption to continuing the game, and weight these comparisons by neighbor adoption probabilities. The neighbor's adoption probability σ_j determines the relative importance of static and strategic incentives. The following proofs in Lemma 2 and Proposition 2 adapt war-of-attrition and single-crossing intuition from Bliss and Nalebuff (1984)'s dynamic public good provision game.

Lemma 2 (Best Responses are a Cutoff Strategy)**.** Player i's best response to a neighbor adoption probability σ_j is a cutoff, a threshold $V^*(\sigma_j)$ such that player i adopts if and only if $V_i \geq V^*(\sigma_j)$. When the cutoff is interior,

$$\frac{dV^*}{d\sigma_j} = -\frac{\partial \Delta_i / \partial \sigma_j}{\partial \Delta_i / \partial V_i}, \tag{8}$$

is positive under strategic substitutes and negative under strategic complements. An increase in the neighbor's adoption probability σ_j under strategic substitutes decreases own adoption by raising the cutoff. Under strategic complements, it increases adoption by decreasing the cutoff.

Proof. See Appendix subsection B.3.

15. Asymmetric charges would produce player-specific cutoff strategies. The results continue to hold conditional on charges.

Proposition 1 says the strategic regime determines whether prospective customers adopt jointly or wait in response to higher neighbor adoption probabilities. While it follows directly from the definitions of strategic substitutes and complements, the lemma's form helps to prove Proposition 2.

Proposition 2 (Higher value players connect first.)**.** Given a fixed continuation-value function $W(V)$, a symmetric cutoff equilibrium exists. If strategic substitutes hold, it is unique. If strategic complements hold, there may be more than one symmetric cutoff equilibria. In any such equilibrium, players with $V_k \geq V^*$ connect first, while players with $V_k < V^*$ wait. Therefore, adoption decisions are weakly increasing in connection value, so high-connection-value players select into adopting earlier at the first-mover charge, while low-connection-value players wait to face the second-mover charge.

Proof. See Appendix subsection B.4

The proposition sharpens Lemma 1.[16] While there, multiple equlibrium did not necessarily indicate price discrimination, here, there may still be multiple equilibria, but it shows that every cutoff equilibria price discriminates. Specifically, whenever a lower-value player adopts, their higher-value neighbor will too, so the higher-value neighbor is weakly more likely to adopt first. Proposition 2 shows that price discrimination is an equilibrium outcome.[17]

Finally, Proposition 3 describes what happens in an event study around a foreseeable neighbor's connection. It forms the basis of testable hypotheses for my primary reduced-form exercises. I compare two versions of the war of attrition. I refer to the game described in Proposition 2 as the *baseline war* and compare it against a *foreseeable event war*. The foreseeable event is that a third player k will connect at time T^k, and both players i and j foresee it. Player k provides both i and j with an externality, decreasing their symmetric connection charge from c to $c^k \in [c', c]$. By comparing the option value of waiting for the foreseeable event against that of the baseline war, the proposition describes an adoption pattern shaped like a heartbeat.

Proposition 3 (Heartbeat Adoption Hazards)**.** I assume strategic substitutes, that k is more upstream relative to the existing infrastructure than the bilateral pair i, j, so that the externality is of intermediate size $c^k \in [c', c]$ and $\kappa^k > \kappa$, and that the continuation values of the baseline war $W(V)$ are fixed so that beliefs about neighbor values do not update.

Both wars have a unique symmetric cutoff equilibrium because of Proposition 2. Let V^*_{t,T^k} and V^* denote the cutoff strategies of the foreseeable event war and the baseline war respectively; the latter is stationary. And let $\sigma_{t,T^k} = 1 - F(V^*_{t,T^k})$ and σ_t denote the corresponding adoption probabilities.

I refer to the inequalities $\beta^{T^k - t}(\overline{V} - c^k) \geq \overline{V} - c$ and $\beta(\overline{V} - c') > \overline{V} - \kappa$, collectively as a freeze condition, because these make waiting a dominant strategy in both wars for even the upper-bound private value. The

16. There, in the one-shot game with complete information, high-value players will weakly tend to connect first because a unique equilibrium may exist where they do so. However, when multiple chicken equilibria exist, the lemma says nothing about who adopts first.

17. In Bliss and Nalebuff (1984), low-cost, high-net-benefit players voluntarily provided a shared good first because they had more to lose when the game continues into the next period. Here, although charges are not voluntary, the intuition is the same. High-value players select into early adoption because they have more to lose from delaying connection, selecting into earlier, more expensive connections that provision the shared infrastructure. This cross-subsidizes low-value players, who wait for the lower second-mover charge. If they had no choice, it expands connections.

first inequality holds if at time t, waiting for the anticipated event beats unilateral adoption. The second inequality holds if free-riding beats cooperation in the stationary baseline. By comparison to the baseline, the foreseeable event war affects adoption as follows:

1. *Baseline Unaffected.* If the freeze condition does not hold, and $\beta^{T^k-t}(V^* - c^k) \leq \beta W(V^*)$, the anticipated event does not change current payoffs and k's anticipated adoption has no effect. Cutoffs are equal in both wars $V^*_{t,T^k} = V^*$, and adoption is unaffected.

2. *Increasing Delay ahead of the externality.* If k's anticipated adoption makes waiting more attractive $\beta^{T^k-t}(V^* - c^k) > \beta W(V^*)$ but is not strong enough to make the freeze condition hold; then cutoffs increase $V^*_{t,T^k} \geq V^*$. Players reduce adoption to wait for the event $\sigma_{t,T^k} = 1-F(V^*_{t,T^k}) \leq 1-F(V^*) = \sigma^*$. And $\beta^{T^k-t}(V^* - c^k)$ is increasing into the event, adoption decrease into it.

3. *Frozen adoption, an optional phase.* When the freeze condition holds, waiting is optimal for every type $V \in [\underline{V}, \overline{V}]$ because holding out for player k weakly dominates unilateral or joint adoption. Both players wait, and current adoption is zero.

4. *Joint adoption with k at time T^k.* When player k adopts, player i can join k's connection request and pay the Shapley cost share κ^k. Joining k's coalition to adopt at time-T^k is preferred to waiting one more period and connecting at the post-event charge c^k if $V_i - \kappa^k \geq \beta(V_i - c^k)$. This means that i has a higher value than a threshold $V_i \geq V^{\text{joint}}_{T^k} \equiv (\kappa^k - \beta c^k)/(1-\beta)$, which will allow me to segment the private value distribution. Let a *survivor cutoff*, the lowest adoption cutoff players faced before T^k, be denoted $V^{\text{surv}}_{T^k} \equiv \min_{s<T^k} V^*_{s,T^k}$. The remaining players just before T^k is $V_i < V^{\text{surv}}_{T^k}$. Joint adoption with k, conditional on reaching the event, occurs with positive probability $\left(F(V^{\text{surv}}_{T^k}) - F(V^{\text{joint}}_{T^k})\right)/F(V^{\text{surv}}_{T^k})$ when $V^{\text{joint}}_{T^k} < V^{\text{surv}}_{T^k}$.

5. *Realization of free-riding after time T^k.* A player who delayed or froze adoption for k's externality $c - c^k \geq 0$ can free-ride, connecting at charge c^k after k adopts. Let $W^k(V_i)$ denote the continuation value of the war that resumes after the event. As in the baseline war, adopting beats waiting if $V_i - c^k \geq \beta W^k(V_i)$. These sequential adopters have connection values $V_i < V^{\text{joint}}_{T^k}$, so that waiting one additional period is preferred to joining k.

6. *A new war at a higher baseline.* If neither i nor j specifically delayed adoption to receive the free-riding benefit, then a lower c^k replaces the original connection charge c, and a bilateral war resumes at lower cost. Holding $W(V)$ and c' fixed, let $V^{*,k}$ denote a new symmetric cutoff. As $c^k \leq c$, the new cutoff is weakly lower, $V^{*,k} \leq V^*$, and the new baseline adoption probability is weakly higher, $\sigma^{*,k} = 1 - F(V^{*,k}) \geq 1 - F(V^*) = \sigma^*$. In the resumed war, players adopt exactly when $V_i \geq V^{*,k}$.

Overall, when the cutoffs are ordered $V^{*,k} \leq V^{\text{joint}}_{T^k} \leq V^{\text{surv}}_{T^k} \leq V^*$, the phases sort connection values into regions by cutoff. Highest-value players $V_i \geq V^{\text{surv}}_{T^k}$ adopt unilaterally. High-value players in $[V^{\text{joint}}_{T^k}, V^{\text{surv}}_{T^k})$ adopt jointly with k at time T^k. Lower-value players in $[V^{*,k}, V^{\text{joint}}_{T^k})$ free-ride and connect after. Lowest-value players $V_i < V^{*,k}$ remain unconnected and resume the war.

Proof. See Appendix subsection B.5

Proposition 3 shows a heartbeat-like pattern around a neighbor externality. Graphically, adoptions decline steeply into the event, spike as joint adoption occurs, and regress to a higher baseline. I show an event study with this pattern in section 5. Empirically, it provides the basis for reduced-form predictions regarding neighbor externalities. I will also present a dyadic regression that shows the empirical magnitude of Proposition 1, and a cross-section that observes Proposition 2's price discrimination and selection. I now turn to typical connection costs and charges in electricity distribution, before turning to the case study.

3 — Electricity Connection Costs and Charges

Electric utilities participate in at least four markets. Just as the wholesale electricity market is split into transmission interconnections and electricity generation, retail electricity markets are split into distribution connections and electricity consumption. For brevity, I refer to the market for retail electricity connections as the infrastructure *connections market* and to retail electricity consumption as the *electrons market*. Because these bookends of flat fees and distance-based charges reflect the choice of whether to pass on supply-side costs, I describe the supply-side costs of the connections market here.

3.1. Physical Spatial Marginal Connection Costs

Table 2 — Connection materials and marginal costs increase with the distance to existing infrastructure

Distance	Description	Additional materials	Marginal cost
Within 30–50 m	Service drop	Drop, meter, wiring	$<$100 USD
50–200 m	Under-grid / within coverage	Several poles, lines	100–2,500 USD
200–600 m	Last mile / transformer communities	More poles, lines	2,500–10,000 USD
Further	Off-grid (new transformer required)	Transformer	$\gg$10,000 USD

Notes. Indicative costs for new marginal consumers at each distance bracket. Exact distances depend on utility choices such as pole heights, voltage choice, pole costs, exchange rates, and other factors. Terms are sourced from the existing energy and development economics literature: Blimpo and Cosgrove-Davies (2019) for coverage, Lee et al. (2016) for transformer communities and under-grid, and Chakravorty, Emerick, and Ravago (2016) for "last-mile."

The cost of marginally expanding the distribution grid is typically a nondecreasing, stepwise function of distance from the closest existing infrastructure.[18] The existing academic literature has used jargon such as "service drop," "under-grid," "last mile," and "off-grid" to denote different distance-cost ranges. I summarize these in Table 2.[19] There are two edge cases in which supply-side costs may not increase with distance. Existing infrastructure may reach capacity and require a non-marginal reinforcement to serve an additional customer. Alternatively, line losses may make a new separate grid the least-cost technical

18. Capacity constraints can generate additional jumps in marginal cost.

19. Throughout, "off-grid" as a distance category means a location far enough from existing infrastructure that connection requires a new distribution transformer, distinct from "off-grid providers" such as solar home systems and microgrids, a competing technology considered in the counterfactual analysis.

solution. Otherwise, marginal connection costs increase with the materials required that correspond to the prospective customer's distance from an access point.

Expanding the grid in a particular direction lowers the utility's supply costs of connecting consumers in that direction. Supply-side costs are poorly summarized by a stable one-dimensional marginal-cost or supply curve because costs depend on the grid state and direction of expansion. Zvoleff et al. (2009) attempts to model the supply curve by assuming the planner's least-cost sequence of connections. The resulting curve is nonlinear and poorly behaved, visually appearing as a random walk, because the next consumer in a least-cost sequence can be near or far from the existing network depending on spatial clustering. The intersection of demand and such a supply curve in one dimension is meaningless because ordering consumers by declining WTP would not match the supply-side geographic order. I use evolving isocost regions to provide a tractable description of the relationship between supply costs, distances, spatial configuration, and the state of the grid.

Figure 2 shows that isocost regions evolve as the distribution network expands. From the high-voltage backbone network, a substation steps the high voltage down to medium-voltage feeders, from which transformers step the voltage down further to low voltage.[20] When a utility installs a *transformer*, consumers within range fall into the yellow "last mile" where connections are feasible.[21] Electricity distribution networks expand with *network extensions* that emanate from the transformer. In an overground distribution network, the extensions are discrete because wires hang from poles spaced 30–60 meters apart. The distance between poles depends on terrain, obstacles, pole height, and how much wires droop. As network extensions emanate outward, consumers are marginally cheaper to connect. Ultimately, the *minimum supply cost* is reached when customers sit within a 30–50 meter *service drop* of a pole. At this point, the utility needs to drop a wire from the pole to a meter, from which the household wires its appliances. Service-drop coverage is represented by the date-colored regions surrounding poles in Figure 2.

3.2. Recovering Distribution Capital Costs

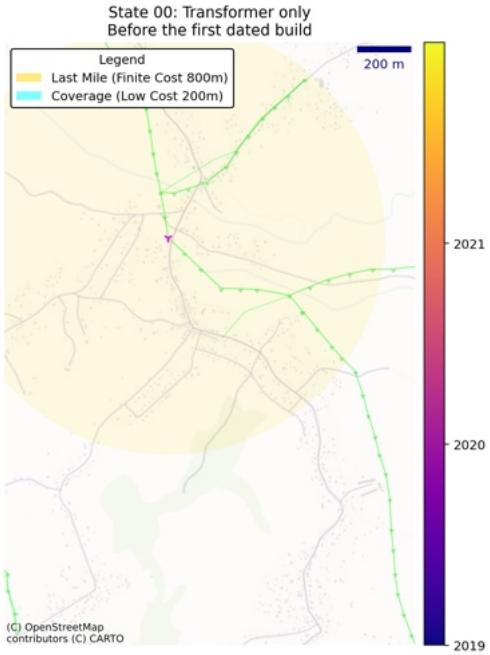
Utilities fund distribution network expansion through three channels: cross-subsidies from the electrons market, direct subsidies, and connection charges. If they are vertically integrated, they can "recover" revenue from the downstream electrons market to cross-subsidize "fixed costs" in the connections market. Direct subsidies from government finance, multilateral development grants, or concessional loans are typically allocated as a flat subsidy per consumer or piece of infrastructure. Connection charges cover the remainder; how distribution network capital costs are funded determines whether consumers pay for distribution capital upfront, through subsequent electricity consumption, or through future taxes.

Most countries outside Sub-Saharan Africa primarily subsidize connections with cross-subsidies from the electrons market. In most developed countries with built-out grids, connections are effectively free within a coverage area. Distribution network expansions are typically treated as systemwide fixed costs, recouped via "fixed cost recovery" in recurring retail tariffs (Joskow 2007; Borenstein 2016).[22] Such cost

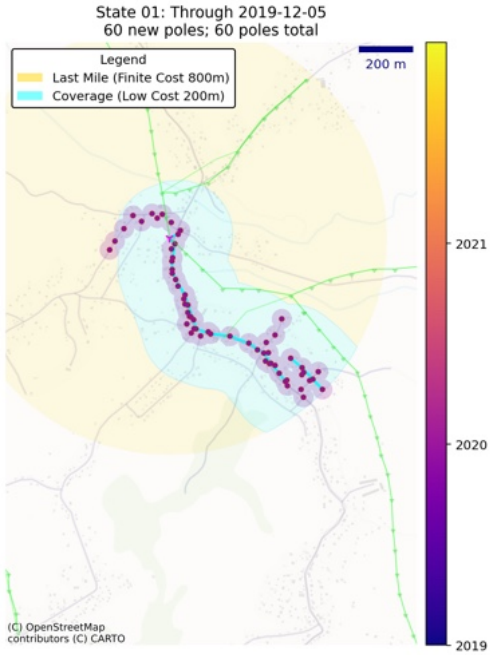
20. Development banks and external donors typically fund "backbone" transmission grids across Sub-Saharan Africa.

21. "Last mile" is a misnomer: a typical distribution network has a half-mile radius. By feasible, I mean without unilaterally paying for transmission infrastructure costs, which are several times higher than distribution ones per unit distance.

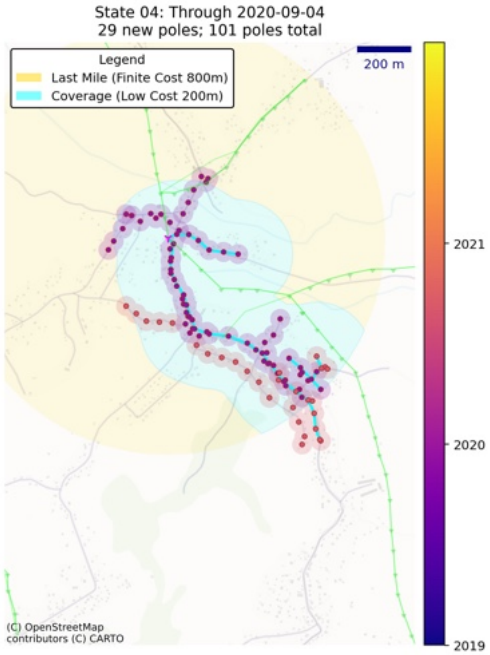
22. While out of scope for this paper, one could show that increasing fixed costs increases economies of scale, entrenching the



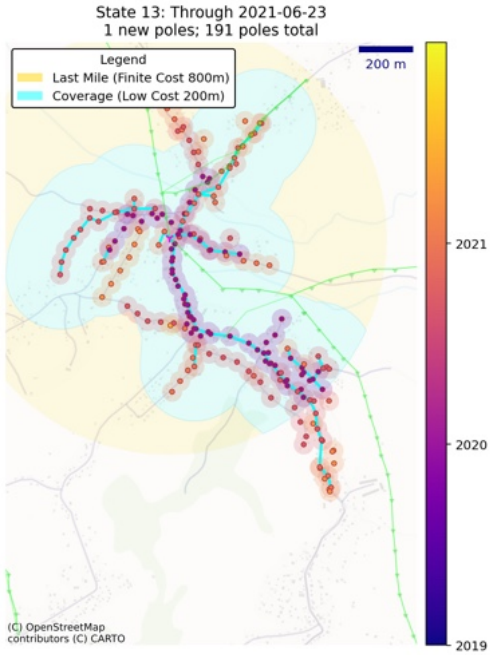
(a) Initial transformer state



(b) December 2019



(c) September 2020



(d) June 2021

Figure 2 — Spatial rollout and supply isocost regions, 2019–2021

Notes. The initial transformer-only state and three dated states of the sampled Gakoro Health Center distribution network in Musanze. Light yellow is the last-mile region; light blue is the under-grid region; the smaller date-colored regions show 37-meter service-drop coverage around poles. The dated panels show December 5, 2019; September 4, 2020; and June 23, 2021. Their expanding blue and date-colored regions show how successive extensions make new locations cheaper to connect. Data sources: Google Open Buildings, OpenStreetMap, CartoDB/Contextily, Rwanda Energy Group, and Rwanda Ministry of Infrastructure GIS portals.

recovery is not necessarily out of reach for all developing countries. Golumbeanu and Barnes (2013) describes how rural electrification cooperatives in Bangladesh provided loans to finance infrastructure and recovered the loans from consumers via electricity tariffs. Upfront, consumers only needed to pay a small service connection fee (USD 23). Most utilities worldwide forgo connection market revenue and collect it elsewhere.

However, Sub-Saharan African utilities have substantially less scope to recover distribution capital costs through electricity consumption. Trimble et al. (2016) reports that all but two Sub-Saharan African utilities charge consumption tariffs below the marginal cost of electricity generation. Blimpo, McRae, and Steinbuks (2018) shows that low and elastic electron demand forces utilities to charge high connection charges to generate revenue. It explains why connection charges tend to be higher in Sub-Saharan Africa than anywhere else in the world (Golumbeanu and Barnes 2013).

3.3. Connections Bundle Private and Public Infrastructure

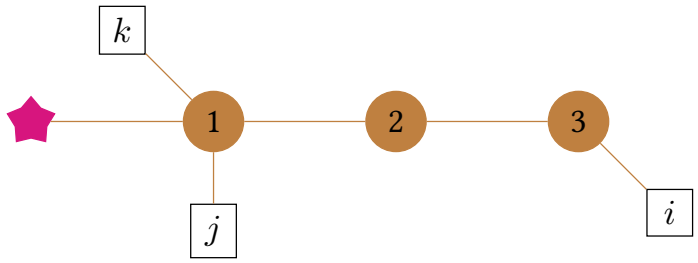
Electricity distribution networks, encompassing the infrastructure from the transformer to the consumer, are neither purely private nor purely public for three reasons. First, distribution infrastructure becomes congestible and rivalrous as it approaches capacity. Second, the utility can exclude users from service, so the network is technically a congestible club good. But excludability is limited in practice. Utilities typically have legal obligations and responsibilities to serve on a non-discriminatory basis under a universal service mandate. Further, once service has commenced, utilities are typically obligated to continue providing service and can exclude only consumers who do not pay their bills (Rossi 2000). Fundamentally, customers cannot exclude each other. While hypothetically one could exclude a neighbor by cutting their electricity line, laws prohibit this. So shared distribution infrastructure is effectively a public good: non-rival among its users and excludable by the utility, though not by customers against each other. Third, and most importantly for this paper, *upstreamness* determines how widely each specific component is shared. The most downstream meters and service drops are customer-specific private goods. The most upstream components—transformers, feeders, and transmission infrastructure—serve many users. The network components between these endpoints lie on a public good spectrum.

Public-ness depends on the distribution tree's network structure. Consider the following infrastructure network, where service emanates from the source through the brown junction nodes to users i, j, and k. Figure 3 shows that downstream poles are more private and less public because of the neighborhood's spatial configuration relative to a network access point. Users j and k require only the first junction node, but i requires two additional poles. Although all three nodes are non-rival once constructed, nodes 2 and 3 are *de facto* private to i because no other user's path includes them. Meanwhile, pole 1 is much more public. Heuristically, downstream components tend to be more exclusive, and upstream components will tend to be more widely shared.

The figure shows that flat fees cross-subsidize across space, and distance-based charges endogenously cross-subsidize across time. An average-cost flat fee divides the costs (1,1,1). Because i is the only customer using nodes 2 and 3, a flat fee would implicitly transfer part of i's marginal network extension cost to j

natural monopoly. Regulation becomes more essential; it raises regulated consumption tariffs, especially under rate-of-return regulation.

Figure 3 — Distribution infrastructure as a hierarchical public good



Notes. A single branch emanates from the closest point on the existing network (magenta star). Nearby users j and k attach at junction node 1, while reaching the distant user i requires nodes 1–3. Nodes 2 and 3 are *de facto* private to i because no other user's path includes them. In an electricity distribution network, the source may be a transformer or an existing pole, and the junction nodes represent poles.

and k. A materials-based charge instead assigns the cost of infrastructure components according to who connects first. If i connects first, i pays for all three nodes, and j and k then do not have to pay for any poles. If, by contrast, j or k connects first, the first mover pays for the first pole. The second-mover then does not need to provision a pole to connect, and the downstream customer i would still need to pay for poles 2 and 3. The order of a connection endogenously creates externalities and determines who pays.

3.4. Average-Cost and Marginal-Cost Network Connection Policies

This subsection compares which costs utilities pass through in distance-based charges, the rationales utilities and policymakers give for choosing average- or marginal-cost policies, institutional variation within distance-based connection charges, and accompanying policies or clauses that acknowledge the market failures studied in this paper. I compiled the connection policies and strategy documents, where available, from utilities and regulators across Sub-Saharan Africa and other developing countries. I describe qualitative findings below.

Regardless of policy, all Sub-Saharan connection charge policies charge consumers for private electricity meters. Except via community-led adoption policies, they do not charge grid electricity adopters for transmission, feeder, or transformer infrastructure.[23] The substantial disagreement was whether to charge consumers for low-voltage distribution network extensions. I found that about a third of countries pool distribution network infrastructure costs and charge an average-cost policy. In contrast, about half of countries pass these costs on to customers in a distance-based policy. The remainder use hybrid policies or apply different policies in different areas.[24]

I call a policy an *average-cost policy* or *flat-fee policy* if supply-side distribution network connection costs are not passed on to consumers through connection charges. In Lee, Miguel, and Wolfram (2020b), for example, all consumers by default would pay a flat rate of USD 398 at baseline, as long as they were within 600 meters of a distribution transformer. Today, Kenya uses a flat USD 120 connection charge. Other countries that use a flat-fee connection charge include Botswana, Morocco, Vietnam, Ecuador, Guyana, Angola, and Mozambique. To set an average-cost policy, utilities choose an exogenous "penetration rate"

23. Funding from multilateral banks typically targets transmission infrastructure (Blimpo and Cosgrove-Davies 2019).
24. I will provide a full appendix in a future version of the paper.

constraint in their least-cost planning algorithm, subtract the share of cross- or direct capital subsidies, and then pool the remaining cost among the expected adopters (Parshall et al. 2009). A significant limiting factor for utility planning is demand estimates over space, which allow the planner to spatially allocate infrastructure (Lee et al. 2021). The clearest economic arguments for a flat fee are lower administrative costs and ex ante fairness through a single market price.

Distance- or materials-based charges pass through supply-side connection costs to the consumer. In policy documents, the main rationale given for distance-based policies is to make connections "cost-reflective." Other policy design concerns include simplicity, transparency, consistency, transaction costs, and early network expansion (Woods 2019; Zambia Electricity Regulation Board 2023). Cost reflectivity decentralizes decisions about whether and where the network expands to customers. As remote customers see higher costs, they choose whether it is worth it. Hence, the spatial extent of distribution networks becomes a function of revealed preference and should improve infrastructure allocation.

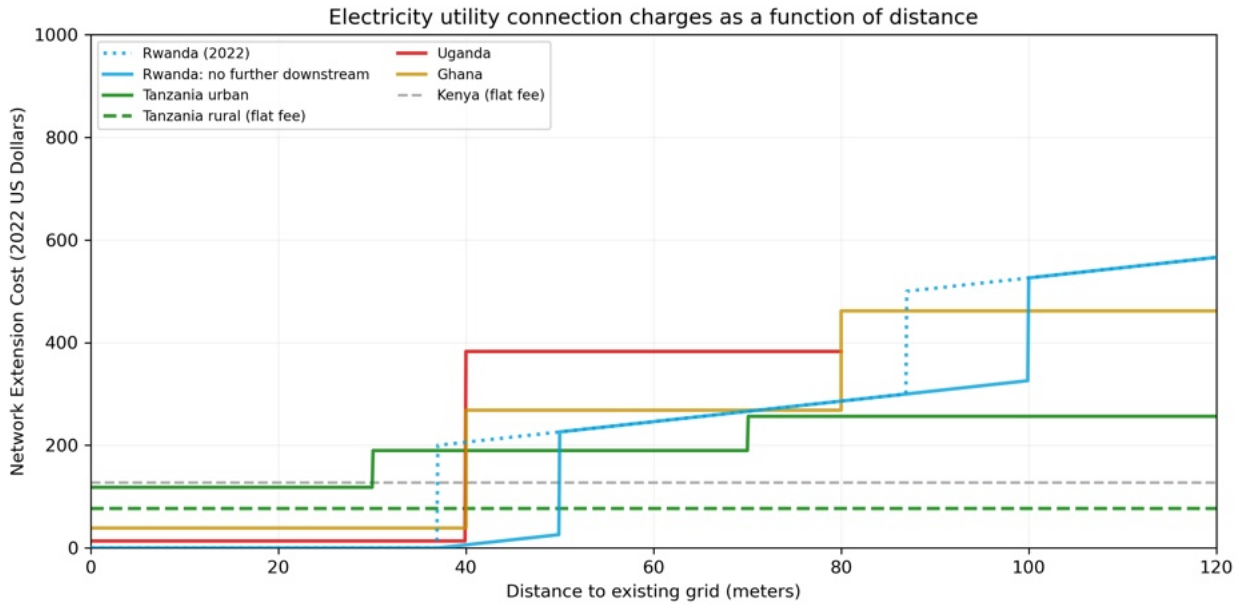
I plot connection charges as a function of distance for selected countries in Figure 4. The visual takeaway is that distance-based charges tend to be piecewise step functions. They start from a low value within the service-drop range, covering the cost of meters, wiring, and service-drop cables. Discrete jumps occur when additional poles are necessary to extend the grid, with hop size depending on material costs and per-pole subsidies. Unlike other countries, Rwanda's policy slopes upward because it charges consumers two dollars per meter of line starting from the service drop.

I observed three additional stylized facts about distance-based charge design when constructing Figure 4. First, published distance-based connection-charge schedules were not the norm. Instead, most countries depend on an ad hoc engineer's survey of the required network extension. Second, how charges are calculated may depend on distance or materials. Ghana and Uganda index the cost to the number of poles, whereas Rwanda and urban Tanzania index it to distance. Finally, all cost-distance functions found on utility websites were piecewise linear, even though utilities theoretically could smooth them while retaining some degree of cost reflectivity. I found no academic literature on these design tradeoffs, but policy and consulting documents divulged the following.

Why do some countries require quotations while others publish distance- or pole-based schedules? The apparent advantage of quotations is that they pass construction-cost risk to consumers. Quoting each network extension separately allows the utility to pass through any heterogeneous supply-side costs from transportation, terrain, procurement, eminent-domain, or wayleave costs to new consumers (Woods 2019). By contrast, a schedule offers transparency, allowing consumers to make connection decisions without requesting a utility quotation.

Why do some utilities that pass on costs base connection charges on materials, while others base them on distance? Policy documents highlight that utilities should charge customers for the *minimum technical solution* needed to connect the household, not the real infrastructure that the utility builds (Woods 2019; Zambia Electricity Regulation Board 2023). Utilities have incentives to place poles that follow a least-cost electrification plan that minimizes the total cost of providing electricity to the entire neighborhood, rather than to connect the adopting consumer at their least cost. Using a distance-based charge over a material-based charge ensures that utility decisions that warp the grid towards neighbors are not passed through to consumers.

Figure 4 — Connection charges in Sub-Saharan African countries with transparent distance-based connection schedules.



Notes. Connection charge policies collected from utility websites in 2022, converted to USD using calendar-year average exchange rates from the World Bank API (PA.NUS.FCRF). For Ghana and Uganda, the geographic component is an explicit function of poles required rather than distance; I assume these are required at 40-meter intervals. Where credit loans were available directly from the utility, the above shows the upfront minimum charge. Kenya and rural Tanzania use flat fees, shown for comparison. Rwanda's policy is partially transparent online, with distance categories and materials denoted but not exact prices. I supplemented this with information from the Rwandan Ministry of Infrastructure, which shared that the utility charges consumers for the extension above 37 meters at $2 per meter of line and $200 per pole. Whether a pole is needed within 50 meters depends on whether there are downstream customers. Additional Rwandan contextual and cost details are provided in subsection 3.5.

Why do utilities use stepwise rather than smoothed connection charge schedules? Utilities face a substantial political economy challenge when using distance-based policies because they must justify to consumers why different consumers should pay different connection charges for the same service (Woods 2019). Using a stepwise function makes material costs tangible, promotes transparency, and therefore helps justify connection costs to consumers. Perceived fairness would seemt to be an important constraint on consumers' private provision of shared infrastructure.

I now describe how institutional choices are implemented in Rwanda, the setting for my empirical case study.

3.5. Rwanda's Connection Market

In 2009, only 6% of Rwandan households had access to on-grid electricity. Substantial electrification has occurred since then; REG recorded 55.9% on-grid access by June 2023, and the World Bank's World Development Indicators show 74% electricity access (including off-grid technologies) at the end of 2023. I briefly review Rwanda's distance-based policy design and other pertinent institutional features, focusing exclusively on policy and institutional facts used to calibrate connection charges from the data, motivate empirical design choices, or implicitly nod at strategic incentives.

The 2017 "New Connection Policy" codified Rwanda's distance-based connection policy, though it used a distance-based charge informally before that (Rwanda Energy Group (REG) 2017). The online connection process portal shows that distance boundaries are transparent to new customers but describes the specific charge as the "cost of extension." It indicates a four-working-day response time for connections within 100 meters of the grid and a ten-working-day response time beyond 100 meters. I use these service standards to motivate an otherwise unobserved weekly delay between connection request and installation in section 4. The discussion below draws from the "New Connection Policy" unless otherwise noted.

The baseline minimum charge, which the utility calls the "Connection Fee," is 56,000 Rwandan francs (RWF) or USD 56.[25] It covers the cost of a meter, service cable, and "connection accessories." However, the effective minimum upfront connection charge is zero (Dabalen et al. 2025), because customers can take up to RWF 56,000 in credit, repaying it at a 100% markup on their electricity-consumption tariff with zero interest. However, the credit only applies to the cost of meters, cables, and connection accessories, not poles or lines, so upfront connection charges are zero only if the prospective customer is within a 37 meter service drop from an existing pole. From 37 to 50 m, customers pay for the "cost of extension" but are nevertheless connected "with existing infrastructure." Rwanda's Ministry of Infrastructure reported that this is a 2,000 RWF/m electricity-line charge. When it would help downstream customers, the utility gives a quotation for an electricity pole within this range. This discretion is explicit in Figure 4.

From 50 to 100 m, the utility passes through part of the supply cost of one pole. Poles are spaced 55 m apart in engineering reticulation standards (Rwanda Energy Group (REG) 2020). In the cleaned geospatial data I describe in the next section, I observed a median pole-to-pole span of 49 m. REG places poles to follow roads, engineering standards, local geography, and anticipated demand. Roads are followed to ensure

25. I study the period from October 2018 to November 2021. In mid-2022, the exchange rate was exactly 1,000 RWF to 1 USD. The Rwandan franc depreciated during this period, and exchange-rate exposure is the basis of my secondary shift-share instrument.

that poles are placed on public land. The connection charge is based on distance rather than materials, so customers pay for the minimum technical solution. The policy emphasizes that customers "will not stop new applicants from being connected on such extension lines," indicating that the utility is aware that it is asking customers to privately provision public goods.

Beyond 100 m, the policy states that "For the applicant or a *group* of applicants beyond 100 meters from the point of electricity supply, a detailed survey will be conducted to determine the most appropriate extension line and point of connection, and a related bill of quantities will be produced." The emphasis on *group* is mine. Consistent with Proposition 1, it suggests that the utility is accustomed to cooperation when the costs of shared infrastructure are extremely large relative to income. Two poles cost 400 dollars, and Rwandan annual income in 2022 was 935.9 (in constant 2015 USD per capita); five poles are more than annual income, and so will generally require cooperative provision. The policy does not prescribe nor suggest how costs might be shared within the group. The last-mile region, indicating a finite cost for customers, is also described by the policy when it says that "the Low Voltage line cannot be extended beyond 800 meters from the Distribution transformer." This gives me geographic bounds for which potential adopters are relevant when transformer placement starts the network game.

Except for the service drop line and meters, Rwanda's distance-based charge is almost the mechanical distance-based policy described in Definition 1. About a third of the supply-side connection costs gets passed through to connection charges. Rwanda's Ministry of Infrastructure reported in 2022 that customers are charged USD 200 (200,000 RWF) for a pole, in addition to the aforementioned line charge of 2,000 RWF per meter. I consolidate documentary evidence I collected reflecting supply-side costs, demand-side connection charges, and subsidies in Appendix A. I found that supply-side connection costs for a pole fluctuate between USD 600 and USD 1,000 over time, and that USD 6 per meter for lines appears to match aggregate financial reports. I assume a one-third pass-through rate or a two-thirds subsidy on supply-side connection costs when calibrating the structural model and evaluating welfare in my counterfactual simulation.

Consumer requests primarily drive grid extension, except when the utility subsidizes "Massive Connections" funded by the Electricity Access Rollout Program (EARP). The EARP pools financial resources from the World Bank and other multilateral development agencies to fund specific extension projects confined to specific localities. Typically, the EARP targets schools and other public infrastructure and electrifies the entire route from transformer to public service. Customers in areas funded by massive connections pay the fill-in upfront rate of zero. These massive connections are largely not observable except by assumption, though I hope to receive more information from the Ministry of Infrastructure or the World Bank about this for a future draft.

Information provision helps customers understand the connection policy (Rwanda Energy Group (REG) 2017). The utility disseminates the policy through markets and district offices, radio and television broadcasts, religious ceremonies, and Umuganda. Umuganda is a monthly required community service that brings together a representative from each household at the village level to do a day of community service each month. These monthly meetings may decrease coordination costs in Rwanda. If cooperation costs are more about coordination than baseline cohesion, this would predict more strategic complements and less strategic substitutes.

Via the Rwandan Least-Cost Electrification Plan, Rwanda effectively spatially delineates off-grid and on-

grid regions to private mini-grid providers and the utility. While not an explicit policy, licenses for off-grid solar providers are limited to designated off-grid zones. This spatial delineation coordinates monopolistic on-grid and off-grid regions, diffusing electrification efforts over space. Therefore, the empirical exercises reflect a context with little competition between the grid and alternative sources of electricity.

4 — Data

I transform timestamped geospatial infrastructure data from Rwanda Energy Group and Rwanda's Ministry of Infrastructure into a prospective customer-week panel of distances, costs, externalities, and peer effects. The pipeline has three stages: geospatial cleaning, a least-cost planning algorithm, and a dynamic grid rollout algorithm. The geospatial cleaning removes duplicated transformers, poles, lines, and customers, matches each utility customer meter to a remotely sensed building polygon, and links them into a network graph. The least-cost planning algorithm adds connection paths for the remaining buildings in the vicinity of a distribution network. Non-adoption is a revealed preference, so all remotely sensed buildings are included in the reduced-form panel and are players in the structural model. The dynamic grid rollout algorithm then constructs a panel from customer installation dates by chronologically adding customer-induced network expansions.

4.1. Raw Data

The primary raw data are from Rwanda Energy Group (REG)'s and the Rwandan Ministry of Infrastructure's Geographic Information Systems (GIS) Portals (Rwanda Energy Group (REG) and ESRI 2022; Ministry of Infrastructure (MININFRA) Rwanda and ESRI 2024).[26] Secondary raw data come from Google Open Buildings, OpenStreetMap, the Rwandan Land Use Authority, the World Bank Multi-Tier Framework Survey, and the Global Economic Monitor. I describe each of the datasets and their uses:

- **REG Customer GIS data.** Two shapefiles of customer *meter* geolocations and associated metadata were available on Rwanda Energy Group's GIS portal. Each point has geospatial coordinates, administrative plot numbers, customer characteristics, meter serial numbers, and transformer name. Connection dates were available only in the substantially smaller file, which covers 4.96% of unique metered customers after deduplication (34,087 of 686,891). I limit the analysis to the 176 distribution networks where the smaller file accounted for more than 80% of installed meters. The customer data appear to be entered when the utility fulfills a customer's connection request and a technician installs a meter.
- **REG LV network GIS data.** I use the GPS locations and characteristics of LV poles, lines, and transformers to construct each distribution network's initial graph and train the least-cost planning algorithm. Across the datasets, the infrastructure metadata includes the primary and secondary voltages and currents, capacities, materials, years of manufacture, manufacturer, lines, upstream substation and feeder, phases, utility branches, and more. The infrastructure data are collected by survey, and the data-generating processes for LV lines and poles are related. A surveyor can either create a dataset

26. Each World Bank funding round has targeted several million dollars of grant funding at geospatial data collection, as this helps with utility planning and modeling.

of poles and sequence them in a collection to form a line or collect a line directly, with the phone or tablet recording location every few seconds as the surveyor walks along the infrastructure. I use the combined universe of line and pole data to construct the infrastructure network. The infrastructure data cover the near-universe of electricity infrastructure across Rwanda in a snapshot that lags behind the customer data.

- **REG Branch boundaries** are not perfectly aligned with administrative district boundaries. Branch membership is used to construct branch-by-week fixed effects that control for correlated shocks in the reduced-form regressions.

- **Google Open Buildings** provides a temporal "2.5D" raster and polygon shapefile. The .5D refers to building heights (Sirko et al. 2023). I spatially join the raster to polygons via polygon centroids to observe building heights and areas and create a binary panel of construction dates at yearly frequency.[27] A bipartite minimum-cost matching algorithm matches each meter to a building polygon within 50 meters using point-to-polygon distance (Virtanen et al. 2020). A large 50-meter range accommodates the rare cases where the utility places meters on electricity poles rather than buildings to deter theft. Matched meters and buildings are 3.7 meters apart on average, with 1.25% of matches above 37 meters. I use each building's "Google plus code" as a customer identifier. Matched buildings are adopters, and unmatched buildings are non-adopters.

- **The Rwandan Land Use Authority** provides geographic terrain data. I use rivers and a 10 m digital elevation model (DEM) (UNDP 2023). I construct terrain ruggedness indices and slope measures from the DEM. I use terrain ruggedness as a customer characteristic. Slope and waterways enter the cost surface of the least-cost planning algorithm.

- **OpenStreetMap roads.** I use roads primarily in the least-cost planning algorithm because the utility builds poles along roads to keep them on public land (OpenStreetMap Contributors 2026).[28] I use distance to the nearest road and the nearest road class as customer characteristics in the structural model.

- **Earth Observation Group Nightlights.** I use 2019 Visible Infrared Imaging Radiometer Suite (VIIRS) nighttime lights as consumer characteristics in the structural model and as a contextual control in the reduced form (Elvidge et al. 2021).

- **Exchange-rate data from World Bank Global Economic Monitor.** I construct a monthly imported-pole cost shifter for each utility branch using official exchange rates from the World Bank Global Economic Monitor. I interact exchange rates lagged by 18 months with each branch's pre-2019 shares of LV poles by manufacturer country to construct a shift-share instrument. I motivate the lag in section 6.

- **The World Bank Multi-Tier Framework (MTF) Surveys** conducted by the World Bank in Rwanda in 2016 and 2022 are used to check the distributions of connection charges paid in Appendix A and

27. I do this by observing when heights and precision monotonically increase, then plateau.

28. OpenStreetMap provides substantially better coverage of Rwandan roads than Rwandan government sources.

assess price discrimination with alternative data. Survey respondents reported their stated willingness to pay and revealed payments for a connection to the grid, alongside many other variables on electricity demand, appliances, and alternatives (World Bank 2018a, 2022).

4.2. Geospatial Graph Construction

The first part of the pipeline converts raw geospatial data into a graphical representation and data structure (Hagberg, Swart, and Schult 2007). The main goal is to encode the lines as edges and poles, customers, and transformers as nodes, satisfying Definition 4.

Definition 4 (Physical Hierarchical Network)**.** Electricity distribution networks are hierarchical. If a low-voltage pole or consumer is electrified, it must have a physically installed and electrified network path connecting the transformer to it.

Consumers require a transmission path from a generator to the distribution transformer and an LV path from the transformer to their premises to be connected to the grid. By hierarchical, I mean that there is a tier for transformers, poles, and meters: on their step-down side, transformers connect only to poles. Poles can connect upstream to transformers, connect horizontally to other poles, or connect downstream via service drops to meters. But meters do not connect to each other, nor do transformers.

I clean the geospatial data to create a network graph for each distribution network that fits Definition 4. I remove duplicates, cluster points, impute missing line data from pole data and pole data from line data, and check network validity. I use methods similar to those electrical engineers use to prepare infrastructure networks for power-flow simulations. Specifically, I clean metadata by normalizing null strings, ensuring consistent engineering units, and standardizing categories. I clean geospatial coordinates by harmonizing coordinate reference systems (CRS), dropping invalid geometries for all points in the infrastructure data, and substantially deduplicating them.[29]

Deduplication is challenging because poles are often recorded multiple times by different surveyors and because floating-point errors propagate through CRS conversions; the Ministry of Infrastructure and Rwanda Energy Group had versions of the data in different CRS projections. I cluster low-voltage points within 7 meters and medium-voltage points within 12 meters, using the resulting centroids as representative points. To speed up lookup, I spatially index each deduplicated point on a hexagonal grid using H3, Uber's geospatial indexing library.[30] I spatially index lines as a sequence of H3 indices. Then I add each pole, transformer, and meter as a node to the graph, and low-voltage lines as edges between nodes.

Some infrastructure data are missing but necessary to ensure Definition 4. Where line data exist but pole data do not, I impute poles from points along the line, accounting for kinks and the distance between poles.[31] Where pole data exist but line data do not, I greedily create synthetic edges between unconnected poles, then prune them while preserving each pole's shortest path to the transformer. Service-drops are

29. I use EPSG:3857 (Web Mercator) throughout. This is a projected coordinate reference system where longitude-latitude coordinates are expressed in meters. As Rwanda is close to the equator, at approximately 2° south, WebMercator distances are negligibly inflated by 0.06% to 0.14% across the country.

30. Hexagonal cells perform better than a square grid because they ensure equal distances to every neighboring cell.

31. The data structure of a line is just a sequence of points.

not observed as data, so as the last stage of graph construction, I link customer meters using a waterfall approach. I connect each building to the closest pole within 10 meters, then to the most upstream pole within a 37-meter buffer, and then use a maximum distance of 50 meters.

After geospatial cleaning, I observe that more than 93% of low-voltage distribution networks outside Kigali are radial trees. I convert the remainder into trees by pruning cycles while preferring transformer-pole shortest paths. I check that each meter has a network path to the distribution transformer. For each matched or unmatched building, I observe nightlights, elevation, terrain ruggedness, building height, building area, and construction year. For matched buildings, I also observe residential, commercial, industrial, or public-service categories, as well as installation dates. The resulting graph has transformers, poles, and meters as nodes, and LV lines and service drops as edges.

4.3. Least-Cost Planning Algorithm

I develop and calibrate a least-cost planning (LCP) algorithm that constructs a network path to each unconnected building within range of the transformer and existing network. It plays several roles in my analysis: First, it allows me to touch up the infrastructure data for customers without a viable observed LV network path because of lags in infrastructure data collection. Second, it allows me to construct the minimum technical solution for each customer at every point in time. The corresponding connection charges are what customers should actually pay under Rwanda's connection charge policy. Third, I use it to extend the observed network to all non-adopters in the sample, simulating the utility's least-cost plan from the existing network. This only affects the structural model, where the least-cost plan is the state space. Finally, I also use it in a demand-stripped way for robustness checks and to consider an alternative instrument in subsection 6.4.

The least-cost planning algorithm's objective is to connect *demand points* to the existing LV network feasibly and at minimum cost. Essentially, it is a wrapper around a Steiner tree algorithm. Relative to a minimum spanning tree, a Steiner tree additionally allows for the construction of Steiner points — poles — that can connect demand points in space to the existing grid. In addition to material costs, the wrapper adds cost surfaces that account for terrain and engineering parameters, including reticulation standards. It takes three primary inputs: the network graph of existing infrastructure, which at minimum consists of a transformer; the demand points, which are buildings and/or meters to be connected; and a set of geographic features and parameters. The algorithm outputs a geospatial network graph, extending from the existing network, that represents the least-cost infrastructure necessary to connect the input demand points.

I trained the algorithm to simulate how Rwanda's realized grid behaves. I developed it iteratively. In a manual outer loop, I tinker with functional forms of the cost surfaces and network by using the algorithm to construct a synthetic network from each transformer to the meters it connects. I visually compare the synthetic network against the real network, satellite backdrops, and various inputs to the cost surfaces, assessing whether the algorithm's output geospatial infrastructure qualitatively matches how the real distribution infrastructure is organized around terrain, other infrastructure, and buildings.

In the current version of the LCP, I assume poles cost 600 USD, and lines cost 6 USD/meter on the supply side, following subsection 3.5. Feasibility constraints from the utility's reticulation standards are softened to a spline around reticulation-standard thresholds. I also embed the catenary equation to measure

clearances and impose a cost based on the clearance of hanging lines over terrain. For terrain, two H3 hexagonal cost surfaces are assembled from flexible spline and kernel functions over geographic features: roads, slope, waterways, and (optionally) buildings. Splines measure cost; kernels measure ideal setbacks. Separate cost surfaces are necessary because line and pole costs vary across space; for example, lines can cross waterways, but poles may not be constructed in them.

In an inner loop, I use gradient descent to optimize how parameters interact with geographic features or engineering standards to construct the cost surface and reticulation parameters. The objective compares the synthetic network against the ground truth by minimizing the root-mean-squared error between bipartitely matched poles, with a penalty for mismatches. Across distribution networks, the least-cost planning algorithm results in a 1.00 (median 0.93) pole-count ratio and a root-mean-square distance of 45 (median 39) meters between bipartitely matched poles within each transformer. The algorithm is imperfect, but it creates a least-cost plan that replicates the average number of poles needed for each network.

After training it on real data, I use the least-cost planning (LCP) algorithm to extend the grid to nearby buildings to create a least-cost plan for non-adopters. I include an unmetered building as a potential electricity adopter only if it lies within 800 meters of its assigned transformer, following the utility's service-radius policy described in subsection 3.5, and if the building is within 150 meters of that transformer's observed grid, which implicitly assumes that far-off customers have no strategic effects. I assign each unmetered building to the transformer associated with its nearest observed-infrastructure node, and remove buildings on multiple networks by assuming they participate in the network game where they have the shortest building-to-transformer distance. Non-adopters become part of the game and estimation panels via Assumption 1.

Assumption 1 (Potential Customers and Network Assignment)**.** Each remotely sensed building footprint is a potential electricity connection location, assumed to be a player in the network game that treats its assigned transformer network as the one through which it would receive a grid connection. Matched buildings adopt, but unmatched buildings are treated as non-adopters through the end of the sample period. Adoption and non-adoption are both assumed to be revealed preferences.

This assumption determines the size of networks and assumes every remotely sensed building could potentially adopt electricity and can provision externalities.

4.4. Date Imputation Algorithm

I convert the static network graph into a panel by iterating through customer adoptions that induce network extensions. Assumption 2 is based on utility operations and allows me to impute each pole's latest possible installation date from customer adoptions.

Assumption 2 (Consumer-Driven Monotone Grid Expansions)**.** When not driven by EARP massive connections, low-voltage distribution infrastructure within each transformer network is constructed in response to customer connection requests. Once installed and electrified, infrastructure remains active, so the grid monotonically expands.

To verify Assumption 2, I quantify the number of raw poles that are not on a network path from a transformer to a consumer. I find that 6.9% of the poles exist without serving a consumer, but of these, 62% have a

public service en route, indicating that the EARP likely funded them. So only 2.6% of poles detectably break Assumption 2 at endline. Given that the geospatial data are a snapshot of the network at the end of 2021, I have no reason to believe this would be higher within the panel period. When Assumption 2 breaks because the physical infrastructure predates the observed installation of the customer assigned to induce it, it biases estimates of strategic behavior to zero; I explain this at the end of the pipeline. Nevertheless, Assumption 2 is critical because it tells me which utility customer or group of customers induced each en-route pole, allowing me to assess who provisions infrastructure, and when each pole existed, allowing me to recompute connection charges and observe who benefits from it.

The last assumption is a convention that links connection requests to observed installations. This is necessary because the GIS data indicate installation dates, and I do not observe when customers make requests. Weekly periods provide a natural break because utility policy provides connections within four working days for consumers within 100 meters (2 poles) of the grid.[32] Within weeks, adoptions are also clustered: among transformer-weeks with multiple adopters, 73.6% of adopters connect on the same day. Most installations are on weekdays, and Sundays only account for 3.8% of installations. Therefore, I make Assumption 3.

Assumption 3 (Week-Ahead Connection Requests)**.** For network extensions that occur at t, I assume they were induced by connection decisions at $t-1$, and enter the grid state at the beginning of week t. I treat customers whose installations share infrastructure and occur in the same week as a coalition. For structural accounting, I assume the coalition shares the total bill using Shapley cost shares.

Assumption 3 rounds the four-working-day lag to five working days, or one weekly period. Given a four-working-day response time, customers that connect within the same week must have submitted an extension request just before the utility constructed the infrastructure. If they did so simultaneously, either they coordinated adoption, or the utility must batch customers because of scale economies. A placebo will rule out scale economies.

The date imputation algorithm applies Definition 4, Assumption 2, and Assumption 3 to estimate when each piece of infrastructure was constructed. It follows the following steps:

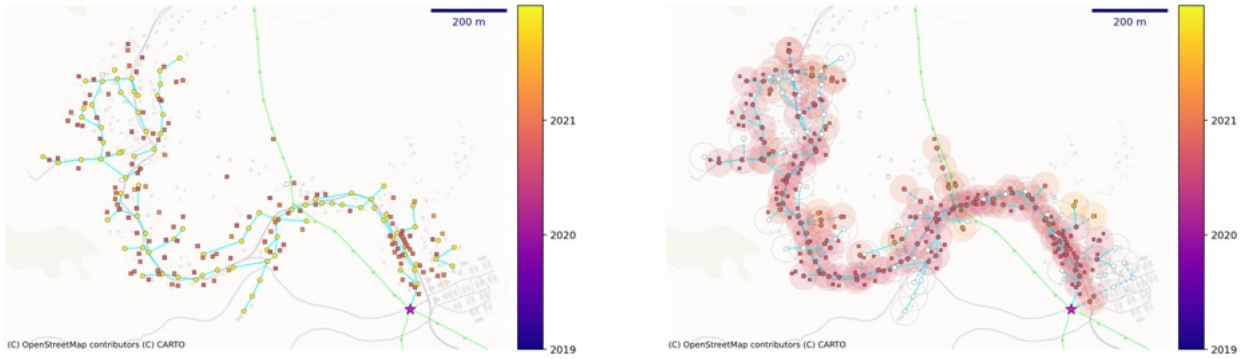
1. Initialize the grid from the transformer.
2. Iterate through each week in chronological order. If there is a network extension,
    i. Select the poles and edges necessary for each coalition of adopters.
    ii. Assign the coalition's minimum installation date to its induced network path infrastructure.
    iii. Save the set of infrastructure that existed at each discrete point in time as a grid state.
3. At each state, measure all consumer adoption incentives and instruments.

Figure 5 shows how the algorithm proceeds. In the left panel, undated poles are colored by end-line date. The uncolored meters lack raw dates. Not all poles were connected to the network in the raw data.

32. For consumers more than 100 meters from the grid, the standard is ten working days.

Geospatial cleaning connects dangling networks; for example, it reaches the poles in the bottom right that were initially unconnected by lines. Least cost planning extends the grid, for example, to reach the customer in the middle living under the LV line. The grid rollout algorithm imputes pole dates. In this grid, most of the infrastructure was built concurrently, but many more customers around the periphery connected and induced poles later.

Figure 5 — Raw inputs and date-imputation output for a sampled distribution network



(a) Raw observed data **(b) Completed graph and imputed dates**

Notes. Comparison of raw and algorithm-completed data for the Nkomero distribution network in Burera. The left panel shows observed low-voltage lines and poles, potential buildings, raw pole record dates, and observed meter-connection dates. The right panel adds the completed tree topology and colors poles by their imputed activation dates. Translucent circles show the 37-meter service-drop region around each dated pole. Squares are meters, circles are poles, colored by their installation date. The magenta star is the transformer; solid cyan lines are observed low-voltage paths, dashed cyan lines are completed paths, and green lines and triangles are medium-voltage infrastructure.

The date imputation algorithm mechanically creates one-sided systematic measurement error. The imputed date for each pole is the latest possible installation date; it represents the first time I observe the pole to be used by customers. So connection charges c_{it} are overstated early in the network rollout, and mechanically shrink as imputed dates catch up to reality. It biases estimates in predictable directions. When estimating the effect of static incentives, this shrinks the effect of static costs toward zero. But it biases tests of strategic incentives against the model's predictions. Consider a cluster of consumers who share poles that physically exist but have not yet entered the imputed grid state. In truth, everyone is close to the grid, and so true externalities are nil. The estimated exposure to externalities is large for whoever connects first and zero after that. This generates a spurious positive bias between future externalities and adoption, against my hypothesis that future externalities would delay adoption. It also generates a spurious negative bias against cooperation, as I see phantom unbuilt shared infrastructure that provides no incentive to cooperate. When the first consumer in the cluster connects, I see an externality that should be followed by sequential adoption if free-riding exists. Without this externality, there would be no incentive to adopt sequentially after the first mover.[33] All these effects run opposite to my hypotheses: anticipation should delay adoption, cooperation should increase it, and free-riding sequential adoption should follow realized externalities.

33. I control for peer effects.

4.5. *Sample Selection*

I select a sample because of data quality concerns in the raw data, geospatial cleaning, LCP, and date-imputation pipeline. The raw data include 687,065 utility customer records over 7,079 total transformers (including HV-MV and step-up transformers); the current reconstruction algorithm retains 7,105 customers, of which 6,992 have dated connection events across 177 transformer networks. The limiting factor for sample selection is that connection dates were available only in the smaller of the two consumer files when I accessed the data in March 2022: 34,125 records carry a connection date, or 34,087 unique metered customers after deduplication—only 4.96% of the 686,891. The game is unobservable for customers without connection dates because it is dynamic. I report how many consumers and networks I exclude at each step in Table 3.

Table 3 — Data pipeline and sample-selection

Step	Networks dropped	Buildings	Metered customers	Dated customers	Networks remaining
Raw meter records			*687,065*	*34,125*	*7,079*
1. Duplicate records			174	38	7,079
2. Failed network assembly	508		60,690	2,539	6,571
3. Transformer filtering and unmatched meters	4,137		289,321	0	2,434
Reconstructed distribution networks		*951,047*	*336,880*	*31,548*	*2,434*
1. Kigali	359	208,818	84,153	1,768	2,075
2. Date coverage $< 80\%$	1,629	652,346	239,584	16,928	446
3. Outside height-raster footprint	8	396	362	358	438
4. Unclear institutional details	14	4,486	1,489	1,409	424
5. Building beyond 1,200 m	18	7,192	1,327	1,259	406
6. Fewer than three consumers	120	17,482	163	163	286
7. Rollout window $<$ 1 week	109	18,040	2,697	2,671	177
Selected sample		*42,287*	*7,105*	*6,992*	*177*

Notes. Italicized rows report totals at that stage of the pipeline, and numbered rows report counts dropped at that step.

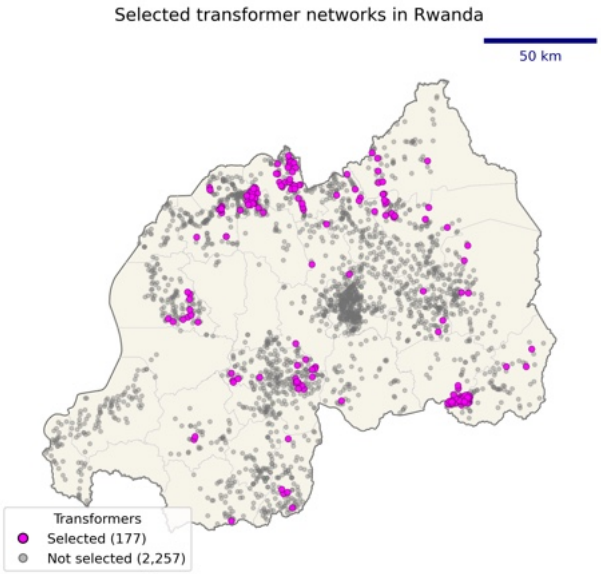
The top half of the table reports excluded data based on raw geospatial data quality. In the first row, I remove duplicate customer records (174 records, 38 of them dated) across the two datasets, leaving 686,891 unique metered customers. In the second row, I drop the 508 network components that fail assembly: components without any customer data or without any pole or line data.[34] Although the LCP algorithm could construct networks for customer-only clusters, I keep these out of my sample to minimize the LCP algorithm's impact on the reduced-form results. This drops 60,690 customers, 2,539 of them dated. This row also includes networks with multiple transformers that I could not separate. In the third row, I deduplicate transformers via spatial clustering, filter out step-up and transmission network transformers, remove unmatched meters, and drop distribution networks that have no customers with installation date data.

34. The transformer dataset is not distribution-specific; it includes transmission network transformers.

Together, these filters drop 4,137 of the 6,571 network components and 289,321 metered customers.[35]

At this point, there are 336,880 observed customer meters, 31,548 of which are dated, and 2,434 transformer networks. I join to Google buildings, and under Assumption 1, there are 951,047 buildings in the vicinity of the reconstructed distribution networks. First, I drop Kigali to focus on rural electrification. Second, I drop distribution networks with less than 80 percent date coverage among metered customers. This is a critical step, as I would not know who induced which pieces of infrastructure or who provided or received externalities. This drops 1,629 distribution networks.

Figure 6 — Sample distribution networks relative to all LV transformers



Third, I drop eight networks outside the footprint of Google's 2.5D building-height raster. Fourth, I drop fourteen transformer networks with unclear institutional details: eleven with undated public-service customers, one refugee camp, and two resettlement zones. As previously discussed, the EARP fully subsidizes health, education, and local administration public services, and electrifies everyone en route (World Bank Independent Evaluation Group 2019). Because customers on these networks do not pay the full connection charge, their connection decisions may not reveal their willingness to pay the charges implied by distance-based policies. Further, if these connections induced poles earlier than observed, this would substantially inflate the systematic measurement bias discussed at the end of the last subsection. The same applies to refugee camps and resettlement zones, which may reflect government policy rather than decentralized consumer demand. Fifth, I drop 18 distribution networks whose farthest assigned building is more than 1,200 meters from the transformer. These reflect imperfect geospatial network cleaning, which is a work in progress and can be improved. Sixth, I drop 120 networks with fewer than three consumers, for which the peer instrument cannot be constructed. Finally, I drop 109 transformer networks with rollout windows shorter than one week. With a weekly time period, these are degenerate in panel regressions, so I

35. This number may seem large, but it includes transformer upgrades where additional transformers are added to reinforce the existing network.

include them here for clarity. These sample selection choices leave 177 distribution networks in the sample out of the 2,075 non-Kigali reconstructed networks (2,434 including Kigali). The networks sit in 18 utility branches out of 33 total.

Figure 6 shows the geographic distribution of the selected networks. Table 4 provides sample stats comparing the sample against non-Kigali reconstructed (but undated) networks, which are the remaining networks after the end of row 1 in the second panel of Table 3.

Overall, the sample consists of 6,992 dated connection events among 7,105 metered customers and 42,287 potential buildings across 177 networks. Date coverage among selected metered customers is 98.4 percent. The sample is primarily selected toward rural distribution networks with unusually complete installation date metadata. My sample focuses only on these high-coverage networks, which represent a little more than 8.5% of non-Kigali distribution networks and 1% of utility customers across Rwanda. Although the sample-selection threshold requires only 80 percent date coverage, the mean coverage among metered customers in the selected networks is 98 percent. Relative to other reconstructed non-Kigali networks, the selected networks are smaller, less electrified, less illuminated, and more residential. They also require a smaller share of synthetic poles to connect never-adopter buildings, so the networks indicate better coverage of buildings in the vicinity. Their average distance from Kigali is similar to the other reconstructed networks. By the end of the observed rollout, approximately 16.8 percent of potential building adopters in the sample were connected. The estimates therefore primarily describe relatively early-stage residential electrification between October 2018 and November 2021. The next section uses the reconstructed weekly grid states to measure the incentives facing each remaining unconnected building.

5 — Static Incentives and Motivating Evidence

I construct various measures of individual static incentives and dynamic strategic incentives at each weekly grid state. This section focuses on the former: the derivation of unilateral connection charges, social network peer controls, and grand coalition Shapley cost shares from the grid states. At the end of the section, I present an event study showing the heartbeat-like pattern predicted by Proposition 3. This pattern motivates the measures of strategic incentives that I construct and instrument in section 6.

5.1. Grid States and Adoption Status

Let $\mathcal{T}$ be the set of in-sample distribution transformers. Each $g \in \mathcal{T}$ has a static planned graph $\mathscr{G}_g = (V_g, E_g)$, where V_g contains transformer, pole, and building nodes and E_g contains distribution lines and service drops. Let $\mathcal{A}_g = V_g \cup E_g$ be the set of all node and edge elements in this planned network. Each element $e \in \mathcal{A}_g$ has a nonnegative supply-side cost w_e. The graph is the observed network for utility customers, augmented by the least-cost plan to connect non-adopting buildings, defined in Assumption 1. It is a directed tree network rooted at the transformer following Definition 4. The *grid state* $\mathcal{E}_{gt} \subseteq \mathcal{A}_g$ is the set of nodes and edges electrified at the beginning of week t. An edge $(u, v) \in E_g$ is active if and only if the two nodes it connects are active. The pair $(\mathscr{G}_g, \mathcal{E}_{gt})$ is the *current state of the grid* at time t. Because all paths, coalitions, and externalities are defined only within distribution tree networks, I suppress g and write the state of a distribution network as $(\mathscr{G}, \mathcal{E}_t)$.

Table 4 — Sample statistics for the selected sample versus all non-Kigali reconstructed networks

Measure	Selected sample	Non-Kigali reconstructed
Totals		
Transformer networks	177	2,075
Branches represented	18	30
Buildings	42,287	742,229
Metered consumers	7,105	252,727
Dated consumers	6,992	29,780
Mean date coverage	0.98	0.30
Per network		
Metered consumers (mean)	40.1	121.8
Metered consumers (median)	27	78
Buildings (mean)	238.9	357.7
Building connect rate	0.17	0.34
LV line km (mean)	5.5	9.3
Poles (mean)	119.6	203.8
Synthetic-pole share	0.60	0.70
Max building distance, m (mean)	781	1,046
Distance to Kigali, km (mean)	43.8	45.0
Building covariates (mean)		
Building area (m^2)	46.6	61.3
Building height (m)	2.44	2.88
Nightlights (mean)	0.039	0.648
Terrain ruggedness index	3.99	3.18
New construction share	0.160	0.127
Consumer types (share of labeled)		
Labeled consumers	6,891	249,336
Residential	0.977	0.888
Commercial	0.020	0.092
Public service	0.002	0.018
Industrial	0.0004	0.001
Apartment	0.000	0.0005

Let $\mathcal{I} \subset V$ be the set of customers and $\text{tx} \in V$ be the distribution transformer. The *network path* of customer i, $\text{path}(i, \mathscr{G})$, is the unique alternating sequence of nodes and edges from tx to i inclusive,

$$\text{path}(i, \mathscr{G}) = (v_0, e_1, v_1, \ldots, e_K, v_K), \qquad v_0 = \text{tx}, \quad v_K = i.$$

Under Assumption 2, an electrified customer's network path must be electrified: $i \in \mathcal{E}_t \implies \text{path}(i, \mathscr{G}) \subseteq \mathcal{E}_t$. The first installations observed within each network are fixed to occur at $t = 1$. The transformer is assumed to be installed the week prior $\mathcal{E}_0 = \{\text{tx}\}$. Let $y^*_{it} = 1$ if customer i is already connected to the grid at the beginning of week t. The *connected set* is $Y^*_t = \{i \in \mathcal{I} : y^*_{it} = 1\}$, and the *unconnected set* is $R_t = \mathcal{I} \setminus Y^*_t$. The panel has one row for each weekly period a customer could make a connection decision. Each network contributes observations from the week before its first adopter until its last adopter,[36] and each customer contributes observations for as long as they are unconnected at the beginning of the week.

5.2. Connection Decisions and Connecting Coalitions

At each observation, the customer $i \in R_t$ chooses whether to make a connection request $y_{it} \in \{0, 1\}$. Under Assumption 3, a connection request in week t means they are connected in week $t + 1$. Connections are permanent and move the adopting customer from R_t to Y^*_t at the end of the week: $y_{it} = 1 \implies y^*_{i,t+h} = 1$ for all $h \geq 1$. The beginning-of-week state $\mathcal{E}_t$ is predetermined by previous connection requests and is taken as given by all consumers when they make period-t choices. The *connecting set* of customers is $Y_t = \{i \in R_t : y_{it} = 1\}$. Customer i's *required network extension* at time t is the ordered subsequence of unbuilt elements on its network path, $\mathcal{B}_t(i) = \big(a \in \text{path}(i, \mathscr{G}) : a \notin \mathcal{E}_t\big)$.

Define customer i's *neighborhood* $\bar{Q}_{it} = \{j \in R_t : e^a(j, \mathcal{E}_t) = e^a(i, \mathcal{E}_t)\}$ as the maximal set of remaining customers who share unbuilt infrastructure with i. The collection of neighborhoods $\bar{\mathcal{Q}}_t$ partitions the set of unconnected customers R_t; each neighborhood is an unbuilt branch of the distribution-network tree. Each neighborhood shares an *anchor node*, which is the first unbuilt pole $e^a(i, \mathcal{E}_t)$ on their path. Via Assumption 3, customer installations within the same week whose required extensions share infrastructure are assumed to have requested a connection together, so let $Q = \bar{Q}_t \cup Y_t$ be called a *connection coalition*, and let $\mathcal{Q}_t$ be the set of all such connection coalitions at time t that partition Y_t. Any neighbor or coalition collectively providing an externality to i is therefore a subset of i's neighborhood $Q \subseteq \bar{Q}_{it}$, because customers in different neighborhoods do not share infrastructure. The required full path and network extension of coalition Q are $\text{path}(Q, \mathscr{G}) = \bigcup_{i \in Q} \text{path}(i, \mathscr{G})$ and $\mathcal{B}_t(Q) = \text{path}(Q, \mathscr{G}) \setminus \mathcal{E}_t$. The total supply-side cost of the coalition's network extension is $C(Q, \mathcal{E}_t) = \sum_{e \in \mathcal{B}_t(Q)} w_e$. The utility extends the grid just enough to connect each customer by aggregating the extensions required by all adopting coalitions in $\mathcal{Q}_t$, $\mathcal{E}_{t+1} = \mathcal{E}_t \cup \bigcup_{Q \in \mathcal{Q}_t} \mathcal{B}_t(Q)$. As grid expansion is monotone under Assumption 2, connection charges weakly decrease for all consumers as the grid expands.

36. Across all transformers, the last observed adoption occurs in November 2021, but this is not uniform across utility branches. I want to avoid assuming that weeks at the end of the recorded installations are zero-adoption weeks, so I truncate each distribution-network panel at its last observed installation.

5.3. *Anticipated and Realized Network Extensions*

A *network-extension method* $m \in \{\text{straight}, \text{lcp}, \text{path}\}$ assigns poles and lines to unconnected customers at a given state. I refer to these as unilateral *straight-line* network extensions, unilateral *least-cost-planning* (LCP) network extensions, and the *realized* or *network-path* extensions, respectively. I define the three extension methods as follows:

- The straight-line network extension draws a line from the nearest activated pole in $\mathcal{E}_t$ to the adopting customer i. The line $\mathcal{L}_{it}^{\text{straight}}(\{i\}, \mathcal{E}_t)$ is the geometric object from the existing grid to i's service-drop pole, which has raw straight-line distance d_{it}, and represents the shortest possible distance over which the utility could connect customer i.

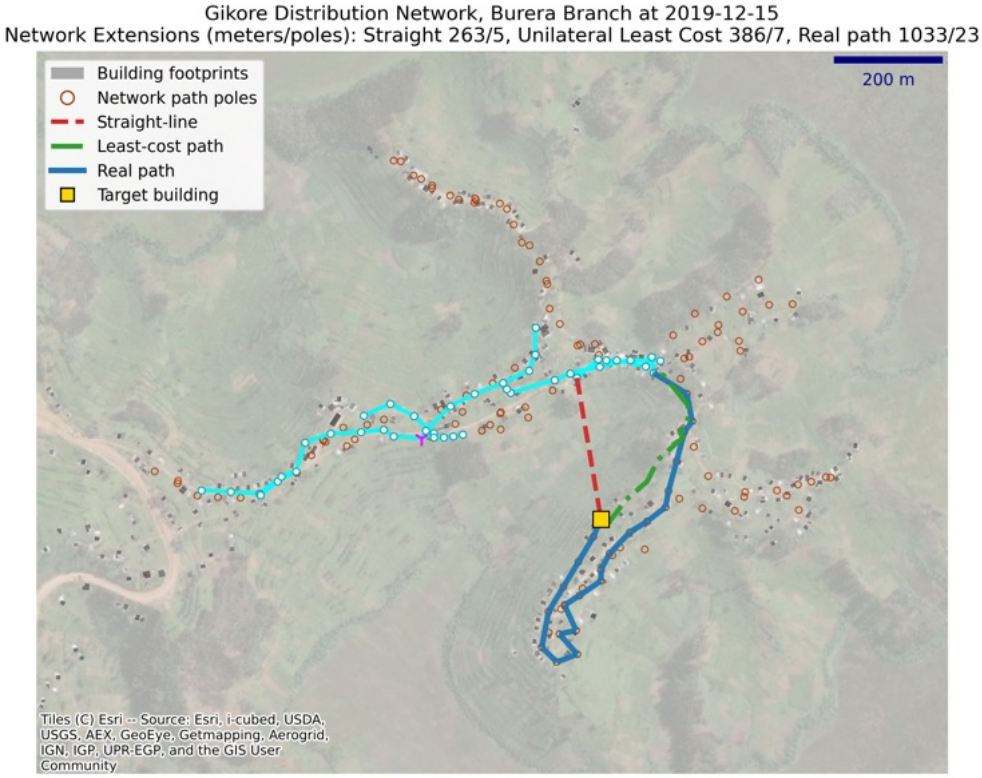
- The unilateral least-cost-planning (lcp) extension is the least-cost network path from the existing grid $\mathcal{E}_t$ to an "individual" customer singleton coalition $\{i\}$. By contrast to the straight-line network extension, it incorporates the effects of geography and terrain on the network topology. Its construction uses the least-cost-planning algorithm described in subsection 4.3. Excluding other buildings outside of the coalition from the algorithm removes their demand pull on the distribution network topology planned by the utility. I measure it as $\mathcal{L}_{it}^{\text{lcp}}(\{i\}, \mathcal{E}_t)$ and $\mathcal{P}_{it}^{\text{lcp}}(\{i\}, \mathcal{E}_t)$.

- The *network-path* extension follows the realized or planned distribution infrastructure that the utility would use to connect customer i. This method calculates the shortest path on the distribution tree and returns the lines and poles $\mathcal{L}_{it}^{\text{path}}(\{i\}, \mathcal{E}_t)$ and $\mathcal{P}_{it}^{\text{path}}(\{i\}, \mathcal{E}_t)$, together denoted $\mathcal{B}_{it}(\{i\}, \mathcal{E}_t)$, that the utility would actually build to connect i at each point in time. Unlike the previous extensions, the network path incorporates utility planning based on projected demand from other customers in the neighborhood.

Expectations over how the utility would expand the network to provide electricity can be very similar or very different from the utility's realized network path. Both the straight-line and LCP paths remove the demand pull of other customers on the realized utility network topology. Relative to the straight-line method, the unilateral LCP rule allows customers' expectations to incorporate terrain and other geographic construction constraints. Figure 7 illustrates a selected case in which a target adopter's hypothetical network extension substantially differs because of hilly terrain, which the straight-line ignores. The least-cost path goes around the obstacle but takes a shortcut relative to the real network path used by the utility, which goes around a switchback and uses substantially more infrastructure.

The observed network $\mathcal{P}_{jt}^{\text{path}}(\{j\}, \mathcal{E}_t)$ is not observed by customers ex ante, before a neighbor connects, nor are they charged for it. The straight-line and unilateral LCP network extension methods are inspired by the geographic instrumental variable methods often seen in the infrastructure and development literature. For example, Faber (2014) uses straight lines and minimum spanning trees to instrument highway infrastructure in China, while Lipscomb, Mobarak, and Barham (2013) uses a least-cost plan connecting hydroelectric dams to instrument for electricity transmission infrastructure in Brazil. However, while the above papers use network extensions to instrument for realized infrastructure *ex post*, I use network extensions to measure *ex ante* connection charges and, later, the externalities from neighbor connection requests.

Figure 7 — Alternative network-extension methods on hilly switchback terrain



Notes. Grid state at 2019/12/15 of the Gikore distribution network in Burera. The magenta shape is the distribution transformer, cyan lines with white circles are energized poles and lines, gray polygons are Google Open Buildings footprints, brown circles are the raw network-path poles in the data before geospatial cleaning, so duplicates are observable. For an adopting customer i, represented as a yellow square, the figure shows three network-extension methods.

- The red dashed straight-line extension draws the shortest line from the existing grid to the new adopter. This network extension would require 263 m of line and 5 poles for the connection.
- The green dash-dotted unilateral least-cost-planning extension incorporates the effects of terrain and obstacles to simulate the least-cost path of connecting i to the grid. This network extension would require 386 m of line and 7 poles.
- The blue solid network path follows the switchback around the hillside. This is the path the utility actually constructed, and it required 1,033 m of line and 23 poles. The utility built this path to extend the network to en-route customers.

5.4. Connection Charges

Customers get charged for the material costs in their "minimum technical solution" or unilateral least-cost plan extension $\mathcal{L}_{it}^{\text{lcp}}(\{i\}, \mathcal{E}_t)$ and $\mathcal{P}_{it}^{\text{lcp}}(\{i\}, \mathcal{E}_t)$. If the utility chooses to distort from the minimum solution to enhance the long run effects of i's provision, they pay the cost of building that additional infrastructure. Customers can assess connection charges either by asking the utility for a quotation or by applying the connection-charge policy themselves.

Applying Rwanda's connection charge policy function to the poles and lines of the unilateral LCP network extension gives the *scheduled unilateral charge* each consumer faces over time, because they represent the minimum technical solution for i's network expansion. Suppose the unilateral least-cost plan to connect i at a given moment t is $\mathcal{L}_{it}^{\text{lcp}}(\{i\}, \mathcal{E}_t)$ and $\mathcal{P}_{it}^{\text{lcp}}(\{i\}, \mathcal{E}_t)$. The total material requirements are $P_{it}^{m}(Q, \mathcal{E}_t) = |\mathcal{P}_{it}^{m}(Q, \mathcal{E}_t)|$ poles and $\Lambda_{it}^{m}(Q, \mathcal{E}_t) = \sum_{\ell \in \mathcal{L}_{it}^{m}(Q, \mathcal{E}_t)} d_\ell$ meters of line. Then Rwanda's connection-charge policy, described in subsection 3.5 and visualized in Figure 4, charges RWF 2,000 per meter of line and RWF 200,000 per pole, so that the *unilateral charge* is

$$c_{it}^{\text{lcp}}(P, \Lambda) = 2\Lambda_{it}^{\text{lcp}} + 200 P_{it}^{\text{lcp}}$$

in RWF 1000.[37] This assumes full take-up of the available credit for the 56,000 RWF meter and service drop.

5.5. Grand Coalition Cost Shares

The unilateral charge c_{it}^{lcp} is an upper bound for each adopter. But at each point in time, a customer may also adopt as part of a coalition. I use *grand coalition Shapley cost shares* $\kappa_{it}^{\bar{Q}}$ to control for the maximal benefit of cooperation. It measures the cost if every remaining customer in a neighborhood $\bar{Q}_{it}$ were to adopt at the same time and share the charges of its common infrastructure according to the Shapley Rule. When applied to cooperative infrastructure-provision games in settings ranging from airports and irrigation systems to generic rooted trees like this one, the Shapley value divides the cost of each component among its users (Littlechild and Owen 1973; Littlechild and Thompson 1977; Márkus, Pintér, and Radványi 2012). Let n_e be the number of customers sharing each network component $e \in \mathcal{B}_t(\bar{Q})$, then applying the Shapley rule piecewise and Rwanda's scheduled element charges gives

$$\begin{aligned} \kappa_{it}^{\bar{Q}} \equiv \kappa_i(\bar{Q}_{it}; q^R, \mathcal{E}_t) = &\sum_{\ell \in \mathcal{L}_{it}(\bar{Q}_{it}, \mathcal{E}_t)} \frac{2d_\ell}{|\{j \in \bar{Q}_{it} : \ell \in \mathcal{L}_{jt}(\bar{Q}_{it}, \mathcal{E}_t)\}|} \\ &+ \sum_{p \in \mathcal{P}_{it}(\bar{Q}_{it}, \mathcal{E}_t)} \frac{200}{|\{j \in \bar{Q}_{it} : p \in \mathcal{P}_{jt}(\bar{Q}_{it}, \mathcal{E}_t)\}|}. \end{aligned} \tag{9}$$

I use $\kappa_{it}^{\bar{Q}}$ as a control because I will later construct and instrument a measure of the incentives for joint adoption following the model. Here, I want to control for cooperation by other means, such as collective

37. As Rwanda uses a distance-based rather than a material-based policy, c_{it}^{lcp} is the appropriate charge. If instead a utility used a material-based rather than a distance-based charge, customers would pay c_{it}^{path}. This charge would not be known before an engineer's survey revealed the path and its cost. The network-path materials provide the engineering supply-side costs for the utility.

bargaining with the utility, coordination of rights-of-way, information sharing, and other unobserved means via which customers may cooperate. Conditional on c_{it}^{lcp}, a smaller $\kappa_{it}^{\bar{Q}}$ implies a stronger cooperation incentive. The unilateral charge c_{it}^{lcp} and the grand-coalition share κ_{it}^{R} together provide upper and lower bounds, respectively, for connection charges. Together with adopter fixed effects that absorb willingness to pay, they provide static controls for the primitive adoption incentives described in section 2 except the externality $c - c'$.

5.6. Social Peer Exposure

Another identification challenge is that social peer network effects may correlate with grid infrastructure network expansions through common geographic primitives. I use a bilateral inverse-distance network interacted with connection status to control for *social peer exposure*, following Blume et al. (2015). I intend for it to control for spatially local social peer effects, including learning, imitation, status, appliance availability, and electricity sharing. Bernard and Torero (2015) provides a framework for the social peer effects of electricity, separating them into *expectation interactions*, *preference interactions*, and *constraint interactions*. Expectation interactions include peer effects that change how customers value electricity, such as social learning. Preference interactions include social status effects such as conspicuous consumption. Constraint interactions include grid externalities, such as cooperation and free-riding that I will explicitly model, and peer effects that change the demand for connections, such as sharing electricity appliances like fridges, which increase the value of non-adoption, or external economies of scale in the supply of complementary electric appliances. Except for the strategic incentives I describe in section 2, the remaining electricity peer-effect channels that Bernard and Torero (2015) catalogs affect grid electricity adoption mainly via consumption rather than through the act of connection. While some potential channels can affect connection directly, such as information about the request process, quotations, credit, eligibility, or coordinated trips to the utility branch, these channels affect coordination costs ρ in the theory and therefore determine strategic responses.

I follow the functional form of Bernard and Torero (2015)'s social peer effects. They assume that social peer effects decline with distance and are a function of the connected set of customers y_{jt}^*,

$$\left[\mathbf{D}^{(-\zeta)}\mathbf{y}_t^*\right]_i = \sum_{j \neq i, D_{ij} \leq 1000} D_{ij}^{-\zeta} y_{jt}^*. \tag{10}$$

I follow Bernard and Torero (2015) in using this functional form with $\zeta = 1$, but perform robustness checks using alternative parameterizations of decay. In the reduced form, I let peer effects occur across transformers, but in the structural model, independence across distribution networks requires confining social peer effects to each distribution network.

5.7. Event-Study Evidence

I present an event study to motivate the remaining strategic measures. I estimate how adoption hazards vary in response to the first neighbor externality that changes the connection charge each customer faces. Specifically, this is the first week in which the grid endogenously expands toward customer i. I omit

adoptions that extend the grid at $t = 1$ because they indicate infrastructure supply rather than endogenous neighbor provision.[38]

$$\tau_i = \min\{t \in \{2, \ldots, 53\} \mid d_i(\mathcal{E}_t) < d_i(\mathcal{E}_{t-1})\}.$$

In response to a neighbor's adoption request at $\tau_i - 1$, whose network extension is realized between periods $\tau_i - 1$ and τ_i, the customer i's distance d_i falls at τ_i.

I regress adoption hazards on bins for each time period $T' = \{-12, \ldots, 12\}$ around a neighbor's first externality. I include consumer fixed effects, branch-by-week fixed effects, and optionally cost measures, social peer exposures, and additional controls:

$$y_{it} = \sum_{t' \in T'} \beta_{t'} \mathbf{1}\{t - \tau_i = t'\} + \theta_c c_{it}^{\text{lcp}} + \theta_\kappa \left(c_{it}^{\text{lcp}} - \kappa_{it}^{\bar{Q}} \right) + \theta_D [\mathbf{D}^{(-1)} \mathbf{y}_t^*]_i + \mathbf{w}_{it}' \theta_w + \mu_i + \psi_{b(i)t} + \varepsilon_{it}. \quad (11)$$

In Figure 8a, I show four estimates: one with only fixed effects μ_i and $\psi_{b(i)t}$, one adding the static cost control c_{it}^{lcp}, one adding the cooperation surplus $c_{it}^{\text{lcp}} - \kappa_{it}^{\bar{Q}}$ and social peer exposure $[\mathbf{D}^{(-1)} \mathbf{y}_t^*]_i$, and finally, one that matches the full specification of Equation 11. This includes the vector $\mathbf{w}_{it}$, which contains additional controls in the full specification I add to ensure exclusion of my instrumental variables.[39] Because the event is induced by endogenous neighbor adoption, I consider these estimates a descriptive pattern rather than a causal effect.

Under a null hypothesis of no strategic incentives, customers should adopt electricity based only on static differences in their willingness or ability to pay. The event study would show a step function at event time 0, when a lower charge is realized without anticipation of the event nor joint adoption. Under this null, controlling for the connection charge should attenuate the step if behavior were unstrategic.

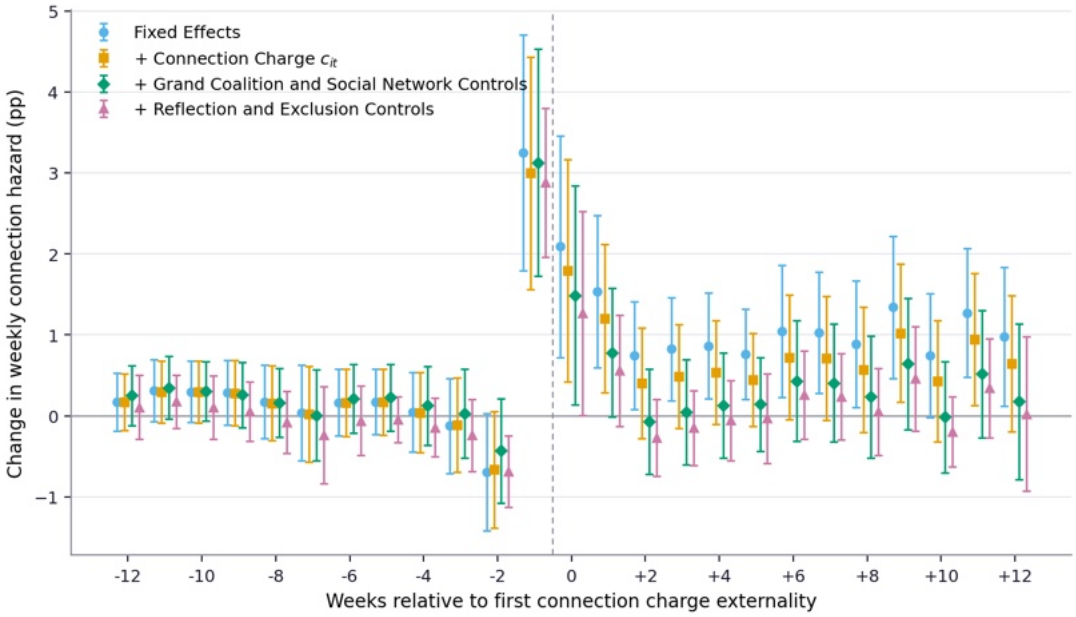
Instead, Figure 8a looks like a heartbeat, a pattern consistent with the theoretical predictions of Proposition 3. Adoption is relatively stable far from the event. It decreases as the event approaches, consistent with the proposition's delay phase. Connection hazards spike as consumers make connection requests that coincide with their neighbor, consistent with the cooperation the theory predicts. However, the spike could also reflect demand-side economies of scale or supply-side batching, which can mimic joint requests. After the externality occurs, connection hazards decline back to a higher equilibrium hazard. Rather than spuriously creating the event-time patterns, adding more controls iteratively tends to attenuate the strategic adoption effects.

Figure 8b does a heterogeneity analysis that descriptively compares events with large and small externalities. It classifies the event occurring at τ_i according to whether the reduction in the connection charge is above or below the median. The same patterns consistent with Proposition 3 occur in both groups, with larger coefficients around larger externalities. These patterns are consistent with potential utility customers adjusting adoption behavior not only in anticipation of an event on the extensive margin but also in response to externality size on the intensive margin.

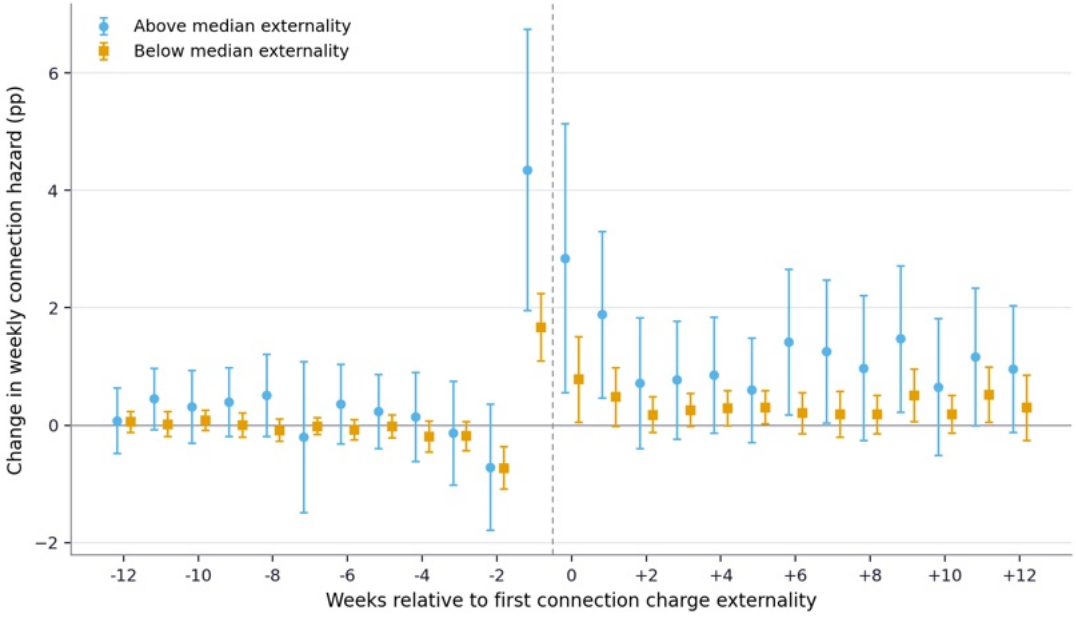
The next section constructs measures that decompose these descriptive patterns before, during, and

38. Transformers are constructed at $t = 0$. Therefore, connections made at $t = 1$ respond to the supply-side initial transformer placement rather than to endogenous externalities in connection charges.

39. Reflection controls include time-varying contextual peer effects and cooperative side-payments, and exclusion controls for exchange rate exposure and spatial autocorrelation.



(a) First cost-relevant peer externality



(b) Above- and below-median peer externalities

Figure 8 — Connection hazards around cost-relevant peer externalities.

Notes. The event week τ_i is the first week in which a neighbor provides an externality. Standard errors are two-way clustered at the transformer and week level. The reference group is the same set of adopters from periods more than 12 weeks before the event and potential adopters who never receive an externality. Observations 12 weeks after the first event are dropped. In panel A, I show a fixed-effects-only specification and add additional controls iteratively. At baseline, joint adoption at event time -1 is about six times the baseline hazard of 0.448%. The second specification adds a connection charge control, the next adds the cooperation surplus and social peer exposure; the final specification adds reflection and exclusion controls discussed in the next section. These include inverse-distance neighbor-characteristic time trends. Panel B uses the connection charge and fixed-effects specification, and separates events by whether the externality is above or below the median size. Larger externalities induce large effects, but this is also an underestimate. Differences in baseline adoption hazards mask some of the effect, as the baseline hazard is 0.592% for the above-median-externality group and 0.265% for the below-median-externality group. Binwise fixed effects control for baseline adoption hazards.

after neighbor externalities into anticipation, cooperation, and free-riding. Unlike the event study, I aim to use all externalities over the course of the grid rollout, not just the first one. It also describes how I instrument the measures to estimate causal effects.

6 — Strategic Incentives, Identification, and Instruments

The reduced-form section primarily tests empirically whether electricity adoption responds to strategic incentives. From the perspective of the structural model, the purpose of the reduced form is to check that the state variables I choose to measure strategic incentives predict adoption in the theoretically expected directions.

Identifying the strategic incentives described in section 2 is a substantial challenge for four reasons. First, utilities place infrastructure to minimize costs over the whole network, not for any one customer. Infrastructure placement is endogenous to the spatial organization of customer demand. It is also not observed ex-ante. This motivates the straight-line and least-cost-planning network extensions that provide boundedly rational measures of how customers anticipate externalities. Second, delayed, joint, and sequential adoption all result from strategic incentives, and these need to be separated. Third, grid externalities may be correlated with social peer effects, which can confound estimation via omitted variable bias. Finally, adoption decisions are made in equilibrium with neighbor actions. Therefore, adoption decisions are simultaneously determined and generate a reflection problem. Combined, the latter three empirical challenges require the empirical strategy to separate seven channels of external network effects from static incentives:

1. endogenous anticipation of future connection charge externalities,
2. endogenous cooperation via contemporaneous connection charge externalities,
3. endogenous free-riding on realized connection charge externalities,
4. endogenous social peer effects,
5. contextual grid peer effects,
6. contextual social peer effects,
7. correlated effects.

Only the first three endogenous peer effects are of theoretical interest in section 2. Disentangling *endogenous, contextual, and correlated* effects constitutes the reflection problem (Manski 1993). Applying Manski's terminology to this context, *endogenous* peer effects occur when adoption behavior depends on neighbors' adoption behavior, *contextual* peer effects occur when behavior depends on neighbors' characteristics, and *correlated* effects occur due to random local shocks. The additional challenge of my setting is that *grid* infrastructure externalities must also be separately identified from *non-grid* social or geographic peer effects, and further, that there are several endogenous responses to strategic incentives I wish to decompose.

In this section, I describe how I construct measures of each of the above *network channels* to identify them separately or control for them. First, I construct an anticipated externality matrix $\mathbf{G}_t(\mathcal{E}_t)$ which

encodes how users anticipate how each of their neighbors would induce infrastructure if they were to adopt. Each cell of the externality matrix simulates a neighbor's path extension from the current grid state with an LCP or straight-line network extension, and then measures the decrease in distance or connection charge each user receives from it.

Second, I temporally decompose the heartbeat into anticipation, cooperation, and free-riding using Proposition 3 and combine neighbor adoptions with the externality matrix to construct three measures: F_{it}, S_{it}, and L_{it}. In advance of the event at $t' < -1$, reduced adoption associated with a future externality is measured by future leads of the externality F_{it}. In the neighbor's installation week at $t' = -1$, same-week adoption is measured by contemporaneous externalities S_{it}. After the externality is realized, at $t' \geq 0$, sequential adoption is measured by lagged differences in the connection charge L_{it}. Separate demand-side peer-of-peer Z_{it}^F and supply-side shift share instrumental variables Z_{it}^S instrument F_{it} and S_{it} respectively.

Third, in addition to the externality matrix, as previously described $\mathbf{D}^{(-\zeta)}$ provides an alternative network structure. Its goal is to control for social peer effects that may be correlated with grid externalities via common geographic primitives. After introducing these variables, I proceed to present the estimating equation, and then state how the construction of each variable maps to specific network channels. I then turn to identification assumptions and explain the construction, relevance, and exclusion of the instruments.

6.1. Externality Matrices

The theory in subsection 2.1 predicts that positive externality $c_i - c_i'$ upon neighbor adoption would lower i's incentives to adopt relative to a static willingness to pay model. With more than two players, an index is needed for each neighbor j's hypothetical network extension. In the present notation, the realized externality that i would receive if j were to adopt can be written $c_i^{\text{lcp}}(\mathcal{E}_t) - c_i^{\text{lcp}}(\mathcal{E}_t \cup \text{path}_{\mathscr{G}}(j))$. As in the strategic game, the positive externality affects i's behavior via the prospect of whether j may adopt in the future via backward induction, rather than only ex-post upon its realization.

Some further restrictions on the externality become relevant in an $N = |\mathcal{I}|$-player game. By definition, customers cannot receive externalities from themselves: $i \neq j$. And connection charges are non-negative, so the maximum externality reduces the connection charge to zero. With these restrictions, the *realized externality* if j were to adopt at time t and provide an externality to i is:

$$G_{ijt}^{\text{path,c}}(\mathscr{G}, \mathcal{E}_t) = \begin{cases} \dfrac{c_i^{\text{lcp}}(\mathcal{E}_t) - c_i^{\text{lcp}}(\mathcal{E}_t \cup \text{path}_{\mathscr{G}}(j))}{c_i^{\text{lcp}}(\mathcal{E}_t)}, & c_i^{\text{lcp}}(\mathcal{E}_t) > 0 \text{ and } i \neq j \\ 0, & c_i^{\text{lcp}}(\mathcal{E}_t) = 0 \text{ or } i = j. \end{cases} \tag{12}$$

Equation 12 uses relative shares of i's current connection charge that would be cross-subsidized if j adopted unilaterally. Using shares is the simplest way to ensure that strategic incentives do not also measure static ones. It allows me to control for the static cost without collinearity, so that variables based on Equation 12 can be pure measures of strategic effects.

I stack the elements $G_{ijt}^{\text{path,c}}(\mathscr{G}, \mathcal{E}_t)$ across each column inducer j and each row receiver i within each distribution transformer in a square network path *realized externality matrix* $\mathbf{G}_t^{\text{path,c}}$. Each row of the realized externality matrix summarizes the inward externalities that customer i would receive from each of

its neighbors, and each column describes the outward externalities that adopter j would confer on each of its neighbors through unilateral adoption. The externality matrix is inspired by Elliott and Golub (2019) and can be viewed as an applied version of their adjacency matrix for networked public goods.

The realized externality matrix can be sparse. It contains zeros for four reasons. First, the diagonal of $\mathbf{G}_t^{\text{path,c}}$ is zero because a customer cannot confer an externality on itself. Second, the rows and columns corresponding to customers in the connected set Y_t^* are zero because connected customers can no longer receive or confer connection-charge externalities through subsequent connections. Third, the rows are zero for unconnected consumers under a service drop already facing the minimum connection charge, as their adoption would induce no public infrastructure. Finally, the externality matrix is block-hierarchical, inheriting this property from the adjacency matrix of a tree graph. Customers on different branches of the tree that do not share any unconstructed infrastructure, such as customers on opposite sides of the transformer, never provide realized externalities to each other.

The realized externality matrix inherits the properties of the externality itself; it is non-negative, local, directed, dynamic, and policy-dependent. It is also asymmetric because externalities are expressed as shares of recipients' current charges. A more downstream customer with a higher $c_i^{\text{lcp}}(\mathcal{E}_t)$ receives a smaller relative externality from an upstream neighbor than the upstream customer would receive from the downstream neighbor because the shared infrastructure makes the numerator approximately equal, but the downstream customer's denominator is larger. Some of Elliott and Golub (2019)'s theoretical results apply to this setting, but not all, because here the externality matrix is neither differentiable nor necessarily irreducible.[40]

An *anticipated externality matrix* uses straight-line and unilateral least-cost-planning extension rules instead of the realized network path. These omit potential endogeneity in the neighbor's extension path $\text{path}_{\mathscr{G}}(j)$ by simulating unilateral network extensions, and are defined in cost or distance measures as:

$$G_{ijt}^{\text{lcp,c}}(\mathscr{G},\mathcal{E}_t)=\begin{cases}\dfrac{c_i^{\text{lcp}}(\mathcal{E}_t)-c_i^{\text{lcp}}\left(\mathcal{E}_t\cup\mathcal{P}_{jt}^{\text{lcp}}(\{j\},\mathcal{E}_t)\right)}{c_i^{\text{lcp}}(\mathcal{E}_t)}, & c_i^{\text{lcp}}(\mathcal{E}_t)>0 \text{ and } i\neq j\\ 0, & c_i^{\text{lcp}}(\mathcal{E}_t)=0 \text{ or } i=j.\end{cases} \tag{13}$$

$$G_{ijt}^{\text{lcp,d}}(\mathscr{G},\mathcal{E}_t)=\begin{cases}\dfrac{d_i(\mathcal{E}_t)-d_i\left(\mathcal{E}_t\cup\mathcal{P}_{jt}^{\text{lcp}}(\{j\},\mathcal{E}_t)\right)}{d_i(\mathcal{E}_t)}, & c_i^{\text{lcp}}(\mathcal{E}_t)>0 \text{ and } i\neq j\\ 0, & c_i^{\text{lcp}}(\mathcal{E}_t)=0 \text{ or } i=j.\end{cases} \tag{14}$$

$$G_{ijt}^{\text{straight,c}}(\mathscr{G},\mathcal{E}_t)=\begin{cases}\dfrac{c_i^{\text{lcp}}(\mathcal{E}_t)-c_i^{\text{lcp}}\left(\mathcal{E}_t\cup\mathcal{L}_{jt}^{\text{straight}}(\{j\},\mathcal{E}_t)\right)}{c_i^{\text{lcp}}(\mathcal{E}_t)}, & c_i^{\text{lcp}}(\mathcal{E}_t)>0 \text{ and } i\neq j\\ 0, & c_i^{\text{lcp}}(\mathcal{E}_t)=0 \text{ or } i=j.\end{cases} \tag{15}$$

$$G_{ijt}^{\text{straight,d}}(\mathscr{G},\mathcal{E}_t)=\begin{cases}\dfrac{d_i(\mathcal{E}_t)-d_i\left(\mathcal{E}_t\cup\mathcal{L}_{jt}^{\text{straight}}(\{j\},\mathcal{E}_t)\right)}{d_i(\mathcal{E}_t)}, & c_i^{\text{lcp}}(\mathcal{E}_t)>0 \text{ and } i\neq j\\ 0, & c_i^{\text{lcp}}(\mathcal{E}_t)=0 \text{ or } i=j.\end{cases} \tag{16}$$

40. Indeed, it is neighborhood-wise reducible.

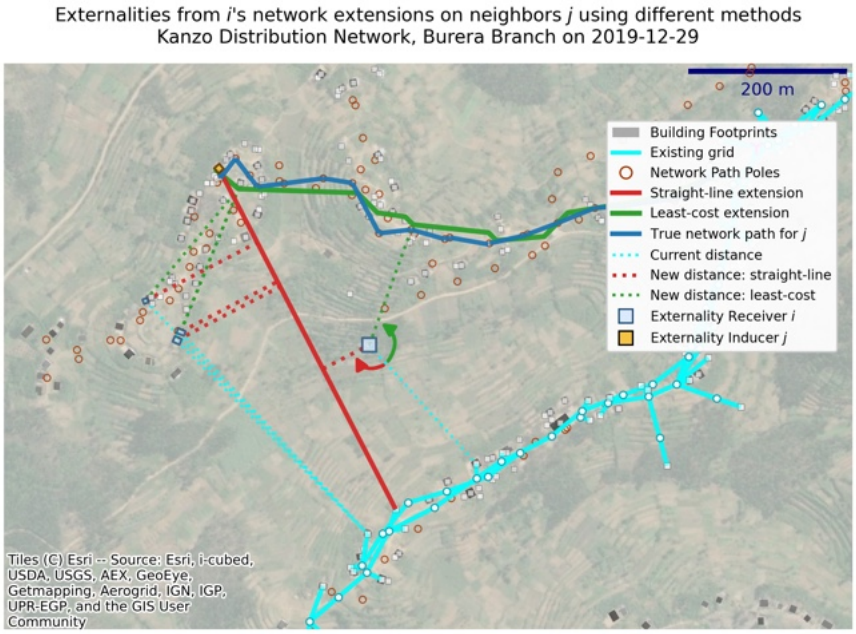
The first uses the unilateral LCP network extension to measure what poles would be constructed if j were to adopt, and the corresponding relative cost change on i as a result of their extension. The second does the same but proxies the impact on i's connection charge with the reduction in distance. The remainder are analogous connection charge and distance measures, but instead of using the LCP network extension, they base the externality on the geometric union of the poles in the grid state and the straight line from the closest pole in that graph to j. Similarly to the $\mathbf{G}_t^{\text{path,c}}$, these measures are stacked across i and j into matrices $\mathbf{G}_t^{\text{lcp,c}}$, $\mathbf{G}_t^{\text{lcp,d}}$, $\mathbf{G}_t^{\text{straight,c}}$, and $\mathbf{G}_t^{\text{straight,d}}$.

The main specifications of the paper use the anticipated externality matrix $\mathbf{G}_t^{\text{straight,d}}$ in the main specification to reflect what customers know as they make decisions. While the utility's least-cost plan $\mathscr{G}$ minimizes total infrastructure costs and therefore incorporates not only i's demand but every other potential demand in the area, customers do not observe planning documents. To correctly project neighbor externalities, they must understand the policy, anticipate who will adopt in the village and when they do so, estimate how the utility optimizes infrastructure over space to project the geographic route the utility would extend the grid in response to each neighbor's request, calculate how each extension would affect their distance to the grid and subsequent connection charge, and best respond to neighbors' mixed strategies. This is a demanding computational requirement. Network expansions are not observed at state $\mathcal{E}_t$ and are only anticipated.

Further, for potential utility customers, anticipating the geographic realization of network expansions is difficult because of asymmetric information. Each customer only adopts electricity once and observes limited utility network extensions in their neighborhood. Understanding how the utility extends the network requires potential customers also to understand to what extent the utility constructs infrastructure to incorporate unrealized demand, what information the utility has on potential customers, and what weights they place on their anticipated customers as they build out the grid, and how fluctuations in multilateral funding, supply of materials, and exchange rates may affect distribution network construction. I consider this too challenging. Instead, I consider a boundedly rational strategy in this local interaction game, following a tradition emanating from (Simon 1955; Sadler 2014). I use anticipated externality matrices $\mathbf{G}_t^{\text{straight}}$ and $\mathbf{G}_t^{\text{lcp}}$ as a boundedly rational heuristic for the realized externality matrix $\mathbf{G}_t^{\text{path}}$. The heuristic assumes that unilateral network extensions encode a reasonable expectation that if j adopts, the utility builds the minimum technical solution to reach j, which I measure using the LCP network extension or a straight line. It assumes that i does not count on the utility to necessarily extend the network towards them. Instead, i considers how much closer they would get to the grid if the utility built a network extension directly to j and considers how it would affect their connection charge. I compare these strategies visually in Figure 9.

In the primary specification of the paper, I use the straight-line distance based externality matrix $\mathbf{G}_t^{\text{straight,d}}$ to construct F_{it} and S_{it}, use the network-path externality matrix $\mathbf{G}_t^{\text{path,c}}$ to construct O_{it}, and use $\mathbf{G}_t^{\text{lcp,d}}$ and $\mathbf{G}_t^{\text{lcp,c}}$ in robustness tests to allow for richer expectations over network expansion. I opt to use distance rather than the connection charge partially for the same boundedly rational reason, but also because a charge-based externality matrix is sensitive to measurement error in the geospatial network data and simulated network extensions, as there are non-linearities in the connection charge. That said, the straight-line network extension is more naive than the unilateral LCP network extension, as it also does

Figure 9 — Externality measurement with network path, unilateral straight-line, and unilateral least-cost plan network extensions



Notes. This picture shows how adopting an externality inducer j, pictured by the yellow diamond, would extend the grid. At this state, the existing grid is colored in neon. The induced network extensions in red and green represent the boundedly rational, hypothetical, and unilateral straight-line and least-cost network extensions j would induce. Eventually, if j connects, the utility uses the blue path. The picture focuses on four customers in blue squares who receive externalities. While their initial straight-line distances from the grid are in dotted neon, their new distances are in dotted red or green. The differences in length from the neon to the red or green dotted lines are $G_{ijt}^{\text{straight,d}}(\mathscr{G}, \mathcal{E}_t)$ and $G_{ijt}^{\text{lcp,d}}(\mathscr{G}, \mathcal{E}_t)$ respectively. While the three receivers on the left have similar externalities under both methods, the central receiver has different externalities, as indicated by the red and green arrows.

not incorporate the effects of roads or terrain.

The network path externality is in general slightly larger than the straight-line externality. The mean ratio of G_{ijt}^{path} to its corresponding $G_{ijt}^{\text{straight}}$ is 1.16, while the median ratio is 1.02. I deflate estimates using $G_{ijt}^{\text{straight}}$ by 1.16, to ensure that they match the size of the real externality, while also ensuring that the estimates reflect ex-ante anticipation and contemporaneous exposure to strategic incentives.

6.2. Measuring Anticipation, Cooperation, and Free-Riding

From the externality matrices, I derive three adoption incentives for the main reduced-form specification from their rows, and another measure from their columns.

$$F_{it}(\beta, H) \equiv \max_{j} \max_{1 \leq h \leq H} G_{ijt}^{\text{straight,d}} \beta^h y_{j,t+h}, \tag{17}$$

$$S_{it} \equiv \max_{j} G_{ijt}^{\text{straight,d}} y_{jt}, \tag{18}$$

$$L_{it}(H) \equiv \frac{c_i^{\text{lcp}}(\mathcal{E}_{t-H}) - c_i^{\text{lcp}}(\mathcal{E}_t)}{c_i^{\text{lcp}}(\mathcal{E}_{t-H})}, \tag{19}$$

$$O_{it} \equiv \sum_{j} \left[G_{jit}^{\text{path,c}} \, c_j^{\text{lcp}}(\mathcal{E}_t) \right] \qquad = \sum_{j \neq i} \left[c_j^{\text{lcp}}(\mathcal{E}_t) - c_j^{\text{lcp}}(\mathcal{E}_t \cup \mathcal{B}_{\mathscr{G}}(it)) \right]. \tag{20}$$

Current inward externalities S_{it} and *future inward externalities* $F_{it}(\beta, H)$, which I refer to collectively as *in-externalities* are the primary variables of interest in the reduced form. They are each motivated by the heartbeat in Proposition 3 and the event study in Figure 8a. The in-exteranlities measure the externality that i receives from neighbors j in the current period y_{jt} or in the next H periods $y_{j,t+h}$. They interact neighbor adoption with the resulting externality that each neighbor adoption would induce, and take the best neighbor overall.

Future in-externalities F_{it} measures the strategic *anticipation* response to future neighbor externalities, corresponding the anticipable externalities $y_{j,t+h}$ from j that occur in period $t+h$. They correspond to the reaction to the upcoming event in periods $t' \in [-H, -2]$ in the event study, and to the delay and freeze phases of Proposition 3. A negative response to this would be consistent with the strategic delay of electricity adopton. The externality $G_{ijt}^{\text{straight,d}} y_{j,t+h}$ is discounted weekly by β^h and truncated at horizon H to measure the present value benefit from waiting for a neighbor to connect in the next H periods. Parameters β and H encode how much customers discount future adoptions and how far in the future they can anticipate neighbor adoption. A discount rate $\beta < 1$ would mean that i discounts externalities further in the future, while a finite horizon H would encode limited foresight. In my baseline specifications, I assume $\beta = 1$ and that $H = 8$. The former imposes no weekly discounting to benchmark anticipation and weights all H future weeks equally. The latter appears to be non-binding as the event study in Figure 8a shows no effect more than eight weeks ahead of the externality.[41] A larger H has an empirical cost because it censors observations in the last weeks of the panel, as adoptions past the last adoption in the data are not observed.

41. Also note that myopic response to neighbor play is rational when discounting is strong relative to the arrival of revision opportunities (Blume 1993).

There is one form of reflection that F_{it} does not control for. When i adopts, their extension may induce j to adopt at $y_{j,t+h}$, so realized future exposure can contain feedback from current own adoption inducing others to adopt. This creates attenuation bias against the theoretical prediction that anticipated neighbor externalities reduce adoption hazards because F_{it} would be artificially large and would increase contemporaneous adoption rather than decrease it.

Current in-externalities S_{it} are an analogous measure but only include externalities in the joint adoption phase of Proposition 3 or during the spike in adoption at $t' = -1$ in Figure 8a. They measure the *joint-adoption cooperative* response to strategic incentives in neighbor externalities. They test the reduced-form demand-side coordination economies of scale and infrastructure cost sharing incentives that enter ρ in the theory, aggregating coordination costs and benefits.

Both future and current in-externalities are constructed using a *maximum* over neighbors $j \neq i$ because, at a fixed state, neighbors provide *mutually exclusive* externalities to potential adopter i.[42] For example, consider a potential adopter in between a Y-shaped fork along a road. If neighbors on both branches of the Y extend the grid to the top of both branches, the realized connection charge for those in the middle depends on the minimum distance to the grid, and only on whichever side's infrastructure is realized to be closer to them. Therefore, if multiple connecting neighbors create alternative routes to the grid, i's remaining distance and charge are determined by the ultimate minimum connection charge to the nearer extension and the maximum cumulative externality.[43] If multiple inducers j and k adopt sequentially and both provide externalities to i, the aggregate effect is also as the closer extension at time t, induced by the more downstream of j and k.

As both measures depend on contemporaneous and future adoption, both F_{it} and S_{it} are endogenous. I instrument them using beginning-of-period-t information to only use the adoptions anticipable at that time. The former uses demand-side censored peer-of-peer effects to avoid the reflection problem, while the latter uses supply-side shift shares to estimate joint adoption. I return to the instrumental variables after presenting the estimating equation and controls.

The *past externality* $L_{it} = \frac{c^{\text{lcp}}_{i,t-H} - c^{\text{lcp}}_{it}}{c^{\text{lcp}}_{i,t-H}}$ measures the total externalities that were realized in the last H weeks. It measures the strategic *free-riding* response to a neighbor externality in the war of attrition described in Proposition 3, and the post-event decline to a new equilibrium in Figure 8a. After the connection charge c^{lcp}_{it} controls for static incentives to adopt, L_{it} measures who has delayed their adoption in order to free-ride. I construct it analogously to $G^{\text{path,c}}_{ijt}$ but using realized rather than prospective adoption, since the past is observable to prospective adopters at the beginning of week t. Under a null hypothesis of no strategic behavior, unlike the other measures, it would still produce a positive coefficient because it would capture selection effects (i.e., adopters in the new baseline war in the proof of Proposition 3); it is therefore neither identified nor a primary test of strategic incentives. I currently am using it only to decompose how often a customer would free-ride or cooperate conditional on a neighbor's adoption in a secondary dyadic regression.

The theory is that it encodes cooperation as joint adoption, but joint adoption is not the only way

42. The condition $j \neq i$ is not explicit here because $G_{iit} = 0$ for all externality matrices.

43. As multiple neighbors y_{jt} would increase F_{it} by the most downstream adopter with largest G_{ijt}, this is reflected in F_{it}.

potential grid adopters may cooperate. In particular, cooperation in the act of connecting to electricity can occur through unobserved side payments such as informal loans, contributions of money, materials, or labor, and reciprocal transfers, which would internalize the positive externalities provided by inducers. I control for potential side payments with the *outward externality* or out-externality O_{it}, which measures the total benefits that customer i's hypothetical adoption would confer on its neighbors. I construct it by interacting the numerator of the externality matrix with own adoption because side payments should be related to the monetary size of the externality that i confers on others, operating via sums instead of maxima because side payments can be received from multiple downstream neighbors. Any inferential conclusions stemming from O_{it} are indirect because it does not directly measure side payments, and the measure does not depend on endogenous adoption behavior and is predetermined at the beginning of week t, so it will not need to be instrumented.

Together S_{it}, O_{it}, and κ_{it}^R decompose three ways that customers may cooperate with neighbors to adopt electricity: joint adoption, transfers and side payments, and aggregate gains from grand-coalition adoption.[44]

6.3. Estimating Equation

I estimate the adoption hazard on the set of remaining unconnected customers. Each potential adopter contributes one observation per week until it connects, then exits the panel. Customers who adopt in the first week of a distribution network contribute only one observation and are implicitly dropped in estimation because they have singleton fixed effects. After the sample selection filters in Table 3, the panel consists of customer-week observations with complete data on the instruments and controls.

The main specification is the linear probability model,

$$\begin{aligned} y_{it} =& \theta_c c_{it}^{\text{lcp}} + \theta_\kappa (c_{it}^{\text{lcp}} - \kappa_{it}^{\bar{Q}}) + \theta_F F_{it}(\beta, H) + \theta_O O_{it} + \theta_S S_{it} + \theta_{Dy}[\mathbf{D}^{-\zeta}\mathbf{y}_{-it}^*]_i \\ &+ \theta_{Dx}[t\mathbf{D}^{-\zeta}\mathbf{x}_{-i}]_i + \theta_{Gx}[\mathbf{G}_t\mathbf{x}_{-i}]_i + \theta_{\text{FX}}\text{FX}_{it} \\ &+ \mu_i + \psi_{b(i)t} + \varepsilon_{it}. \end{aligned} \tag{21}$$

The first line includes the aforementioned variables. The second line includes controls for instrumental variables, and time-trend contextual effects. The final line adds adopter and branch-by-week fixed effects and the error term: own characteristics such as willingness to pay, static characteristics x_i, and static peer effects of all kinds are controlled for by adopter fixed effects μ_i. Branch-by-week fixed effects $\psi_{b(i)t}$ controls for correlated shocks, where $b(i)$ denotes the utility branch containing customer i. And two-way clustered standard errors by transformer and by week, because customers on the same distribution network share transformer-specific variation in the externality matrix.

Each of the terms in the first and second lines target how specific grid and nongrid endogenous and contextual peer effects affect adoption incentives via network channels:

- Static, non-strategic incentives are measured by the connection charge c_{it}^{lcp}.

44. I struggle to think of other cooperation channels that would not be captured by these measures, but I do not necessarily think that this is exhaustive.

- Static individual non-pecuniary adoption costs and benefits are measured by adopter fixed effects μ_i.
- Endogenous anticipation is measured by future inward externalities $F_{it}(\beta, H)$.
- Endogenous cooperation is measured by current inward externalities S_{it}, outward externalities O_{it}, and potential gains from a grand coalition wave of adopters $c_{it}^{\text{lcp}} - \kappa_{it}^{\bar{Q}}$.
- Endogenous free-riding is optionally measured by past externalities in L_{it}.
- Endogenous social peer effects are controlled for by inverse-distance-weighted exposure to connected customers $[\mathbf{D}^{-\zeta}\mathbf{y}^*_{-it}]_i = \sum_{j \neq i} D_{ij}^{-\zeta} y^*_{jt}$.
- Contextual grid peer effects are controlled for by externality-weighted peer characteristics $[\mathbf{G}_t\mathbf{x}_{-i}]_i = \sum_j G_{ijt}x_{jt}$. They are assumed not to vary with time except via grid expansion, and support the peer-of-peer instrument's exclusion restriction.
- Contextual social peer effects that are static are controlled for by customer fixed effects μ_i,[45] while time-varying contextual peer effects are controlled for by $[t\mathbf{D}^{-\zeta}\mathbf{x}_{-i}]_i = \sum_{j \neq i} tD_{ij}^{-\zeta}\mathbf{x}_j$.
- Static correlated effects are captured by customer fixed effects μ_i.
- Correlated shocks are partially absorbed by branch-by-week fixed effects $\psi_{b(i)t}$.
- Own procurement exposure FX_{it} supports the shift-share instrument's exclusion restriction.

The primary hypothesis test of strategic incentives is $F_{it}(\beta, H)$. An estimated $\theta_F < 0$ would indicate that anticipated endogenous positive externalities from neighbors decrease connection charges and delay current adoption. Contemporaneous inward externalities S_{it} combine cooperation and within-week free-riding incentives, so the sign of θ_S is theoretically ambiguous, but on net I hypothesize $\theta_S > 0$, noting that it may be imprecise from multiple equilibria.

Cooperation variables O_{it} and $c_{it}^{\text{lcp}} - \kappa_{it}^{\bar{Q}}$ only indirectly observe realized transfers and the surplus from grand coalition formation. Nevertheless, if neighbors internalize outward externalities via side payments or cooperate in other ways aside from joint adoption, then $\theta_O > 0$ and $\theta_\kappa > 0$ should be expected. Downward-sloping demand implies $\theta_c < 0$, and the remaining measures are pure controls.

The most involved control is for endogenous social peer effects. Using $[\mathbf{D}^{(-\zeta)}\mathbf{y}^*_{-it}]_i$ as a control is based on a two-network identification logic from Blume et al. (2015). They show that when endogenous and contextual social peer effects are transmitted via different networks, including both sources of network variation can help identify those specific network channels.[46] In my setting, I want to control for endogenous non-grid social and geographic peer effects with $[\mathbf{D}^{(-\zeta)}\mathbf{y}^*_{-it}]_i$, and contextual social and geographic peer effects with $[\mathbf{D}^{-\zeta}t\mathbf{x}_{-i}]_i = \sum_{j \neq i} D_{ij}^{-\zeta}\mathbf{x}_j t$. The inverse distance matrix proxies for social and geographic peer networks, specifically in how they would correlate with the utility's infrastructure network through common geographic primitives. As geography is a common primitive of social network formation and grid

45. Contextual social peer effects such as $[\mathbf{D}^{-\zeta}\mathbf{x}_{-i}]_i = \sum_{j \neq i} D_{ij}^{-\zeta}\mathbf{x}_j$ are static.

46. This is Theorem 7.

infrastructure, controlling for the endogenous geographic component of social networks is necessary to prevent omitted variable bias.

Where social peer effects exist that affect joint cooperation costs or benefits, such as demand-side economies of scale, they are measured in ρ, which directly affects cooperation costs and benefits, so they do not pose a threat to my identification strategy and should be captured by S_{it}. Creating a hypothesis for the sign of θ_D is outside the motivation of this paper because its job is to control for the expectation, preference, and non-grid constraint peer effects described in subsection 5.6. Nevertheless, most of the peer-effect mechanisms described by Bernard and Torero (2015) imply that electricity consumption should increase the demand for electricity connections, so one should expect $\theta_{Dy} > 0$ as a gut check.

The remaining controls in Equation 21 supports instrumental variable exclusion restrictions. Externality-weighted peer characteristics $[\mathbf{G}_t\mathbf{x}_t]_i$ control for grid contextual peer effects that operate through the infrastructure network or are correlated with it (Bramoullé, Djebbari, and Fortin 2009). Own procurement exposure FX_{it} absorbs the direct effect of an exchange-rate-based shifter on customer i's connection cost. I describe the logic and construction of these controls after explaining their corresponding instruments below.

Sans forthcoming instrumental variables and their respective controls, the primary identification assumption for Equation 21 is sequential exogeneity (Kuersteiner and Prucha 2020).[47] Sequential exogeneity allows for the non-independence across the cross-section of unconnected adopters within the panel, so that adopter behavior can be jointly determined and affected by geography, past adoption, and anticipated neighborhood demand. It requires that the grid state and the set of already-connected neighbors at the beginning of week t are uncorrelated with new idiosyncratic shocks to customer i's adoption decision during period t, after conditioning on controls and remaining in the panel. The timing of the model fits sequential exogeneity because the grid state and the set of unconnected customers are predetermined at the beginning of each period. Neither current nor future adoption retroactively changes the beginning-of-week connection charge c_{it}^{lcp}, grand-coalition Shapley charge $\kappa_{it}^{\bar{Q}}$, outward externality O_{it}, inverse-distance exposure to connected peers $[\mathbf{D}^{(-1)}\mathbf{y}_t^*]_i$, externality matrix $\mathbf{G}_t$, neighbor characteristics $\mathbf{x}_j$ or $\mathbf{x}_j t$. Nor will adoption affect any of the instruments Z_{it} and Z_{it}^S.

The current and future inward externality measures, S_{it} and $F_{it}(\beta, H)$, however, remain sequentially endogenous, because they depend on neighbors' equilibrium adoption decisions. Recall that joint adoption happens contemporaneously in S_{it}, while anticipation of future neighbor externalities happens in $F_{it}(\beta, H)$. These are mutually exclusive because future inward externalities $F_{it}(\beta, H) = \max_j \max_{1\leq h\leq H} \beta^h G_{ijt} y_{j,t+h}$ are constructed to omit contemporaneous connection measured in $S_{it} = \max_j G_{ijt} y_{jt}$. As both measures depend on the externality matrix G_{ijt}, they pose an additional identification challenge because they are correlated.

Common primitives in the externality matrix $\mathbf{G}_t$ mean that identification of either θ_F or θ_S requires a joint instrumental variable strategy. If I solely instrumented F_{it}, the remaining uninstrumented variation would be correlated with S_{it}, creating omitted variable bias in the estimate of θ_S. If I solely instrumented S_{it}, the opposite would also occur. Both must be jointly instrumented with distinct sources of variation to

47. Future versions of the paper may add their interactive fixed effects to control for correlated shocks.

avoid spilling over variation from one term to the other, and further, both instruments must be sequentially exogenous. I instrument them jointly with a peer-of-peer instrument Z_{it}^F and a shift-share instrument Z_{it}^S. The peer-of-peer instrument primarily predicts anticipated future exposure through demand-side variation. Meanwhile, the shift-share provides supply-side variation in contemporaneous exposure. The overall joint IV is just identified, with one instrument per endogenous variable: Z_{it}^F for F_{it} and Z_{it}^S for S_{it}.

The IV reduced form and first-stage regressions are:

$$\begin{aligned} q_{it} = \pi_F^q Z_{it} + \pi_S^q Z_{it}^S + \gamma_c^q c_{it}^{\text{lcp}} + \gamma_\kappa^q (c_{it}^{\text{lcp}} - \kappa_{it}^{\bar{Q}}) + \gamma_O^q O_{it} \\ + \gamma_{\text{D}y}^q [\mathbf{D}^{(-1)} \mathbf{y}_t^*]_i + \gamma_{\text{D}x}^q [\mathbf{D}^{(-1)} \mathbf{x}_j t]_i + \gamma_x^q [\mathbf{G}_t \mathbf{x}_j]_i + \gamma_{\text{FX}}^q \text{FX}_{it} + \mu_i^q + \psi_{b(i)t}^q + u_{it}^q \end{aligned} \tag{22}$$

Each regression includes one of the three endogenous variables $q_{it} \in \{y_{it},\ F_{it}(\beta, H),\ S_{it}\}$, controls, fixed effects, and indirect measures from Equation 21.

6.4. Instrumenting Anticipated Externalities

Testing whether predictable future neighbor adoption changes own adoption is a fundamental challenge because of reflection. I interact the externality matrix with itself to create a second-order externality instrument. This method adapts the social-network peer-of-peer approach developed by Bramoullé, Djebbari, and Fortin (2009) to externalities that operate via maxima. Although I apply this method to externalities, I retain the *peer-of-peer instrument* label for familiarity instead of calling it an unwieldy and unintuitive "externality-of-externality instrument". Their approach instruments a linear-in-means peer effect $[\mathbf{G}_t \mathbf{y}_t]_i = \sum_j G_{ijt} y_{jt}$ with the summed second-order peer contextual effect, or "peer-of-peer characteristics" $[\mathbf{G}_t^2 \mathbf{x}_t]_i = \sum_j G_{ijt} \sum_k G_{jkt} x_{kt}$. Through this explanation, I denote the focal customer i, the peer customer j, and the peer-of-peer customer k, noting that $\{i, j, k\} \subset \mathcal{I}$ without loss of generality. Empirically, x_k is the area of k's remotely sensed building footprint, which is an imperfect proxy for willingness to pay.

I make two modifications to Bramoullé, Djebbari, and Fortin (2009)'s peer-of-peer instrument. The first modification is to use maxima as in F_{it} and S_{it}, instead of linear-in-means. The max analog of Bramoullé, Djebbari, and Fortin (2009)'s peer-of-peer instrument is $Z^F = \max_j G_{ijt} \sum_k G_{jkt} x_{kt}$. Here, the outer externality is a maximum, matching the $\max_j$ term in F_{it} over neighbor externalities j to adopter i. Meanwhile, the inner term is a sum, because it uses contextual peer effects rather than endogenous peer effects. I have no reason to expect contextual effects occur via maxima.[48] A second modification strengthens the exclusion restriction of the standard peer-of-peer instrument. In standard peer-of-peer applications, Bramoullé, Djebbari, and Fortin (2009) argues that an identifying assumption that requires linear independence of I, G, and G^2 is sufficient for identification. Informally, however, an intuitive application of the exclusion restriction would require that peer-of-peer k's characteristics affect focal customer i's behavior only through peer j.

I show that one can censor the instrument by imposing this intuition on G^2. When k's adoption would

48. Recall the customer in the middle of a fork, where the grid is endogenously expanded on either side in subsection 6.1. If the grid were not expanded, it would be contextual because then arguably either downstream-fork-neighbors could provide a contextual peer effect.

provide a grid externality to i, this is recorded in the matrix cell $G_{ik} \neq 0$. Censoring the instrument with $\mathbf{1}\{G_{ik} = 0\} \circ \mathbf{G}^2$ removes peer-of-peers k who have a direct network linkage to i, removing customers who would directly provide an externality to i.[49] Together with the max analog, the *exclusion-censored peer-of-peer instrument* for F_{it} I use is

$$Z_{it} = \max_j G_{ijt} \sum_{k \neq i,\, G_{ik}=0} \frac{G_{jk}}{\deg^i(j)} x_k, \qquad \deg^i(j) \equiv \sum_{k \neq i,\, G_{ik}=0} \mathbf{1}\{G_{jk} > 0\}, \tag{23}$$

where the exclusion-censored degree $\deg^i(j)$ counts the peer-of-peers who satisfy the censoring condition. The censored instrument limits its variation to peer-of-peers whose adoption would not provide a grid externality to i. The inner sum $\sum_{k \neq i,\, G_{ik}=0} \frac{G_{jk}}{\deg^i(j)} x_k$ takes the weighted average of the peer-of-peer characteristics x_k who are not known to violate the exclusion restriction ex ante. When $\deg^i(j) = 0$, no peer-of-peers of j satisfy the censoring condition and so the instrument is set to zero. Censoring prevents variation from peer-of-peers who violate the exclusion restriction from entering the instrument. It does not impose an exclusion restriction, as k may affect i in other ways, but it prevents the instrument from including k's variation when I know that k directly affects i.

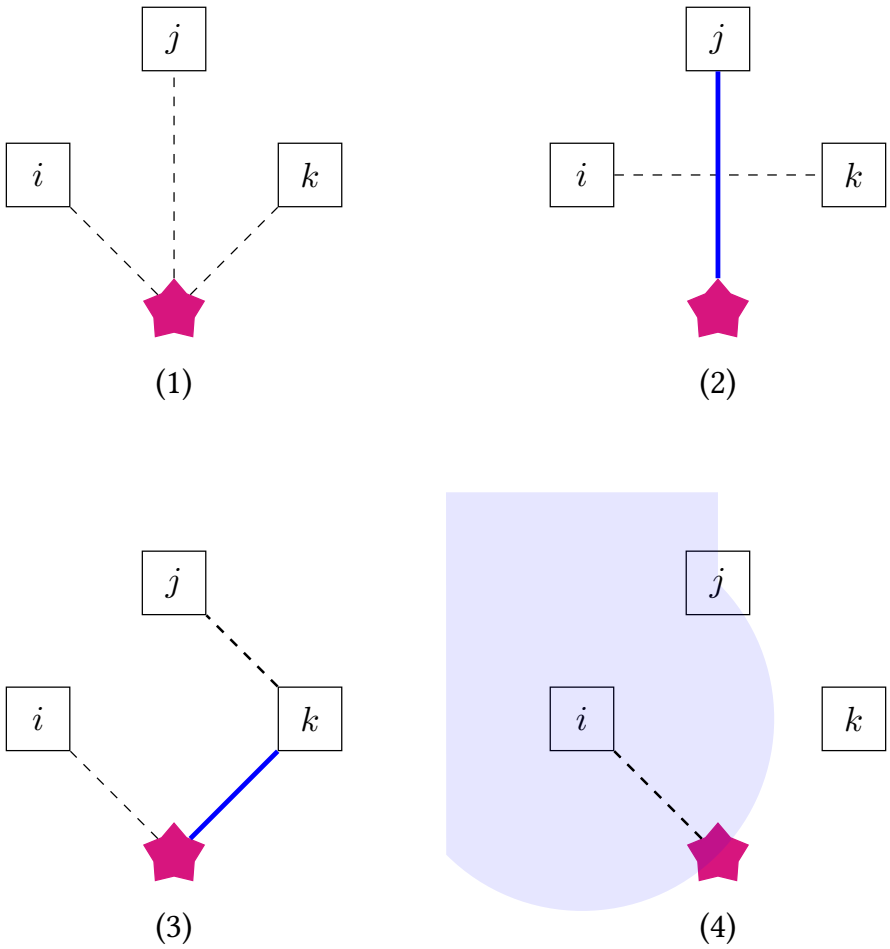
Figure 10 illustrates the instrumental variable logic and the censoring. In panel 1, the magenta transformer is placed, and customers i, j, and k estimate their connection charges based on their straight-line distances to the existing grid indicated by the dashed lines. In panel 2, j connects. This changes the minimum distance for i to connect to the grid, so j's adoption is relevant to i's externality. In panel 3, suppose that k connects first instead. This changes the minimum distance for j to connect to the grid but not i, which reflects $G_{jk} > 0$. Peer-of-peer k is relevant to j's adoption decision, but as they do not reduce i's distance, $G_{ik} = 0$, so they do not violate the exclusion restriction through G. With relevance and the direct grid-charge channel shut down, x_k can instrument for the effect of j's endogenous adoption y_j on i, as long as k does not affect i's adoption via other uncontrolled non-grid channels. Panel 4 shows a shaded region that represents building locations that are en route to or downstream of i's hypothetical network extension. Customers in these regions all violate the exclusion restriction, as their adoption would decrease i's connection charge. They customers fail $G_{ik} = 0$, and are censored in the inner sum of Z_{it}.

The exclusion-censored peer-of-peer instrument is relevant to F_{it} because peer-of-peer characteristics x_k affect the endogenous peer behavior y_j, which in turn affects y_i's decisions. For example, k may have a high willingness to pay, making j more likely to free-ride. So i may then adopt sooner because they cannot free-ride on j, because the best response for i would be to increase adoption if i and j have strategic substitutes, or decrease adoption if they have strategic complements. The instrument is externality-weighted on both legs because j may be more or less likely to free-ride on k based on the size of the potential externality G_{jk}, with the same logic for i free-riding on j.

The exclusion restriction requires x_k to affect i's adoption only through j. The instrument Z_{it} censors k when their adoption would affect i's charge path, but k's characteristics may still affect i without grid extensions, so censoring does not impose an exclusion restriction by itself. This censoring strategy is

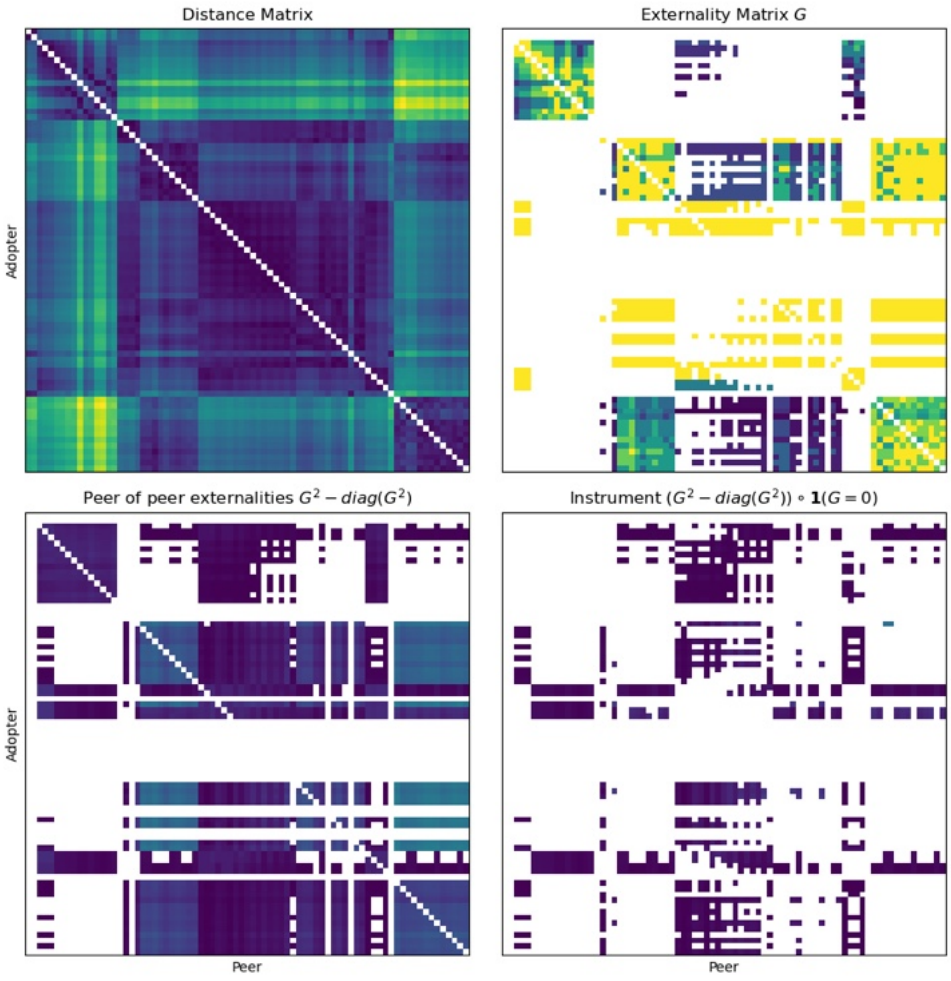
49. It does not fully impose the exclusion restriction because non-grid channels may still allow k to affect i. I include non-grid endogenous and contextual peer-effect controls for these channels.

Figure 10 — Exclusion censored peer-of-peer instrumental variable logic



Notes. This figure illustrates the logic of the peer-of-peer instrument. Panel 1: Initial distances to the existing grid. Panel 2: Endogenous adoption of j that provides an externality to i by decreasing their distance to the grid. Panel 3: Relevance and direct exclusion of k. Relevance via providing j with an externality. Neighbor k does not affect i's initial distance, so they do not provide them with an externality. Panel 4: Censored locations of neighbors who would violate the exclusion restriction. See the text for details.

Figure 11 — Censoring the externality matrix



Notes. Top-left: Distance matrix between customer i and customer j, a static and symmetric matrix. Top-right: Externality matrix between customer i and customer j, representing the change in distance to the grid for customer i if customer j were to adopt grid electricity. It is asymmetric when i and j do not use the same infrastructure. Bottom-left: Square of the externality matrix. It measures second-order peer-of-peer exposure. Removing the diagonal removes reflected variation. Bottom-right: Instrument after censoring cells with a grid path from customer k to customer i, $1\{G_{ik} = 0\}$. All matrices use the same ordering of i and j in both rows and columns. These customers are sorted in the network-path-based optimal leaf order.

limited to cases in which $G_{ik} = 0$ is known. The network must be fully observed, which may be easier in geographic network settings than in social network settings, where data are often censored (Griffith 2022).

I use two measures to control for remaining grid and non-grid channels through which k may affect i. Following the standard practice, *grid contextual peer effects* are controlled for with $\mathbf{G}_t\mathbf{x}_j$ (Bramoullé, Djebbari, and Fortin 2009). In this infrastructure network setting, $\mathbf{G}_t\mathbf{x}_j$ controls for spatial autocorrelation in the utility's remaining least-cost plan. If a component of k's characteristics, such as building area, is correlated with i's immediate neighbors, then the instrument Z_{it} could partly capture first-order neighborhood composition rather than only the effect of k on j's adoption. If buildings were systematically larger in one direction, spatial autocorrelation would make the instrument correlated with errors. The measure $\mathbf{G}_t\mathbf{x}_j$ also captures non-grid contextual peer effects that happen to be correlated with infrastructure network externalities. I also use $[\mathbf{D}^{(-1)}\mathbf{x}_j t]_i$, as described above but which also includes x_k, controlling for *nongrid contextual peer effects* that may not operate via the grid.

I visualize the distance matrix $\mathbf{D}$; the externality matrix $\mathbf{G}$; an uncensored peer-of-peer externality matrix $\mathbf{G}^2 - \text{diag}(\mathbf{G}^2)$, which removes own effects by setting $G^2_{ii} = 0$, and an exclusion-censored peer-of-peer externality matrix $\left(\mathbf{G}^2 - \text{diag}(\mathbf{G}^2)\right) \circ \mathbf{1}\{\mathbf{G} = 0\}$ with cells $\sum_j G_{ij}G_{jk}\,1\{G_{ik} = 0\}$ in Figure 11. The censored peer-of-peer externality matrix is considerably sparser where $\mathbf{G}$ indicates a violation of the exclusion restriction, showing how limiting this variation trades potential bias in the instrument from exclusion-restriction violating peers for variance.

6.5. Instrumenting Contemporaneous Externalities

While future inward externalities are instrumented with demand-side peer-of-peer characteristics, current inward externalities S_{it} are instrumented with a supply-side procurement-based shift-share so they use separate identifying variation. I map supply-side variation to S_{it} and demand-side variation to F_{it} because of an informal argument. Regardless, I will include both instruments in both first stages and will show empirical relevance via first-stage conditional F-statistics (Sanderson and Windmeijer 2016).

Current inward externalities S_{it} measure the net effect of two strategic responses generated by incentives created by an adopting coalition $\mathcal{Q}_t$. First, customer i may join the coalition, making it $\mathcal{Q}_t \cup i$. The incentive to do so is the difference between i's static connection cost less the cost share of the realized coalition $c^{\text{lcp}}_{it} - \kappa^{\mathcal{Q}_t \cup i}$, and i's bargaining costs with that group of neighbors $\rho^{\mathcal{Q}_t \cup i}$. Second, i may also free-ride on the coalition $\mathcal{Q}_t$'s network extension. The size of this incentive depends on β and how much infrastructure i shares with the group. Both incentives depend on the supply of shared infrastructure, so supply-side variation shifts the realizations of network extensions S_{it} via pass-through costs, timing, or utility choices. By contrast, future-inward externalities F_{it} primarily depend on which neighbors i can anticipate will adopt in future weeks, and the externalities they would give to i via their spatial configuration relative to i and the existing grid. Therefore, demand-side variation in the peer-of-peer instrument is more relevant to anticipating extensions in F_{it}, and supply-side variation is more relevant to S_{it}.

Specifically, the supply-side shift share instrument is s

$$Z_{it}^{S} = \underbrace{\xi_{b(i),t-78}}_{\text{Shift}} \underbrace{\max_{j} \left\{ \frac{M_{ij0}^{\text{straight}}}{d_{i0}^{\text{straight}} + d_{j0}^{\text{straight}}} \mathcal{P}_{j0}^{\text{path}} \right\}}_{\text{Shares}}. \tag{24}$$

I justify the shifts and the shares in turn.

The shifts are a mouthful: lagged foreign exchange rates weighted by within-utility-branch LV pole supplier composition. Specifically, let the shifts for branch b be a supplier-weighted procurement-cost shock

$$\xi_{bt} \equiv \sum_{m \in \mathcal{M}} w_{bm,0} \left(\frac{e_{mt}}{e_{m0}} - 1 \right) \tag{25}$$

where $w_{bm,0}$ is branch b's pre-2019 pole-manufacturer composition share $\sum_m w_{bm,0} = 1$ over pole-manufacturing countries $m \in \mathcal{M}$, and e_{mt} is the Rwandan/foreign exchange rate for m's currency in period t. The normalization $\frac{e_{mt}}{e_{m0}} - 1$ indexes exchange rates to July 2019, before the estimation period.

Branch-weighted exchange rates are used because the Rwandan utility imports a substantial share of the poles that it uses, so these are relevant to infrastructure provision. Using country-of-manufacture metadata on the utility's LV poles, I find that 30% are manufactured in Rwanda, with the rest unknown or imported. The imported poles primarily come from Uganda, South Africa, and Tunisia, with a small share from developed countries as well as India and Egypt. Supplier composition varies substantially across REG branches. For example, Gatsibo uses approximately 85% poles manufactured in South Africa, and Gicumbi uses approximately 63% poles manufactured in Uganda. Branch-specific supplier-weighted shifts create variation across branches to the extent that they have specific pole supply arrangements. There are many distribution networks within each branch b, so the shift is uniform not only across all players within each network g, but also across all distribution networks within the same branch b. Foreign exchange rates are assumed to be on an 18-month (78-week) lag. I chose this lag by comparing the pole's year of manufacture to its estimated installation date. The lag between manufacture and installation has a median of 1.4 years and a mean of two years in the data. As procurement sits between manufacture and installation, I select an 18-month lag in exchange rates to measure exposure to the shift.

Shares scale the proportion of infrastructure that i and j would be expected to share by the number of poles $\mathcal{P}_{j0}^{\text{path}}$ on customer j's planned network path when the transformer is placed at $t = 0$. Via Proposition 1, the more they share, the larger the strategic complements. In $\left\{ \frac{M_{ij0}^{\text{straight}}}{d_{i0}^{\text{straight}} + d_{j0}^{\text{straight}}} \mathcal{P}_{j0}^{\text{path}} \right\}$, the numerator $M_{ij0}^{\text{straight}} = \frac{1}{2} \left(G_{ij0}^{\text{straight,d}} \, d_{i0}^{\text{straight}} + G_{ji0}^{\text{straight,d}} \, d_{j0}^{\text{straight}} \right)$ measures the anticipated shared infrastructure savings of i and j, a boundedly rational measure of the shared infrastructure savings M_{ij} in subsection 2.1. Each term $G_{ij0}^{\text{straight,d}} \, d_{i0}^{\text{straight}}$ multiplies the externality share by i's initial distance to recover the reduction in meters of i's straight-line distance to the existing grid if j were to adopt at time $t = 0$. Then it takes the bilateral average and dividing it by the sum of their initial distances $d_{i0}^{\text{straight}} + d_{j0}^{\text{straight}}$, to give the ex-ante network expansion weighted average of how much shared infrastructure would be expected to lie on i and j's individual paths. These shares are bilateral-pair specific and relevant to how

likely each customer is to cooperate with each neighbor j via their shared infrastructure prior to any network expansion performed by either the utility or induced by anyone in the game.

The shift share is relevant because lagged procurement shocks affect utility network expansions and procurement. I control for own exposure to the same exchange-rate shift in $\mathrm{FX}_{it} = \mathcal{P}_{i0}^{\mathrm{path}} \xi_{b(i),t-78}$, which follows the same formula for own exposure to poles.[50] After controlling for FX_{it} and including branch-by-week fixed effects, the identifying excluded variation in Z_{it}^{S} comes from the interaction of bilateral spatial proximity with infrastructure procurement costs. Any branch-level rationing or changes in charges or input prices would not be i-specific and are controlled for in the fixed effects, and any individual infrastructure exposure is controlled for in FX_{it}.

6.6. Hypotheses

I make the following hypotheses following the previous discussion and Proposition 3.

Hypothesis 1 (Anticipation of neighbor adoption)**.** Anticipated inward externalities from future neighbor adoption delay current adoption $\theta_F < 0$.

Hypothesis 2 (Joint cooperative adoption)**.** Current inward externalities increase contemporaneous adoption, $\theta_S > 0$.

I test both hypotheses against the null of no strategic response, $\theta_F = 0$ and $\theta_S = 0$, with alternative hypotheses $\theta_F < 0$ and $\theta_S > 0$. Signed estimates that reject the null would indicate consistency with the strategic mechanism. Analysis of the selection effects predicted by Proposition 2 and heterogeneity across primitives of Proposition 1 that determine strategic regimes will be descriptive and informal.

7 — Evidence for Anticipation, Cooperation, and Free-Riding

This section presents reduced-form evidence on anticipation, cooperation, and free-riding, which constitute the strategic responses through which adopters respond to neighbor externalities. I first present summary statistics of all the measures described in the last three sections. Next, I present first-stage and reduced-form estimates for the joint instrumental variables in Equation 22. I then evaluate the main hypotheses of the paper and present Ordinary Least Squares (OLS) and Two-Stage Least Squares (2SLS) estimates of Equation 21. Then I consider verification of the mechanism and robustness to parametric choices, including a placebo for supply-side economies of scale, choices of network-extension rule variants, and different specifications of the social peer effect control. Finally, I perform a heterogeneity analysis across ex-ante measures of the primitives underlying Proposition 1 that, according to theory, determine the degree of strategic substitutes or complements, to assess the efficacy of the mechanism's price discrimination.

7.1. Summary Statistics

Table 5 provides summary statistics for the panel. The first two columns show the initial summary statistics of the full set of buildings upon transformer placement, reflecting the cross-section of customers at the

50. I also test $\mathrm{FX}_{it} = \mathcal{P}_{it}^{\mathrm{path}} \xi_{b(i),t-78}$, which does not change any results.

outset of the game. The middle columns show summary statistics for the overall estimation panel of 1,245,910 customer-week observations. The right set of columns shows summary statistics for customers at the beginning of the week in which they connect, reflecting buildings that select into adopting electricity.

In the first row, it shows that close to one-third of buildings connect in the first week. They are dropped due to singleton fixed effects. On average across the panel, the remaining customers have an adoption hazard of 0.379 percentage points in a given week, defined as the probability that someone adopts conditional on survival.[51] By definition, adopters have a 100% adoption hazard conditioned on adopting.

In grid state and connection incentives, we see that at the outset of the game, unilateral connection costs and connection charges are extremely high, but these fall through the panel until the connection event. The utility builds more poles than necessary beyond the minimum technical requirement, as it uses 6.07 poles to connect people to the grid when only 4 may have been necessary, spending 529,000 RWF ($\approx$ USD 529) on the median connection and 1,373,000 RWF ($\approx$ USD 1,373) on the mean connection. Just before connection, the median unilateral connection charge observed is 79,200 RWF ($\approx$ USD 79), while the mean is 1,037,000 RWF ($\approx$ USD 1,037). These are unilateral charges and do not reflect cost sharing. The grand coalition surplus is substantial when sharing lots of infrastructure but small for those who do not: connection-event-time customers have a grand coalition surplus of 21,000 RWF. Outward externalities are large, in the hundreds of thousands of USD. This reflects that providing infrastructure can benefit hundreds of downstream potential adopters.

Under strategic variables and instruments, we see that future inward externalities are large upon transformer placement and decline to an average of 12.5% during the panel. When customers connect, some could have received connection events in the future, but they adopt earlier, fewer than half of connection events have a neighbor connection that follows, so the median is zero. Current inward externalities are very high at $t = 0$ because, by construction, the first adopters on the network occur at $t = 1$. Their construction halves the connection charge and reduces the distance by 45.6% for everyone who could be connected to the network. In an average week, however, current distances decline by 2.62% and charges by 2.92%. When customers adopt, the 57.5% mean and 84.7% median inward externalities indicate many joint adoptions. Lagged externalities are zero at baseline by definition. During the panel, customers on average received a 13.8% reduction in connection charge over the previous eight weeks. During connection events, this tends to be larger, as expected, but the zero median appears to reflect the empty ballroom predicted by the war of attrition (Bliss and Nalebuff 1984). Peer-of-peer instruments are externality-weighted log building-footprint areas, so they are always positive. The supply-side shift-share instrument is an index of supplier exchange rates by branch, multiplied by poles on the path, and these are positive because Rwanda's currency declines throughout the period.

The consumer and location characteristics show that adopters have larger building footprints and taller buildings, higher pre-electrification nighttime lights, are on more rugged terrain, live at higher elevations, and are more likely to occupy buildings constructed by 2019 than the cross-section. Endogenous non-grid adoption can be positive at $t = 0$ because of customers on other grids. At $t = 0$, time-varying contextual non-grid peer effects are zero by definition. The contextual grid controls decline over time because the

51. Survival means that they have not yet adopted.

Table 5 — Summary statistics for the reduced-form estimation panel

	Cross-section at $t = 0$		All customer weeks Full panel		Adopters just before adoption $t \mid y_{it} = 1$		
Variable (Units)	Mean	SD	Mean	SD	Mean	SD	Median
Grid state and connection incentives							
Weekly connection hazard $\bar{y}_{it}$ (pp)	32.5	46.8	0.38	6.14	100	0.000	100
Ex-post path distance to grid d_{it}^{path} (m)	464	286	333	297	294	355	148
Ex-post path poles to grid p_{it}^{path}	10.2	6.20	7.08	6.50	6.07	7.51	3.00
Ex-ante LCP distance to grid d_{it}^{lcp} (m)	447	322	1,530	1,437	1,572	1,513	1,018
Ex-ante LCP poles to grid p_{it}^{lcp}	10.3	7.08	24.8	20.4	24.2	21.0	18.0
Ex-ante straight distance to grid d_{it}^{straight} (m)	351	214	233	220	196	240	77.5
Ex-ante straight poles to grid p_{it}^{straight}	5.81	4.25	3.65	4.19	3.02	4.56	0.000
Supply side connection cost (10^3 RWF)	8,694	5,395	6,042	5,636	1,373	2,674	529
Unilateral connection charge c_{it}^{lcp} (10^3 RWF)	1,461	1,581	879	1,369	1,037	2,104	79.2
Grand-coalition surplus $c_{it}^{\text{lcp}} - \kappa_{it}^{R}$ (10^3 RWF)	1,342	1,510	780	1,303	830	1,965	21.0
Outward externality O_{it} (10^5 RWF)	3,298	4,114	1,901	2,756	1,096	2,299	203
Strategic variables							
Future inward externality $F_{it}^{d,\text{straight}}(1,8)$ (%)	28.3	38.8	12.5	28.9	21.8	35.8	0.000
Future inward externality $F_{it}^{c,\text{straight}}(1,8)$ (%)	31.2	42.6	13.9	32.4	23.4	39.7	0.000
Future inward externality $F_{it}^{d,\text{lcp}}(1,8)$ (%)	26.4	36.3	13.4	28.2	20.4	33.1	0.000
Future inward externality $F_{it}^{c,\text{lcp}}(1,8)$ (%)	31.2	42.4	15.4	33.9	22.6	39.1	0.000
Current inward externality $S_{it}^{d,\text{straight}}$ (%)	45.6	38.8	2.62	13.9	57.5	44.4	84.7
Current inward externality $S_{it}^{c,\text{straight}}$ (%)	51.2	43.4	2.92	15.7	53.6	49.0	96.5
Current inward externality $S_{it}^{d,\text{lcp}}$ (%)	43.4	36.8	2.99	14.1	44.9	40.5	46.6
Current inward externality $S_{it}^{c,\text{lcp}}$ (%)	51.9	43.4	3.42	17.0	47.6	47.8	28.9
Lagged externality $L_{it}(8)$ (%)	0.000	0.000	13.8	32.5	33.2	44.1	0.000
Peer-of-peer instrument Z_{it}^{F}	0.10	0.067	0.20	0.30	0.21	0.40	0.080
Supply shift-share instrument Z_{it}^{S}	0.026	0.13	0.015	0.13	0.010	0.18	0.007
Consumer and location characteristics							
Building footprint area (m^2)	45.9	37.8	46.5	38.1	57.1	38.8	51.4
Building height (m)	2.42	1.24	2.38	1.23	2.72	1.24	2.77
Nighttime lights, 2019	0.026	0.16	0.027	0.15	0.039	0.20	0.000
Terrain ruggedness index	3.86	3.99	3.63	3.76	4.57	4.97	3.97
Elevation (m)	1,722	309	1,691	292	1,859	316	1,838
Constructed by 2019	0.91	0.28	0.91	0.28	0.93	0.25	1.00
Peer effect and IV controls							
Endogenous non-grid adoption $D^{-1}y_{it}$ (m^{-1})	0.070	0.29	0.18	0.30	0.48	1.50	0.14
Contextual non-grid $D^{-1}x_j t$: log building area	0.000	0.000	180	176	79.3	128	16.2
Contextual non-grid $D^{-1}x_j t$: log nighttime lights	0.000	0.000	0.82	4.34	1.42	8.50	0.000
Contextual non-grid $D^{-1}x_j t$: log terrain ruggedness	0.000	0.000	63.9	61.3	29.5	47.5	6.84
Contextual non-grid $D^{-1}x_j t$: log elevation	0.000	0.000	367	360	161	257	33.1
Contextual non-grid $D^{-1}x_j t$: log building height	0.000	0.000	57.0	56.7	25.3	41.5	5.30
Contextual non-grid $D^{-1}x_j t$: constructed by 2019	0.000	0.000	45.4	44.6	19.8	31.9	3.99
Contextual grid GX: log building area	1.74	0.57	1.51	0.81	1.44	1.01	1.58
Contextual grid GX: log nighttime lights	0.010	0.054	0.009	0.050	0.019	0.098	0.000
Contextual grid GX: log terrain ruggedness	0.69	0.29	0.58	0.37	0.62	0.48	0.62
Contextual grid GX: log elevation	3.57	1.12	3.09	1.63	2.91	2.01	3.22
Contextual grid GX: log building height	0.56	0.20	0.48	0.27	0.46	0.34	0.48
Contextual grid GX: constructed by 2019	0.44	0.14	0.38	0.21	0.36	0.26	0.39
Own exchange rate exposure FX_{it}	0.043	0.23	0.025	0.23	0.021	0.32	0.009
Observations	38,826		1,245,910		6,992		
Consumers	38,826		38,826		6,992		
Transformer networks	177		177		177		

Notes. Left two columns refer to initial values upon transformer placement. Middle two columns refer to the full panel. Right three columns report values conditional on adoption. The initial hazard is the share of all observed adopters who connect upon transformer placement ($2{,}273/6{,}992 = 32.5\%$). Conditional on a connection event, the hazard is mechanically 100 percentage points. Straight, LCP, and Path refer to straight-line, least-cost plan, and network path network extensions, respectively. Nonlinear connection charges are applied to these paths in $c(.)$ and distance measures in $d(.)$. Peer effects are abbreviated; contextual non-grid refers to time-varying contextual non-grid peer effects, while contextual grid peer effects are assumed to be static.

externality matrix becomes more sparse as customers adopt.

7.2. Instrumental Variable Validity

I estimate the reduced form and first stages of the instrumental-variable system following Equation 22 and report the results in Table 6. In the top half of the table, each endogenous variable, y_{it}, F_{it} and S_{it}, is regressed individually on both instruments Z_{it}^F and Z_{it}^S. The main specification is jointly instrumented and just-identified with two instruments; therefore, I report both conditional and standard two-way cluster-robust Wald F-statistics because there are two endogenous variables (Andrews, Stock, and Sun 2019). In the bottom half of the table, I regress F_{it} and S_{it} on single instruments, so that each cell is a separate regression of F_{it} and S_{it} on a single instrument with the same controls and fixed effects. These regressions help me to assess how well each instrument matches its intended endogenous variable.

Table 6 — Instrumental-variable reduced form and first stages

Endogenous Variable	Reduced-form y_{it} (pp)	First Stage S_{it}	First Stage $F_{it}(1,8)$
Units	pp	10% reduction in current distance	
Joint Instrumental Variable Regressions			
Demand-side peer-of-peer Z_{it}^F (1 s.d.)	+0.221 (0.076)	+0.019 (0.017)	-0.190 (0.054)
Supply-side shift-share Z_{it}^S (1 s.d.)	-0.130 (0.139)	+0.355 (0.091)	+0.489 (0.254)
Endogenous variable mean	0.379	0.26	1.25
Conditional F statistic	–	15.4	17.5
Joint Wald F statistic	–	7.8	9.0
Single Instrumental Variable Regressions			
Standard Z_{it}^S F statistic	–	15.2	3.6
Standard Z_{it}^F F statistic	–	1.4	12.2
Observations		1,245,910	
Distribution networks		177	
Potential Adopters		38,826	

Notes. In the top half of the panel, each column is a separate regression following Equation 22. The first column is the reduced form, the second the first stage for S_{it}, and the final column the first stage for $F_{it}(1,8)$. Standard errors are two-way clustered at the transformer and calendar-week level. The reduced-form outcome is measured in percentage points, while the first-stage estimates are scaled to an externality that decreases i's current distance to the grid by 10%. Conditional and joint Wald F-statistics are reported below the means (Sanderson and Windmeijer 2016; Andrews, Stock, and Sun 2019). Under Single IV regressions, each cell reports the F-statistic of the endogenous variable on a single instrument with all controls.

Recall that the censored peer-of-peer building-area instrument Z_{it}^F is meant to target future exposure to externalities, while the shift-share instrument Z_{it}^S is meant to target current exposure to externalities. Table 6 shows that the instruments match their intended purpose: Z_{it}^F instruments F_{it} well and Z_{it}^S instruments S_{it} well. The single-instrument specifications show that the instruments are primarily relevant for their intended endogenous variables, with diagonal (intended)F statistics of 15.2 and 12.2, and cross-F statistics

(unintended) of 1.4 and 3.6. The conditional F statistic for the Z^S shift-share instrument is 15.4, while that of the peer-of-peer instrument is 17.5 when S_{it} and F_{it} are jointly instrumented. While the joint Wald F statistics of 7.8 and 9.0 for S_{it} and F_{it} appear weak, recall that adding a second instrument with no predictive power approximately halves a singly instrumented F statistic. Because individual instruments are intended to match individual endogenous regressors, the joint Wald statistic being approximately half the conditional F statistic is by design.

The joint-IV first stage for the demand-side peer-of-peer instrument Z_{it}^F reports an estimate of -0.190 on F_{it}, 0.019 on S_{it}, and 0.221 on y_{it}. The first coefficient means that a one-standard-deviation increase in the instrument reduces anticipated externalities by approximately 15% relative to their mean. This is consistent with the mean instrumented peer j delaying adoption to free-ride on peer k with a larger building area x_k, resulting in a negative coefficient on F_{it}. As a result, i no longer expects j to connect in the near term and increases adoption in response, resulting in the +0.221 reduced-form coefficient. This is a sizeable response, as a standard deviation in Z_{it} increases the adoption hazard by 58% of the baseline rate. It has no significant loading on S_{it}, so Z_{it}^F cleanly instruments future neighbor adoption in the delay phase of the heartbeat.

The pole-based shift-share Z_{it}^S reports an estimate of 0.489 on F_{it}, 0.355 on S_{it}, and -0.130 on y_{it}. Only the estimate on S_{it} is statistically significant, indicating that the pole-based shift-share increases current inward externalities by approximately $1.4\times$ their mean. For those interpreting exchange rates as a pure supply-side cost shifter, it would appear to have the wrong sign on S_{it}. However, two countervailing forces matter. Recall that Proposition 1 shows that more shared infrastructure leads to more joint adoption under strategic complements. As S_{it} measures joint cooperative adoption, more expensive connection charges should lead to strategic complements, resulting in the positive sign indicated. An alternative interpretation is that the utility procures poles from foreign suppliers approximately eighteen months in advance using multilateral funding and uses current exchange rates at the time of procurement to pass pole costs on to consumers. Fixing pole costs to procurement-time supply costs means that Rwandan consumers' purchasing power for connection charges counterintuitively increases relative to other goods as the Rwandan Franc (RWF) depreciates smoothly over time.[52] The regression does not distinguish between these possibilities.

7.3. Full Specification Estimates

The main specification in Equation 21 is reported in Table 7. The first two columns report OLS and 2SLS versions that omit S_{it} from the main specification to show that delay estimates in F_{it} are not spuriously created via correlation with S_{it}. The third and fourth columns show OLS and 2SLS estimates of the main preferred specification. The F_{it} and S_{it} coefficients and standard errors are divided by 1.16 to convert to realized-externality units. The table also reports coefficients for the static cost control c_{it}, grand cooperation surplus $c_{it} - \kappa_{it}^{\bar{Q}}$, outward externality O_{it}, and endogenous social peer effect control $[\mathbf{D}^{(-1)}\mathbf{y}^*_{-it}]_{it}$.

Signs are consistent across specifications and match the empirical design. The first two columns demonstrate that the IV strategy does not drive the results. I focus on the preferred estimates in columns 3 and 4. Start with future inward externalities $F_{it}(1, 8)$. The OLS and IV estimates are both negative and

52. At the beginning of the panel in 2019, 1000RWF was worth 1.10USD. At the end of the panel in 2023, 1000 RWF was worth 0.77 USD. The exchange rate decline appears relatively smooth. Conveniently, around the midpoint of the study, 1000 RWF was worth 1 USD.

Table 7 — Main estimates

	$F(1,8)$ only		Full specification	
	OLS	2SLS	OLS	2SLS
	(1)	(2)	(3)	(4)
Future externality $F_{it}(1,8)$	-0.088	-1.007	-0.076	-0.910
(10% reduction in distance)	(0.023)	(0.476)	(0.021)	(0.372)
Current externality S_{it}			+0.655	+0.937
(10% reduction in distance)			(0.104)	(0.838)
Current unilateral cost c_{it}	-0.672	-0.284	-0.718	-0.383
(10^5 RWF $\approx$ 100 USD)	(0.251)	(0.262)	(0.268)	(0.227)
Grand coalition surplus $c_{it} - \kappa_{it}^R$	+0.642	+0.329	+0.684	+0.417
(10^5 RWF $\approx$ 100 USD)	(0.257)	(0.237)	(0.274)	(0.214)
Outward externality K_{it}	+0.017	+0.244	-0.059	+0.116
(10^5 RWF $\approx$ 100 USD)	(0.065)	(0.136)	(0.069)	(0.137)
Endogenous social peer effects $[\mathbf{D}^{-1}\mathbf{y}^*_{-it}]$	+0.861	+0.561	+0.951	+0.716
(weighted peer counts)	(0.139)	(0.285)	(0.078)	(0.146)
Observations	1,245,910			
Potential adopters	38,826			
Transformer clusters	177			
Mean weekly hazard (pp)	0.379			
Fixed Effects	Adopter and Branch-week			
Additional controls	Grid and Social Contextual Peer Effects; IV exclusion controls			
Overall R^2	0.199	–	0.219	–
RMSE	5.497	–	5.427	–
F statistic for F_{it}	–	12.2	–	17.5
F statistic for S_{it}	–	–	–	15.4

Notes. Each column presents a separate regression. The first two columns omit S_{it} to ensure that results are not driven by the joint IV strategy nor the correlation of S_{it} and F_{it}. In Column 2, the 2SLS estimate uses only the single peer-of-peer instrument. The next two columns report the estimates from the main specification in Equation 21. Column 4 is the preferred specification, where S_{it} and F_{it} are instrumented with first-stages on Z^S and Z^F from Equation 22. The future- and current-externality coefficients report the change in the weekly adoption hazard caused by a future or current externality that decreases the customer's current connection cost by 10% in realized terms. Grid contextual peer effects and social contextual peer effects interact $\mathbf{G}_t^{\text{straight,d}}$ and $\mathbf{D}^{-\zeta}t$ with logged variables x_j: building footprint area, 2019 nighttime lights, terrain ruggedness index, elevation, building height; and with an indicator for whether the building was constructed by 2019. The IV exclusion controls include grid contextual peer effects but also own exposure to exchange rate shocks. Standard errors are two-way clustered at the transformer and calendar-week level.

significant, as both reject zero and therefore reject the null of Hypothesis 1. The IV magnitude is about $12\times$ the OLS magnitude, which is consistent with substantial simultaneity and reflection. In the OLS, if customers are waiting for each other in the chicken game, then the reflected effect is absorbed by adopter fixed effects. The 2SLS removes most of the reflection by design, so it estimates the impact of a neighbor providing an externality without the impact of the customer on the neighbor. One might notionally consider the OLS estimate the marginal equilibrium effect and the 2SLS estimate the initial best response. The 2SLS estimates that a 10% anticipated decline in distance over the next eight weeks reduces the adoption hazard by 0.910 percentage points, $2.4\times$ the mean weekly hazard. The adoption hazard conditional on adoption on the transformer is 3.909, so 0.910/3.909 is 23.3% of the intensive margin for a 10% charge reduction. The estimates indicate that potential grid adopters substantially delay adoption when they can anticipate neighbor adoption that would decrease their connection charge.

Current inward externalities S_{it} create a 0.655-percentage-point increase in the adoption hazard in the OLS and a 0.937-percentage-point increase in the IV. These correspond to 1.7x or 2.5 times the baseline hazard, and 16.8% and 24.0% of the conditional active-transformer-week hazard, respectively. The OLS is significant, but the IV is imprecise and does not reject zero. The instrument for S_{it} is a shift share, where the share is the proportion of shared infrastructure customers shared with neighbors, and the shift is foreign exchange rates weighted by supplier composition before the study. As S_{it} measures cooperation, it does not remove reflection, as joint adoption requires both parties to participate in the adopting coalition. Notionally, one might consider the difference to represent whether i is marginal to the coalition. The imprecise 2SLS estimates may reflect that the endogenous variable measures cooperation and, therefore, strategic complements. Jumps between multiple equilibria are inherently noisy and hard to predict or instrument. I will return to this variable in the heterogeneity analysis.

Own unilateral-cost coefficients are approximately halved from the OLS to the IV specifications. A RWF 100,000 ($\approx$ USD 100) reduction in unilateral cost increases the weekly hazard by 0.718 percentage points ($1.9\times$ the mean hazard) in the OLS and 0.383 percentage points ($1.0\times$ the mean hazard) in the preferred 2SLS specification. The decline suggests customers are more elastic when joint adoption is allowed and fully accounted for in the 2SLS. The grand-coalition-surplus coefficients have signs opposite to the unilateral-cost coefficients, with similar magnitudes that tend to be smaller in the OLS and slightly larger in the IV. This is consistent with the instrument measuring joint adoption better: as the instrument identifies more of the instantaneous cooperation incentive, the overall cooperation control declines in magnitude. Finally, the outward externality O_{it} measures how much unilateral adoption in a particular state would benefit neighbors. If consumers believed that the positive externalities they induced would be compensated through side payments or internalized via some other measure, I would expect outward externalities to increase adoption. Both 2SLS estimates of the outward externality are positive but insignificant. The social-peer-effect control is positive and precisely estimated in both the OLS and IV specifications. In the preferred IV column, its coefficient is 0.716 for each unit of weighted peer exposure, consistent with electricity consumption raising adoption via the social network. For scale, it implies one connected peer 100 meters away would raise adoption hazards by 0.007 percentage points.

Overall, I take these results as strong evidence for strategic delay in advance of neighbor adoption in F_{it}, with weaker evidence for cooperation. The effect sizes are large: an anticipated 10% reduction in

distance can increase or decrease adoption by about $2.5\times$ the baseline adoption rate. I conclude from the above that, conditional on the forthcoming robustness and verification checks, strategic incentives strongly affect the *timing* of adoption.

7.4. Ruling Out Alternative Explanations

I now turn to checking robustness to various assumptions in the pipeline and verifying that alternative explanations could not explain the anticipation, cooperation, or free-riding results. First, I use alternative network-extension rules to consider how least-cost pathing that incorporates geography, terrain, and nonlinearities in the connection charge could change the results. Second, I consider alternative specifications of the social network controls by varying how social peer effects decline with distance. Finally, I use a placebo to address the threat that supply-side economies of scale confound estimation.

7.4.1 Alternative Network Expansion Measures

The baseline externality matrix is constructed using the straight-line method for simulating ex-ante anticipated network extensions. Figure 7 shows that a straight-line network extension can differ substantially from the actual network extension, and Figure 9 shows that the resulting externalities can also differ in magnitude. While the preferred straight line does not take into account geography, terrain, nor the endogenous placement of other customers, it does not necessarily simulate the realized network extensions particularly well. In this subsection, I assess robustness to their construction by varying how I construct the endogenous variables F_{it} and S_{it}. Let $\chi \in \{d(.), c(.)\}$ denote distance or cost and $\nu \in \{\text{straight}, \text{lcp}\}$ denote the straight-line or LCP network extension. Then define analogous endogenous variables to the baseline F_{it} and S_{it} constructions, which use $d(.), \text{straight}$, as:

$$F_{it}^{\chi,\nu}(\beta, H) \equiv \max_j \max_{1\leq h\leq H} G_{ijt}^{\chi,\nu}\beta^h y_{j,t+h}, \qquad S_{it}^{\chi,\nu} \equiv \max_j G_{ijt}^{\chi,\nu} y_{jt}. \tag{26}$$

I replace F_{it} and S_{it} in the preferred specification with these network extension variants. As above, I convert the coefficients to realized-externality units using the ratio of the externality on the realized path to the straight-line and LCP externalities. This is 1.16 for the straight-line extension and 0.81 for the least-cost-path extension. The charge-based columns are not converted and are in anticipated units. I report these results in Table 8.

Table 8 replicates the preferred specification in the first column. In the second, fourth, and fifth columns, I use $(\chi, \nu) \in \{(c(.), \text{straight}), (d(.), \text{lcp}), (c(.), \text{lcp})\}$ respectively. In the third column, I add a second censor to the baseline peer-of-peer instrument. In addition to censoring the peer effects of customers who have a direct linkage to i, $G_{ik} \neq 0$, in column 3, I additionally censor peers k of peers j that are within the service drop, a 37-meter radius of the focal adopter i. This limits the instrument's variation to peers of peers k who are located both outside the grid network and the immediate neighborhood of either i or j:

$$Z_{it} = \max_j G_{ij} \sum_{k:G_{ik}=0\cap d_{ik}\geq 37} \frac{G_{jk}}{\sum_k \mathbf{1}(G_{ik}=0\cap d_{ik}\geq 37)} X_k. \tag{27}$$

Proximity censoring addresses the concern that k could affect i not only through the grid network but also

Table 8 — Robustness to network extension method

Network extension method	Straight-line		Least-cost planning route		
Charge measure	Distance	Cost	Distance	Distance	Cost
	2SLS	2SLS	2SLS	2SLS	2SLS
	(1)	(2)	(3)	(4)	(5)
Future externality $F_{it}(1,8)$	-0.910	-0.807	-1.144	-1.553	-1.019
(10% reduction in charge measure)	(0.372)	(0.322)	(0.499)	(0.678)	(0.458)
Current externality S_{it}	+0.937	+0.928	+1.256	+1.978	+1.338
(10% reduction in charge measure)	(0.838)	(0.835)	(1.070)	(1.402)	(1.037)
Current unilateral cost c_{it}	-0.383	-0.406	-0.305	-0.362	-0.305
(10^5 RWF $\approx$ 100 USD)	(0.227)	(0.221)	(0.256)	(0.217)	(0.239)
Grand coalition surplus $c_{it} - \kappa_{it}^R$	+0.417	+0.426	+0.355	+0.411	+0.330
(10^5 RWF $\approx$ 100 USD)	(0.214)	(0.213)	(0.233)	(0.198)	(0.222)
Outward externality K_{it}	+0.116	+0.107	+0.139	+0.120	+0.176
(10^5 RWF $\approx$ 100 USD)	(0.137)	(0.120)	(0.167)	(0.165)	(0.173)
Endogenous social peer effects $[\mathbf{D}^{-1}\mathbf{y}^*_{-it}]$	+0.716	+0.797	+0.682	+0.812	+0.823
(weighted peer counts)	(0.146)	(0.141)	(0.179)	(0.136)	(0.144)
Conditional F for Z_{it}^F	17.5	22.0	9.7	9.6	9.8
Conditional F for Z_{it}^S	15.4	15.6	15.2	13.2	14.3
Proximity Censoring in Z_{it}^F	No	No	Yes	No	No
Direct Linkage Censoring in Z_{it}^F			Yes		
Observations			1,245,910		
Potential adopters			38,826		
Transformer clusters			177		
Mean weekly hazard (pp)			0.379		
Fixed effects			Adopter and Branch-week		
Additional controls	Grid and Social Contextual Peer Effects; IV exclusion controls				

Notes. Columns differ by network extension method and whether externalities use connection charges or distance. Instruments Z_{it}^F and Z_{it}^S are held fixed. The F and S coefficients correspond to a 10-percent reduction in exposure; cost, surplus, and outward-externality coefficients correspond to 10^5 RWF; the social-peer row is in weighted peer counts. Standard errors in parentheses are clustered by transformer and week.

more generally in space, without this effect already being controlled for in contextual or endogenous social peer effects $[\mathbf{D}^{-\zeta}\mathbf{x}_{-i}t]_i$ and $[\mathbf{D}^{-\zeta}\mathbf{y}^*_{-it}]_i$. Specifically, there may be a scenario in which the externality matrix captures the minimum network path overlap rather than geographic proximity: a peer-of-peer k might satisfy $G_{ik} = 0$ even when they are immediate neighbors because their closest points to the grid are different. It simply ensures that k is not spatially close to i, even if their adoption would not affect i's distance to the grid.

The results of Table 8 show that the results are robust to the construction of the endogenous variable and the instrument. If anything, effect sizes seem to increase when adding proximity censoring or using the least-cost-planning route, but decrease when the policy uses connection charges.

7.4.2 Distinguishing Infrastructure from Social Peer Effects

The empirical design relies on controlling for endogenous and contextual social peer effects that may be correlated with infrastructure provision externalities. The main specification uses the weight $\zeta = 1$ and the functional form $D^{-\zeta}y^*_t = \sum_j y^*_{jt}/d_{ij}$ from Bernard and Torero (2015) to control for social and geographic peer effects. I test robustness to this parameterization by changing the denominator's functional form over a grid of $\zeta \in [0.5, 2]$ in increments of 0.1. Lower values of ζ imply a flatter geographic decay and higher weight on distant peers. In comparison, higher values imply a steeper decay in social peer effects and more weight on proximal neighbors. I show a $2 \times 3 \times 15$ grid of estimates for future, current, and social peer effects. The top row varies ζ in both $[\mathbf{D}^{-\zeta}\mathbf{x}_{-i}t]_i$ and $[\mathbf{D}^{-\zeta}\mathbf{y}^*_{-it}]_i$. The bottom row removes $[\mathbf{D}^{-\zeta}\mathbf{x}_{-i}t]_i$ as a control and similarly varies ζ in $[\mathbf{D}^{-\zeta}\mathbf{y}^*_{-it}]_i$. This shifts F_{it} and S_{it} towards zero, and $[\mathbf{D}^{-\zeta}\mathbf{y}^*_{-it}]_i$ away from zero. The columns report the estimates for F_{it}, S_{it}, and $[\mathbf{D}^{-\zeta}\mathbf{y}^*_{-it}]_i$ respectively with 95% confidence bands.

The anticipation estimate changes little across the distance-decay grid. The current-exposure estimate is minimized near $\zeta = 1$, while the estimated social-peer effect peaks near that value. This indicates not only that social peer effects do not drive the results, but also that the main specification attributes the least amount of explanatory power to the externality mechanism I hypothesize. Contextual peer trends $[\mathbf{D}^{-\zeta}\mathbf{x}_{-i}t]_i$ barely affect results across the top and bottom specifications.

7.4.3 Placebo: Grid-Irrelevant Economies of Scale

A substantial threat to identification is supply-side economies of scale in utility operations. Customers may adopt simultaneously or sequentially for reasons other than the strategic incentives in the model. For example, let Assumption 3 be incorrect, so that the utility does not expand networks on its reported timeline. Instead, suppose that it batches local connection requests until they reach a threshold and constructs networks only when they reach a certain scale. This would appear as joint adoption, or even as sequential adoption across consecutive weeks if the construction is large. However, it would reflect utility practices rather than my proposed mechanism. Therefore, I use a placebo to test whether the estimates reflect connection-charge externalities rather than utility-side economies of scale.

Define a *grid-relevant peer* as one who could provide an externality to the consumer at any point during the electrification rollout. Placebo *grid-irrelevant* peers are their complement: they could never provide

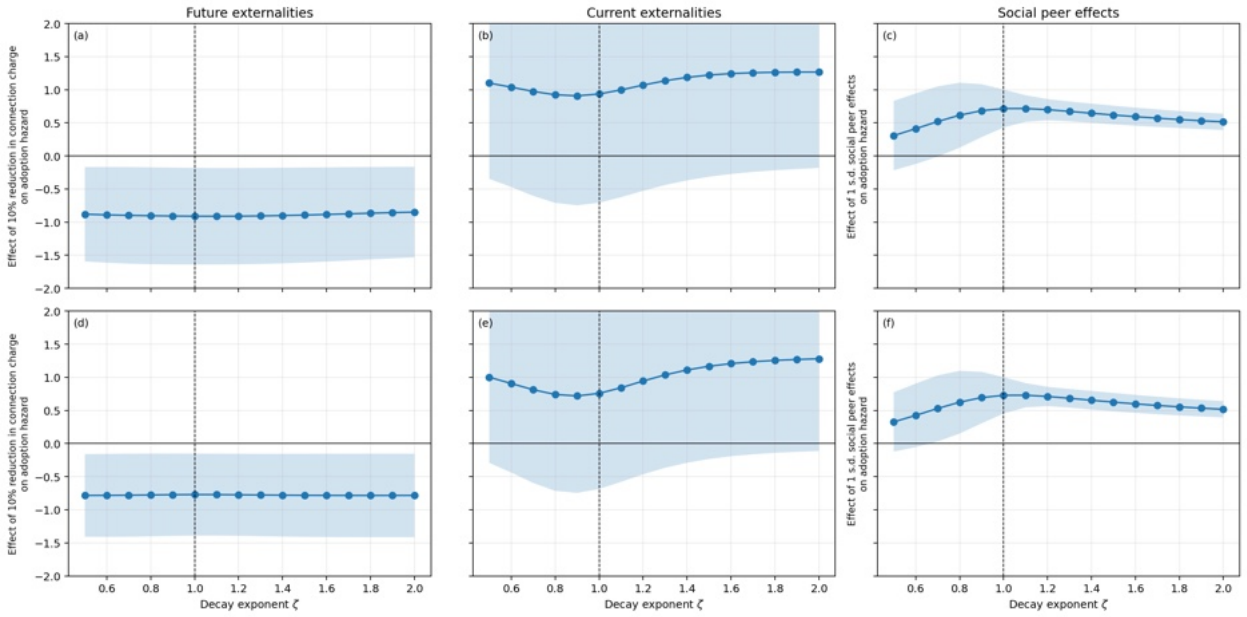


Figure 12 — Sensitivity of estimates to assumed decay in social peer effects ζ.

Notes. Panels (a)–(c) use a common ζ for $[\mathbf{D}^{-\zeta}\mathbf{y}^{*}_{-it}]_i$ and $[\mathbf{D}^{-\zeta}\mathbf{x}_{-it}]_i$; panels (d)–(f) omit the contextual peer controls $[\mathbf{D}^{-\zeta}\mathbf{x}_{-it}]_i$. All other controls and variables follow Table 8.

an externality but sit on the same transformer. Using the externality matrix, this defines the set of viable placebo peers for i to be $\text{Placebo}_i = \{j \neq i : g(j) = g(i),\ G_{ijt} = 0\ \forall t\}$. As the mechanism I propose in section 2 operates through infrastructure provision $G_{ijt} > 0$, placebo peers should not affect i through the externalities I propose. They may still affect i through utility economies of scale because they remain on the same network, are typically within one mile of each other, and their adoption would bring utility supply teams to the local village. If utility distribution, transportation, and logistics create economies of scale that drive my previous results, then the placebo should have the same sign and attenuate estimates in a horse race.

I construct the placebo on the intensive margin by defining

$$F_{it}^{\text{Placebo}} = \max_{j \in \mathcal{P}_i} \max_{t+h \in [1,8]} y_{jt} \min\left\{\frac{37}{d_{ij}}, 1\right\}. \qquad S_{it}^{\text{Placebo}} = \max_{j \in \mathcal{P}_i} y_{jt} \min\left\{\frac{37}{d_{ij}}, 1\right\}.$$

and then use an extensive margin indicator to define, for example, $\mathbf{1}[F_{it}^{P} > 0] = \mathbf{1}\left\{\exists j \in \mathcal{P}_i : 1 \leq t_j - t \leq 8\right\}$. so that the variation in F_{it} and F_{it}^{Placebo} can be easily compared, and similarly for S_{it} and S_{it}^{Placebo}. I then replace the baseline F_{it} and S_{it} with the extensive-margin indicator or rescale them to standard deviations, and enter them into a horse race with the analogous placebo measure. The results of the horse race are in Table 9.

Table 9 shows that in anticipation of a neighbor externality F_{it}, the placebo has the opposite sign and is insignificant. By contrast, having the grid extended on the other side of the network increases adoption rather than decreasing it. This occurs on both the extensive and intensive margins. For contemporaneous neighbor externalities S_{it}, the placebo matches the sign and is insignificant, but it increases the size of the

Table 9 — Placebo grid-irrelevant neighbor network extensions.

	Extensive margin: $\mathbf{1}[F_{it} > 0]$		Intensive margin: F_{it}	
	Baseline IV	IV + placebo	Baseline IV	IV + placebo
	(1)	(2)	(3)	(4)
Placebo future F_{it}^{Placebo}		+1.76 (0.96)		+0.413 (0.259)
Grid future F_{it}	-9.04 (3.73)	-9.85 (4.87)	-2.694 (1.102)	-2.495 (1.117)
Placebo current S_{it}^{Placebo}		+0.369 (0.270)		+0.396 (0.232)
Grid current S_{it}	+1.631 (1.321)	+1.725 (1.575)	+1.447 (1.295)	+1.739 (1.435)
Observations		1,245,910		
Distribution networks		177		
Conditional first-stage $F(F)$	17.15	9.80	17.54	12.80
Conditional first-stage $F(S)$	14.19	11.24	15.38	13.31

Notes. The first and third columns replace the endogenous variables in the RF headline with extensive margin and standard deviation-scaled versions of those measures. Columns two and four add analogous placebo measures to that baseline.

S_{it} estimate. It indicates that demand-side cooperation may be strengthened when utility-side economies of scale are controlled for. Together, these indicate that utility-side economies of scale do not drive the results.

7.5. Selection and Price Discrimination

Lemma 1 and Proposition 2 predict that high-connection-value or high-willingness-to-pay consumers are more likely to connect earlier, pay higher connection charges, and provide externalities to their neighbors. I assess the empirical magnitude of this proposition by estimating the connection charges of realized adopters as a function of wealth and customer characteristics. These forthcoming regressions are cross-sectional and descriptive, and they reflect game outcomes of the mechanism rather than causal estimates of strategic incentives.

To ensure these cross-sectional results do not arise from spurious bilateral matching between remotely sensed buildings and electricity meters, I also run the same analysis on the World Bank Multi-Tier Framework Survey. This survey focused on energy use in Rwanda and is publicly available (World Bank 2018a, 2022). I use theory-relevant values from the separate rounds conducted in 2016 and 2022 and regress connection charges on them.

The cross-section is limited to customers who adopt an infrastructure connection. In the REG data, I regress the realized connection charges on the log difference between the initial and realized connection charge, adoption week, and rank on log building area, the most predictive customer covariate I can match to meters. In the World Bank data, I regress paid connection charges on dwelling area, dwelling value, electricity consumption, a principal component analysis (PCA) of electric appliance ownership, and Ubudehe category, a social classifier that the Rwandan government assigns to households and social programs.[53]

53. It categorizes households into poorest, poor, self-reliant, and affluent categories, with the bottom tier allowing access to

Dwelling area and value were not available in the 2022 dataset. I treat building area and value as a pure income effect, while electric appliances, consumption, and expenditure reflect willingness to pay for electricity specifically, since connections are used to consume and provide electricity services.

I estimate Poisson pseudo-maximum-likelihood regressions of the form

$$\mathbb{E}\Big[c_i^R \mid A_i, c_{i0}^{\text{path}}, x_i, g(i)\Big] = \exp\left(\iota_{g(i)} + \beta \ln A_i + \gamma \ln\Big(c_{i0}^{\text{path}}\Big) + \eta x_i\right),$$

and matched OLS regressions of the form

$$\ln\big(c_i^R\big) = \iota_{g(i)} + \beta \ln A_i + \gamma \ln\Big(c_{i0}^{\text{path}}\Big) + \eta x_i + \varepsilon_i,$$

These specifications allow me to estimate the elasticity of the realized connection charge with respect to building area and other variables determining willingness to pay. The variables are $\iota_{g(i)}$, which denotes transformer fixed effects; c_{i0}, the initial connection charge upon transformer placement; and c_i^R, the realized connection charge. To put the connection charges I estimate on the same footing as the World Bank connection charges, for which there are few zeros in the connection fee, I add the meter connection fee to my public-good connection charges, without assuming, as before, that there is 100% take-up of the credit available for the private component of the connection fee. I use transformer or cell (adm4) fixed effects, and include two additional controls and utility data: whether the building was constructed in mid 2019, and the initial connection charge upon transformer placement. The former excludes constructed households from entering the price discrimination metric, and the latter conditions the realization of price discrimination on the initial state of the game.

First, note that under a flat fee, any regression of the connection charge on an independent variable after conditioning on adoption would mechanically be zero. The results indicate price discrimination in nearly all OLS specifications and most PPML specifications. Panel A shows that building area itself, a weak proxy for wealth, is correlated with the realized charge and cross-subsidy. The larger the building area, the higher the connection charge and the fewer cross-subsidies customers receive from neighbors. Within the sample, this suggests that the connection charge has an elasticity of 0.04 with respect to building area. At the same time, the proposed mechanisms lead customers with 1% larger building area to receive 0.06% less subsidies. However, the same customers adopt 50% earlier in time. Panel B shows that connection charges increase with dwelling value, an indicator of wealth, with an elasticity of 0.05–0.06. It also indicates that households in a higher Ubudehe category (higher social class) paid larger connection charges in 2022. However, neither Ubudehe category nor dwelling area is correlated with connection charges in the 2016 survey. Across measures of willingness to pay, the elasticities of paid connection charges range from 0.09 to 0.24. Connection charges are most elastic with respect to electricity consumption in 2022; the elasticity is 0.17 for electricity expenditure and 0.10 for the PCA of electrically powered appliances.

Because a flat fee would produce no variation in connection charges by construction, I interpret these regressions as evidence that the proposed mechanism results in modest price discrimination. The results indicate that adopters with a high connection value or willingness to pay, or who are wealthier, both adopt

certain social programs. These are ordinal numbers 1-4.

Table 10 — Connection charges, adoption order, and consumer wealth

Dependent variable	Independent variable	Dataset	PPML	OLS	N_{PPML}	N_{OLS}
Panel A: Realized adoption						
Connection charge	Building area	REG	0.037	0.041	5,813	5,813
			(0.026)	(0.020)		
Gross cross-subsidy	Building area	REG	-0.057	-0.058	5,688	5,688
			(0.021)	(0.021)		
Adoption week	Building area	REG	-0.565	-0.532	5,813	5,813
			(0.267)	(0.254)		
Adoption rank	Building area	REG	-0.019	-0.019	5,813	5,813
			(0.006)	(0.006)		
Panel B: WB MTF survey						
Connection charge	Dwelling area	MTF 2016	0.002	0.012	950	950
			(0.034)	(0.028)		
Connection charge	Dwelling value	MTF 2016	0.063	0.047	932	932
			(0.027)	(0.023)		
Connection charge	Electric appliance PCA	MTF 2016	0.108	0.092	982	982
			(0.036)	(0.029)		
Connection charge	Electric appliance PCA	MTF 2022	0.109	0.110	1,480	1,480
			(0.026)	(0.030)		
Connection charge	Electricity consumption	MTF 2016	0.107	0.102	933	933
			(0.040)	(0.034)		
Connection charge	Electricity consumption	MTF 2022	0.170	0.235	1,449	1,449
			(0.053)	(0.047)		
Connection charge	Electricity expenditure	MTF 2022	0.125	0.171	1,449	1,449
			(0.038)	(0.034)		
Connection charge	Ubudehe category (1–4)	MTF 2016	0.080	0.029	898	898
			(0.052)	(0.040)		
Connection charge	Ubudehe category (1–4)	MTF 2022	0.115	0.101	1,471	1,471
			(0.052)	(0.046)		

Notes. Panel A uses the cross-section of realized adoption. Table 5 provides descriptive statistics for these. These use transformer fixed effects and clustered standard errors, with controls for baseline building existence and initial connection charge. Panel B uses the World Bank Multi-Tier Survey 2016 and 2022 datasets World Bank (2018a, 2022). Fixed effects are at the cell or survey group level, with corresponding clustered standard errors. The PCA is standardized to z-scores; building areas, value, electricity consumption, and spending are logged. PPML and OLS regressions are both reported and rescaled where necessary to report elasticities, except for Ubudehe and the PCA. Standard errors are clustered by transformer in the top panel and by cell (Admin 4-level regions) in the bottom panel. Estimates reflect elasticities conditional on controls. The sample is restricted to realized electricity adopters. In the World Bank 2016 data, connection charges are in panel-column C7, area is B13, value is B9, monthly kWh is C22, and Ubudehe is B29. In the World Bank 2022 survey, connection charges are in panel-column C007, monthly electricity spending is C025, and Ubudehe is B13. PCAs are constructed from lighting, media, heating and cooling, computer, and TV-related appliances.

earlier, pay higher connection charges, and receive fewer cross-subsidies from their neighbors' provision of shared infrastructure.

7.6. Decomposition of Free-Riding and Cooperation

To close out the reduced-form part of the paper, I finally turn to a heterogeneity analysis to decompose free-riding and cooperation. Proposition 1 proposed that the larger the strategic complements, the more customers should join the same adopting coalition. The stronger the strategic substitutes, the more free-riding should result. When applied to the near-mechanical baseline of Rwanda's connection charge policies, one should imagine two close neighbors toward the leaves of the distribution tree network as having strategic complements, while those sharing a trunk but on different brambles have strategic substitutes.

I compare joint adoption $J_{ij} = \mathbf{1}(y_{it} = y_{jt})$, which is defined as i connecting in the same week as j, with sequential adoption, in which i connects between one and eight weeks after j, $L_{ijt} > 0$ or $L_{jit} > 0$, using a logit to estimate the probability that customers jointly connect in a dyadic regression. Each pair is one observation. I take all dyadic pairs that fall into one of the two categories and define the binary variable $\Pr(J_{ij}) = J_{ij}/(J_{ij} + L^8_{ijt} + L^8_{jit})$. which is the probability that customers jointly adopt given that they either jointly or sequentially adopt. This conditions the sample to 56,665 dyadic pairs who adopt within 8 weeks of each other. I divide the result of Proposition 1 by c_i, and estimate this probability on bins of the proportion of overlapping infrastructure $M_{ij}/c_i = G_{ij}/c_i$, which is the share of i's standalone connection charge that can be saved by sharing infrastructure with j, other terms in the regression, and controls.[54] Specifically, the baseline dyadic regression is

$$\begin{aligned}\log\left(\frac{\Pr(J_{ijt})}{1-\Pr(J_{ij})}\right) = \beta_0 + \sum_{b=2}^{6}\beta_b \mathbf{1}\left\{\frac{M_{ij}}{c_i} \in B_b\right\} + \rho\frac{c_j}{c_i} + \gamma f(\frac{1}{c_i}) \\ + \theta_\kappa\left(c^{\text{lcp}}_{it} - \kappa^{\bar{Q}}_{it}\right) + \theta_O O_{it} + \theta_{Dy}[\mathbf{D}^{-1}\mathbf{y}^*_{-it}]_i + \theta_{\text{FX}}\text{FX}_{it} \\ + \boldsymbol{\theta}'_{Gx}[\mathbf{G}^{\text{lcp}}_t \mathbf{x}_{-i}]_i + \boldsymbol{\theta}'_{Dx}[\mathbf{D}^{-1}\mathbf{x}_{-i}t]_i.\end{aligned} \tag{28}$$

I exclude fixed effects to avoid incidental parameters. I bin the shared infrastructure savings as $B \in \{[0, .125), [.125, .25), [.25, .50], (.50, .75], (.75, .875], (.875, 1]\}$, where the middle bins are larger to reflect sample sizes.

I report four pair specifications. The first pair uses connection charges $c(.)$ while the second pair uses straight line distances $d(.)$ Then I consider linear and spline forms for $f(\frac{1}{c_i})$. In a linear specification, $f(\frac{1}{c_i}) = \frac{1000}{c_i}$, while in a spline, I fit a cubic spline to $\frac{1}{c_i}$ for more flexibility. For each interval, I report fitted logit probabilities over characteristics so that $\widehat{\Pr}_b = \frac{1}{N}\sum_{ijt}\Lambda\left(\widehat{\beta}_0 + \widehat{\beta}_b + \widehat{\mathbf{X}}'_{ijt}\widehat{\boldsymbol{\delta}}\right)$ for each bin b, holding the distribution of controls across bins constant.

Table 11 and Figure 13 show that joint adoption becomes much more common as customers share more infrastructure. The figure reports the linear estimates of the table. In the linear specification, the predicted share of active dyadic responses that are joint rises from 32.2% in the lowest shared-infrastructure bin to 66.3% in the highest bin using connection charges, and from 32.2% to 81.0% using distance. The spline

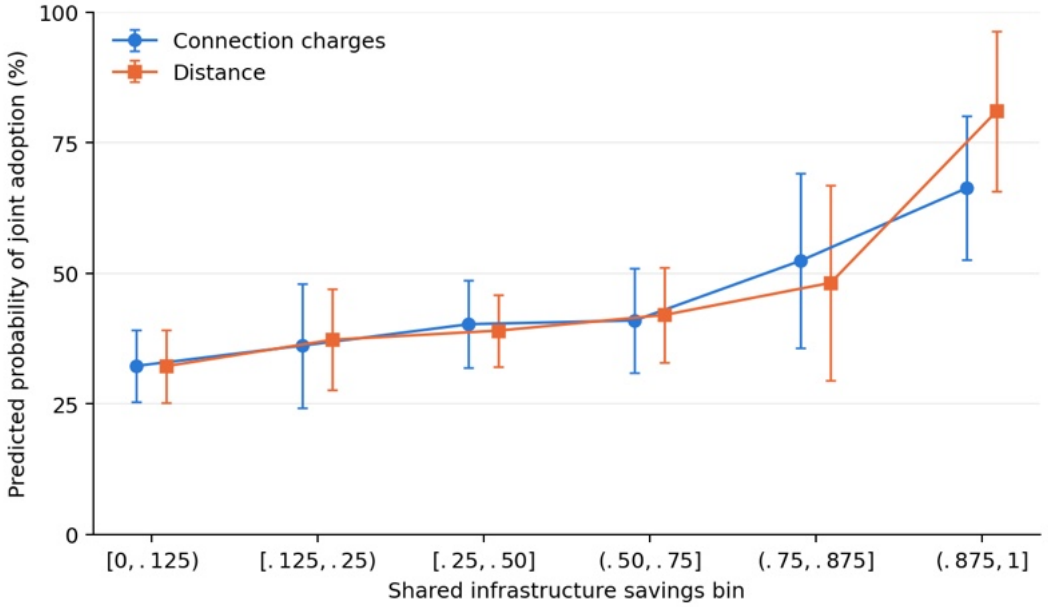
54. I am not sure if I should do the same for j, given that the pair would be counted twice, or put controls for j in the same regression.

Table 11 — Predicted probability of joint adoption

	$[0, .125)$	$[.125, .25)$	$[.25, .50]$	$(.50, .75]$	$(.75, .875]$	$(.875, 1]$
Cost measure						
Linear	32.2	36.1	40.2	40.9	52.4	66.3
	(3.5)	(6.1)	(4.3)	(5.1)	(8.6)	(7.0)
Splined	32.2	37.4	40.4	42.0	54.6	64.6
	(3.5)	(6.5)	(4.6)	(4.7)	(8.3)	(6.1)
Bilateral Pairs	43,009	2,303	3,106	3,123	1,583	3,541
Distance measure						
Linear	32.2	37.3	39.0	42.0	48.2	81.0
	(3.6)	(4.9)	(3.5)	(4.6)	(9.6)	(7.8)
Splined	32.1	38.3	39.8	42.6	48.0	78.2
	(3.6)	(4.9)	(3.5)	(4.8)	(9.3)	(8.4)
Bilateral Pairs	39,993	3,721	4,895	3,501	2,267	1,903

Notes. Each bin cell reports the predicted probability for dyadic pairs of customers who jointly adopt rather than sequentially do so $P(J \mid J + L^8 = 1)$. Each cell reports predicted probabilities after averaging the logit probability over the entire sample (removing selection effects). I use two-way clustered standard errors by transformer and week and use the delta method to report standard errors in percentage points to match the logit probability.

Figure 13 — Predicted probability of joint adoption by shared-infrastructure bin



Notes. Predicted probabilities from the linear specification of Table 11. Delta-method standard errors are reported as 95% confidence intervals in whiskers. The raw SEs were two-way clustered by transformer and week. Cubic-spline estimates and per-bin pair counts are reported in Table 11.

estimates are similar, rising from 32.2% to 64.6% for connection charges and from 32.1% to 78.2% for distance. The gradient is relatively gradual across the middle bins and becomes especially steep as customers share most of their infrastructure. Because these probabilities condition on either joint or sequential adoption, the pattern shows that greater infrastructure sharing shifts strategic responses from free-riding toward joint adoption, consistent with Proposition 1.

The overall reduced-form evidence shows that distance-based charges create local strategic responses consistent with the theory. Potential adopters delay when they anticipate a neighbor-induced reduction in their future charge. Conditional on a neighbor's intended expansion, they either join and cooperate to adopt jointly or free-ride on them sequentially. Higher willingness to pay and wealthier adopters adopt first, provisioning infrastructure that correlates with realized connection charges and willingness to pay. Strategic complements and substitutes result in heterogeneity in the degree of cooperation or free-riding.

8 — A Dynamic Coalition Model of Expanding Infrastructure

The preceding reduced-form estimates are local. They measure how customers best respond to exogenous variation in anticipated neighbor adoption, and show that customers respond in a way that can only be explained by strategic incentives. However, they do not account for how strategic responses propagate to affect aggregate outcomes. I therefore turn to a structural model that traces how these responses propagate into aggregate outcomes.

As in the two-player theory in section 2, the structural model is a dynamic network game in which customers face directed externalities from their neighbors' adoption that depend on their relative spatial locations and the existing infrastructure network. The main difference and modeling challenge is that, with $N = |R(\mathcal{E})|$ remaining customers, 2^N possible coalitions can form. The main structure I put on the problem is already familiar: I use the definitions of graphs, infrastructure states, network paths, supply costs, unconnected adopters, connected customers, connection requests, coalitions, and state transitions from subsection 5.1 and subsection 5.2. The main additional structure is to assume that coalition formation is within-period, following a specific microfounded process.

I nest a participation game within a dynamic optimal-stopping problem. In the inner game, prospective customers endogenously form coalitions within their neighborhood. The coalition splits infrastructure costs via Shapley cost shares, which are supermodular, so each additional member of the coalition strictly decreases individual charges. This creates strategic complements. It also makes the inner participation game a congestion game with an exact potential (Monderer and Shapley 1996). Changes in the exact potential function match the change in any customer's payoff whenever they unilaterally deviate from their current strategy. The process through which the coalition is formed follows Blume (1993). Customers receive random asynchronous revision opportunities to join the coalition; under the remaining conditions described below, the resulting coalition can be represented as a draw from a stationary Gibbs distribution, emulating the statistical mechanics of spin configurations in the Ising model.[55] In the outer game, customers solve an optimal-stopping problem, choosing the best moment in time to adopt as their neighbors expand the grid

55. Potentials replace Hamiltonians. The functional form of the Gibbs distribution is also known as a softmax in machine learning.

around them. Strategic substitutes are created by benefits of waiting for neighbor coalitions to expand the network first. The state space is the physical infrastructure, represented as a tree graph that fits Definition 4. Being connected to the grid is an absorbing state, so connecting is a terminal decision, and continuation values need only be calculated one step ahead. However, the information structure of the game is different from the standard application of conditional-choice-probability methods to dynamic games (Hotz and Miller 1993; Aguirregabiria, Collard-Wexler, and Ryan 2021). I adapt them to a new information structure that conditions on what customers observe: a tentative coalition Q_{-i} within the participation game, breaking the assumption that players move simultaneously and have conditional independence.

The games meet at the margin. Waiting is the outside option in the inner strategic-complements game, while joining the resulting coalition is the value of stopping in the outer strategic-substitutes game. In equilibrium, the marginal adoption payoff are the same in both games. Therefore, differences in the inner game's potential function match the difference between any customer's ex-ante choice values in the outer game, holding neighbor actions Q_{-i} fixed. Hotz–Miller inversion gives $\Pr(y_i = 1 \mid Q_{-i}, \mathcal{E})$, which I call *coalition-conditional choice probabilities (CoCCPs)* to distinguish them from standard conditional choice probabilities $\Pr(y_i = 1 \mid \mathcal{E})$. I use Bayes' rule to invert the CoCCPs into $\Pr(Q_{-i} \mid y_i, \mathcal{E})$, an *action-conditioned coalition distribution* within neighborhoods. I show this is a function of the inner game's potential function. The waiting-conditioned version of this $\Pr(Q_{-i} \mid y_i = 0, \mathcal{E})$ allows me to integrate over the state transitions induced by coalitions of neighbors, so prospective customers can form expectations of how the grid expands if they were to wait. A fixed point between coalition formation and the future value of waiting is the equilibrium.

I estimate this fixed point sequentially (Aguirregabiria and Mira 2007). First, I estimate marginal CCPs. The reduced-form state space includes row maxima and sums that measure inward externalities, as in the reduced form, and the spectral radius of the externality matrix, which measures the severity of the public-goods problem (Elliott and Golub 2019). Second, I estimate a myopic Gibbs participation game that sets the value of waiting to zero, using the CCPs as starting values. This initializes the degree of strategic complementarity. Then I start a sequential loop. I construct one-step-ahead continuation values by integrating state valuations over the waiting-conditioned coalition distribution implied by potential functions. I estimate the structural parameters of the game by maximizing the Gibbs likelihood of the realized coalition conditional on those continuation values. The estimated Gibbs coalition distribution then gives a potential for each coalition, which feeds back into the one-step continuation values and the CoCCPs. This process repeats until it reaches a fixed point.

I present the timing and information structure, payoffs, the household's dynamic problem, the coalition subgame, and their equilibrium in turn, before turning to estimation and identification.

8.1. Timing, Information, and Strategies

Recall that grid-relevant *neighborhoods* $\bar{Q}_t \in \bar{\mathcal{Q}}_t$ partition the remaining unconnected customers R_t by their most upstream unbuilt shared infrastructure. Each neighborhood $\bar{Q}_t$ plays a participation game within each period, which I will call a *meeting*, to negotiate collective procurement of a network extension. The result of the game is a realized coalition $Q_t \subseteq \bar{Q}_t$. The inner game is nested within the outer game as follows:

1. The period starts, and each member of the neighborhood $\bar{Q}_t \in \bar{\mathcal{Q}}_t$ observes the physical infrastructure state $\mathcal{E}_t$.

2. An initial coalition $Q_t^{\ell=0} \subseteq \bar{Q}_t$ is proposed, where $\ell = 0$ indexes the meeting stage.

3. Each $i \in \bar{Q}_t$ has an independent Poisson clock with mean waiting time r. Each ring gives customer i a *revision opportunity*, which advances the *meeting index* from ℓ to $\ell + 1$.

4. At each revision opportunity:

   - Customer i observes the tentative coalition formed by their neighbors $Q_{-i,t}^{\ell} \subseteq \bar{Q}_t \setminus \{i\}$.
   - They draw a Type-I extreme-value shock for each action, $\varepsilon_{it\ell}(0), \varepsilon_{it\ell}(1) \overset{\text{iid}}{\sim} \text{EV1}(0, \sigma_\epsilon)$, which are independent across customers, actions, revision opportunities, weeks, and meetings.
   - They choose to adopt or not by selecting $y_{it\ell} \in \{0, 1\}$. This updates the coalition to $Q_t^{\ell+1} = Q_{-i,t}^{\ell} \cup \{i\}$ if they join the coalition, or to $Q_t^{\ell+1} = Q_{-i,t}^{\ell}$ if they stay out of it.
   - The meeting continues.

5. The meeting suddenly ends at an exogenous, unknown time that is unobserved in advance. Let $\ell_i^*(t)$ denote customer i's last revision opportunity before the deadline. At the deadline, participation, private values, and the coalition are all fixed

$$y_{it} = y_{it\ell_i^*(t)}, \qquad \varepsilon_{it}(y) = \varepsilon_{it\ell_i^*(t)}(y), \qquad Q_t = Q_t^{\max_i \ell_i^*(t)} = \{i : y_{it} = 1\},$$

   and no customer can revise their strategy again.

6. The coalition Q_t submits a joint connection request to the utility, splitting payment of their collective connection charge via the Shapley rule.

7. The utility observes all connection requests within the period across each distribution network and aggregates them $Y_t = \bigcup_{Q \in \mathcal{Q}_t} Q$, building the required network extensions $\mathcal{B}_t(Q)$ for each $Q \in \mathcal{Q}_t$ at the end of the period.

8. This evolves the grid state according to the transition function $\mathcal{E}_{t+1} = T(Y_t, \mathcal{E}_t) = \mathcal{E}_t \cup \bigcup_{Q \in \mathcal{Q}_t} \mathcal{B}_t(Q)$, which builds all the nodes necessary to connect each Q. The connecting customers in Y_t exit the game before $t+1$, so $R_{t+1} = R_t \setminus Y_t$ and $Y_{t+1}^* = Y_t^* \cup Y_t$, and each member of a neighborhood is in a weakly smaller one.

The meeting has an asynchronous Poisson revision opportunity and a logit private value structure. The structure of the meeting results in a stationary distribution through four microfoundations. First, when $r \to 0$ and alarms come quickly enough so that customers can revise many times before the deadline, the coalition distribution becomes well-mixed, so the initial proposal Q_t^0 no longer affects the outcome. Second, a customer's new decision depends on a fresh private value draw and not their historical private shocks, so each draw forces them to reoptimize. Third, both own and neighbor decisions are reversible. Finally,

as the deadline is unanticipated, behavior is stationary through the end of the meeting. Together these properties imply that at each revision opportunity, customers make adoption decisions by observing only their current adopting neighbors $Q^{\ell}_{-i,t}$, the state, and their own private shocks. I assume that $r \to 0$, so the resulting distribution of coalitions is stationary throughout and can therefore be described as a draw from a stationary Gibbs distribution (Blume 1993).

This timing changes the within-period information structure of a baseline Markov-perfect Nash equilibrium dynamic game, such as the one described in the review by Aguirregabiria, Collard-Wexler, and Ryan (2021). There, players only observe the common state and their private shock before *simultaneously* choosing a realized action, so the joint action distribution that determines state transitions is fully described by *independent* individual CCPs $\Pr(\mathbf{y}_t = \mathbf{y} \mid \mathcal{E}_t) = \prod_{j \in R_t} \Pr(y_{jt} = y_j \mid \mathcal{E}_t)$. With this inner meeting, customers additionally observe tentative coalition membership $Q^{\ell}_{-i,t}$ at their revision opportunity before choosing their action $y_{it\ell}$, so their actions are not independent and are instead *correlated*. Therefore, the joint action distribution is not conditionally independent of i's choices. Given the tentative coalition, customers act as Bayesians. They can only noisily best respond because neighbors may receive revision opportunities after them, so they integrate over unobserved neighbor private shocks and anticipate that their neighbors' decisions may change before the deadline. The stationary Gibbs distribution allows me to condition CCPs on coalition realization, allowing me to nest the strategic complements game within the dynamic Markov Perfect Equilibrium.

The customer's information set now includes tentative coalition Q_{-i}, which I use to denote a generic realization of $Q^{\ell}_{-i,t}$ at the beginning of i's revision opportunity. At the start of the revision opportunity, but prior to receiving a private shock, i integrates over their own anticipated private shocks to form a *coalition-conditional choice probability (CoCCP)*,

$$\sigma_i(Q_{-i}, \mathcal{E}_t) \equiv \Pr\Big(y_{it\ell} = 1 \mid Q^{\ell}_{-i,t} = Q_{-i}, \mathcal{E}_t\Big) \in [0,1], \qquad Q_{-i} \subseteq \bar{Q}_t \setminus \{i\}, \tag{29}$$

which describes customer i's strategy at every meeting index ℓ conditional on observing each tentative coalition Q_{-i}. When the meeting terminates at $\ell = \ell^*_i(t)$, the interim choice $y_{it\ell^*_i(t)} = 1$ becomes i's action, and CoCCPs match the notation of the introduction, $\sigma_i = \Pr(y_{it} = 1 \mid Q_{-i,t}, \mathcal{E}_t)$. Let $\boldsymbol{\sigma} = (\sigma_i)_{i \in \bar{Q}_t}$ denote the neighborhood's strategy profile. If each customer maximizes by choosing their optimal action given neighbor strategies, then $\boldsymbol{\sigma}$ represents a Markov-Perfect Bayesian Nash equilibrium in the inner game.

8.2. Payoffs

The structural model simplifies the game by assuming customers pay the *material-based connection charges* for the network-path-based infrastructure, rather than distance-based charges that built the minimum technical solution, and that the graph $\mathscr{G}$ is common knowledge. The total supply-side cost of extending infrastructure to provide access to Q is $C(Q, \mathcal{E}_t) = \sum_{e \in \mathcal{B}_t(Q)} w_e$, recalling the coalition's required extension elements $\mathcal{B}_t(Q)$ and their supply costs w_e from subsection 5.1. The utility passes through supply costs at a constant rate α following the mechanical baseline in Definition 1.

A realized coalition Q divides the resulting connection charge using Shapley values, which is the unique allocation satisfying efficiency, symmetry, the null-player condition, and additivity (Shapley 1953). Within

this private provision of infrastructure game, efficiency means that the connection charge is fully paid for, symmetry means that customers split the cost of a specific edge or node amongst those who will use it, the null-player condition means that adding a customer under a service drop doesn't change the cost, and additivity implies that the customer's overall cost share is the sum of each component's cost share. On a rooted tree, Shapley is applied element by element along the shared network path, so customers only share the costs of the infrastructure they require (Littlechild and Owen 1973; Littlechild and Thompson 1977; Márkus, Pintér, and Radványi 2012). Specifically, customer $i \in Q$ pays

$$c_i(Q, \mathcal{E}_t) = \alpha \sum_{e \in \mathcal{B}_t(i)} \frac{w_e}{|\{j \in Q : e \in \mathcal{B}_t(j)\}|} = \alpha \sum_{e \in \mathcal{B}_t(i)} \frac{w_e}{n_e(Q)}, \tag{30}$$

where the denominator is the number of members of Q who use and share element e. I refer to $c_i(Q, \mathcal{E}_t)$ as coalitional *cost shares*. The coalitional cost share is bounded above by a unilateral or singleton coalition connection charge and bounded below by the grand coalition connection charge $c_i(\{i\}, \mathcal{E}_t) \geq c_i(Q, \mathcal{E}_t) \geq c_i(\bar{Q}_t, \mathcal{E}_t)$. Each additional participant in a coalition weakly reduces $c_i(Q, \mathcal{E}_t)$ implying supermodularity. These charges fund engineering resource costs at rate α.

While the utility receives payments of $\sum_i c_i(Q, \mathcal{E}_t)$ if coalition Q adopts, customers may incur cooperation costs, frictions, or alternatively, demand-side economies of scale. The coordination-adjusted *connection charge* that enters i's payoff is defined as

$$\kappa_i(Q; \rho, \mathcal{E}_t) = \alpha \sum_{e \in \mathcal{B}_t(i)} w_e n_e(Q)^{-\rho}. \tag{31}$$

This matches section 2, where coordination costs ρ decrease the payoff value of cooperating. The exponential form preserves supermodularity as long as $\rho \geq 0$ because another marginal customer in the coalition still lowers cost. I estimate ρ, which will be a key parameter for welfare outcomes. If $\rho = 1$, there are neither coordination costs nor benefits, so the coordination-adjusted connection charge matches the coalitional cost share in (30). I interpret $\kappa_i(Q; \rho, \mathcal{E}_t) - c_i(Q, \mathcal{E}_t)$ as a coordination wedge and deadweight loss. When $\rho < 1$, there are cooperation frictions, creating a welfare loss of size $\alpha \sum_e w_e \left(n_e(Q)^{1-\rho} - 1\right)$. If $\rho > 1$, then there are demand-side economies from coordination, with welfare benefits of the same size. Both cases affect welfare through customer utility and therefore affect consumer surplus. The player will convert between utility and cost units at rate $\lambda > 0$.

In addition to the engineering-based connection charges, I impose a fixed cost $\chi \geq 0$ that accounts for the non-engineering cost of procuring infrastructure. It is meant to represent negotiation, transaction, labor, transportation, and survey costs associated with requesting and coordinating the utility to expand the distribution network. However, it does not include meters, wiring, or other per-adopter costs, which are not separately identified from flow benefits. It helps to explain weeks in which no one adopts, as coalitions form in only 13.6% of transformer-weeks. Computationally, I attach the fixed cost to the coalition's anchor node, and the realized coalition splits the fixed cost equally, $\chi/|Q|$. Service-drop customers pay no fixed cost, since they do not need a pole.

The utility occasionally makes massive expansions that are supply-side driven and subsidized (see

Assumption 2). Large expansions mostly occur in the first weeks after a transformer is placed. I parameterize massive expansions in the model by assuming that multilateral funders or the utility absorb some of the cost in the initial weeks of transformer placement. Let $a_t^g \geq 0$ be the number of weeks since the transformer was placed, and define

$$m(a_t^g; h) = 1 - e^{-h a_t^g}, \qquad h \geq 0. \tag{32}$$

Combining the fixed costs and connection charges, customer i's overall connection charge as part of coalition Q is $m(a_t^g; h)\,[\kappa_i(Q; \rho, \mathcal{E}_t) + \chi/|Q|]$ when their path requires a pole, and $m(a_t^g; h)\,\kappa_i(\{i\}; \rho, \mathcal{E}_t)$ if they are already under a service drop when their neighborhood is a singleton. In week one, the first opportunity to connect, pass-through is $\alpha(1 - e^{-h})$, and the half-life of the "massive expansion" utility funding is $\ln(2)/h$ weeks.

Prospective customers pay utility connection charges to obtain infrastructure access. The point of access is to avoid the costly transportation of electricity, liquids, or gases, which collectively generates *flow benefits*. Let $b_i(x_i)$ denote customer i's per-period flow benefit from utility infrastructure access net of consumption tariffs, where customer characteristics x_i are common knowledge to the other customers in the game. Customers discount the future at rate δ. As a week-t connection request provisions the extension between t and $t+1$, flow benefits begin in the following week, so the primitive willingness to pay for a connection is $\delta b_i(x_i)/(1-\delta)$, the present value of future flow benefits.

As in the reduced form, already-connected neighbors provide social peer effects that can raise the perceived value of adopting. I use the same measure as before, with $\bar{y}_{it}^D = [\mathbf{D}^{(-1)}\mathbf{y}_t^*]_i$ following Equation 10. Its linear functional form allows for the decomposition into current adopters and future adopters $\bar{y}_{i,t+1}^D = \bar{y}_{it}^D + \sum_j D_{ij}^{(-1)}\,\mathbf{1}\{j \in Y_t\}$, and so the former will enter static payoffs and the latter enters waiting values. One assumption is that social peer exposure affects only the terminal adoption value: connection is a terminal action, so exposure changes arriving after connection are payoff-irrelevant to the connected customer. A waiting customer still anticipates exposure growth, because $\bar{y}_{it}^D$ evolves as neighbors connect. This is a reasonable simplification because social peer effects are small on the margin—one connected peer 100 meters away raises the weekly adoption hazard by only 0.007 percentage points—and effects arriving only after connection are further discounted.

8.3. The Household's Dynamic Connection Problem

From the perspective of the inner game, continuation values are an outside option. By contrast, in the outer game, customers choose the best moment in time to connect to utility infrastructure as an optimal stopping problem. Connecting one week earlier gives a customer an additional week of flow benefits, but it commits them to paying their share of the coalition-adjusted connection charge. Waiting for neighbor coalitions to expand the grid, meanwhile, reduces their future connection charge. The model differs from the standard dynamic game because a customer who is pivotal to a coalition accounts for how their marginal choices affect coalition formation. This tradeoff depends on the fixed point of the inner and outer games. I describe the customer's Bellman equation, the coalition subgame, and their equilibrium in turn.

I start with the Bellman. The set of possible coalition configurations faced by customer i is $\mathcal{Q}_{-i}(\mathcal{E}_t) = 2^{\bar{Q}_t \setminus \{i\}}$, the power set of their neighborhood excluding i. These configurations have a stationary distribution

$\eta_i(Q_{-i} \mid \mathcal{E}_t)$, which I derive later. The ex-ante value of a state integrates over both private value shocks and the stationary distribution of coalition configurations:

$$V_i(\mathcal{E}_t) = \int_{\mathcal{Q}_{-i}(\mathcal{E}_t)} \mathbb{E}_\varepsilon \left[\max_{y \in \{0,1\}} \{v_i^y(Q_{-i}, \mathcal{E}_t) + \varepsilon_i(y)\} \right] d\eta_i(Q_{-i} \mid \mathcal{E}_t). \tag{33}$$

Customers compare two ex-ante choice values, $v_i^1(Q_{-i}, \mathcal{E}_t)$ and $v_i^0(Q_{-i}, \mathcal{E}_t)$, given their private shock distribution for each possible neighbor coalition. The customer integrates their optimal choices over the stationary distribution of neighborhood configurations $\eta_i(Q_{-i} \mid \mathcal{E}_t)$ to form state values.

I decompose each ex-ante choice value into its components. In the value of a connection, $v_i^1(Q_{-i}, \mathcal{E}_t)$, the terminal value of infrastructure access can be summarized as

$$B_i(\mathcal{E}_t) \equiv \frac{\delta b_i(x_i)}{1-\delta} + \theta_D \bar{y}_{it}^D. \tag{34}$$

I estimate $B_i(\mathcal{E}_t) = x_i(\mathcal{E}_t)'\theta_\pi$, where $x_i(\mathcal{E}_t)$ includes customer characteristics and social peer exposure, and normalize flow payoffs from being unconnected to zero. While the flow benefits of staying unconnected are non-zero in real life, because customers have kerosene, solar lanterns, wood, and other ways of procuring energy services, those options are unaffected by the grid state. Only the net flow benefits matter for a one-off connection decision.

To define the option value of waiting, I need the transition function for the physical state induced by equilibrium strategies in the inner game:

$$\begin{aligned} \Psi_i(\mathcal{E}_t' \mid \mathcal{E}_t, y; \boldsymbol{\sigma}) &\equiv \Pr_{\boldsymbol{\sigma}} \left(\mathcal{E}_{t+1} = \mathcal{E}_t' \mid \mathcal{E}_t = \mathcal{E}_t, y_{it} = y \right) \\ &= \int_{Y \subseteq R_t} \underbrace{\mathbf{1}\left\{ \mathcal{E}_t' = T(Y, \mathcal{E}_t) \right\}}_{\substack{\text{network expansion} \\ \text{induced by } Y}} \underbrace{d \Pr_{\boldsymbol{\sigma}} \left(Y \mid y_{it} = y, \mathcal{E}_t \right)}_{\substack{\text{conditional equilibrium} \\ \text{distribution of adopters}}}. \end{aligned} \tag{35}$$

Conditional on the current state, an action y, and the strategies in the inner game $\boldsymbol{\sigma}$, the transition function describes how deterministic network expansions $T(Y, \mathcal{E}_t)$ occur. The connecting set of customers $Y_t = \bigcup_{Q \in \mathcal{Q}_t} Q$ is the union of all realized coalitions on the distribution network. Via the inner game, customer i understands that their actions affect which coalitions form in $\Pr_{\boldsymbol{\sigma}}(Y \mid y_{it} = y, \mathcal{E}_t)$.

Now we specify ex-ante choice values in the outer game. Setting $y = 0$ allows me to use the transition function to compute the value of waiting for the grid to expand:

$$\begin{aligned} W_i(\mathcal{E}_t | y_{it} = 0) &= \delta \int V_i(\mathcal{E}_t')\, d\Psi_i(\mathcal{E}_t' \mid \mathcal{E}_t, 0; \boldsymbol{\sigma}) \\ &= \delta \int_{Y \subseteq R_t \setminus \{i\}} \underbrace{V_i(T(Y, \mathcal{E}_t))}_{\text{Continuation value}} \underbrace{d \Pr_{\boldsymbol{\sigma}} \left(Y \mid y_{it} = 0, \mathcal{E}_t \right)}_{\substack{\text{conditional equilibrium distribution} \\ \text{of the connecting set}}}. \end{aligned} \tag{36}$$

The value of waiting is the discounted expectation of future state valuations, conditional on customer i choosing to wait while understanding how that choice shifts the realized adopter set Y. Now return to the

optimization problem in Equation 33. For a specific coalition configuration Q_{-i} observed by customer i at the beginning of the revision opportunity, the coalition-conditional ex-ante choice values are

$$v_i^1(Q_{-i}, \mathcal{E}_t) = B_i(\mathcal{E}_t) - \lambda m(a_t^g; h) \left[\kappa_i(Q_{-i} \cup \{i\}; \rho, \mathcal{E}_t) + \mathbf{1}\{P_{it}^{\text{path}} > 0\} \frac{\chi}{|Q_{-i}| + 1} \right], \tag{37}$$

$$v_i^0(Q_{-i}, \mathcal{E}_t) = W_i(\mathcal{E}_t | y_{it} = 0). \tag{38}$$

The connecting value are the net static incentives for adoption. It depends on Q_{-i} because its members divide the connection charge according to Shapley values. The value of not connecting is the value of waiting for the grid to expand. It depends on Q_{-i} because neighbor coalitions expand the grid. Unlike the standard game, how neighbors expand the grid depends on i's action, and i understands that when they wait, they affect their neighbors conditional equilibrium distribution. However, Equation 38 also shows that the tentative coalition has no further effect on the outside value except through grid evolution. While these ex-ante coalition-conditional choice values have additional information relative to the baseline dynamic game, we will see that they can still be Hotz-Miller inverted, allowing me to form CoCCPs.

8.4. Participation Potential Subgame

In the inner subgame, the state $\mathcal{E}_t$ and customer i's neighborhood $\bar{Q}_{it}$ are fixed at the beginning of the week. Within a revision opportunity, the tentative coalition $Q_{-i} \subseteq \bar{Q}_{it} \setminus \{i\}$ is known, and the ex-ante net payoff from joining the coalition before receiving a private value draw is $v_i^1(Q_{-i}, \mathcal{E}_t) - v_i^0(Q_{-i}, \mathcal{E}_t)$. Within the outer game, the value of waiting is an outside option. The customer compares it to static benefits given the tentative coalition they observe and their private value draws.

The subgame follows the definition of a *congestion model* (Rosenthal 1973; Monderer and Shapley 1996), though payoffs are inverted; more customers adopting and sharing each piece of infrastructure makes it cheaper per person to collectively procure. To make the congestion model analogy precise, the finite player set consists of the neighborhood's prospective customers $\bar{Q}_t$. Facilities are the unbuilt poles and lines of the neighborhood's network extension, the elements $e \in \mathcal{B}_t(\bar{Q}_t)$. Each customer has a unique path $\mathcal{B}_t(i)$ to the transformer that is fixed by the physical state $\mathcal{E}_t$. Therefore they have two routing strategies in the congestion game: they either use the bundle of poles and lines on their path, or a private outside facility with zero payoff. A payoff is associated with each node and edge e. Let k be the number of customers in the coalition using each element. I can then apply Monderer and Shapley (1996)'s notation directly to each edge, fixed cost, and payoff:

- The supply cost of each pole or line is divided among its k users, $c_e(k) = -\lambda m(a_t^g; h)\, \alpha w_e k^{-\rho}$.
- Fixed costs of network extension are split equally among the coalition, $c_p(k) = -\lambda m(a_t^g; h)\, \chi k^{-1}$.
- The benefit of obtaining electricity access, $c_i(1) = B_i(\mathcal{E}_t) - W_i(\mathcal{E}_t)$, is located at customer i's final node, which is the endpoint of their routing strategy.

As Shapley cost shares are symmetric and decline with an additional customer, payoffs have increasing differences with the coalition size, but have diminishing marginal returns.

I describe the inner subgame as a congestion game because every congestion game is a *potential game*: there exists a function $\Phi : 2^{\bar{Q}_t} \to \mathbb{R}$ over coalitions such that the change under any unilateral deviation is exactly that customer's change in payoff (Monderer and Shapley 1996, Theorem 3.1). The coalition subgame I describe has the exact potential

$$\Phi_t(Q; W) = \sum_{i \in Q} [B_i(\mathcal{E}_t) - W_i(\mathcal{E}_t)] - \lambda m(a_t^g; h) \left[\mathbf{1}\{\mathcal{P}^{\text{path}}(\bar{Q}_t, \mathcal{E}_t) \neq \varnothing\} \chi H_{|Q|}(1) + \alpha \sum_e w_e H_{n_e(Q)}(\rho) \right], \tag{39}$$

where $H_n(r) \equiv \sum_{k=1}^{n} k^{-r}$ and $H_0(r) \equiv 0$ are generalized harmonic numbers of order r. The indicator removes the fixed cost for the customers within a service drop, who are always in a singleton neighborhood. I show that differences in the potential are exactly equal to the difference in ex-ante choice values of the outer game in the following proposition.

Proposition 4 (Exact potential ex-ante choice value equivalence). For every $Q_{-i} \subseteq \bar{Q}_{it} \setminus \{i\}$,

$$\Delta_i \Phi_t(Q_{-i}; W) = \Phi_t(Q_{-i} \cup \{i\}; W) - \Phi_t(Q_{-i}; W) = v_i^1(Q_{-i}, \mathcal{E}_t) - v_i^0(Q_{-i}, \mathcal{E}_t). \tag{40}$$

Proof. Adding i to the coalition adds their net static payoff $B_i(\mathcal{E}_t) - W_i(\mathcal{E}_t)$ to Φ's first term. For each edge on i's path $e \in \mathcal{B}_t(i)$, the number of users of each element e rises from $n_e(Q_{-i})$ to $n_e(Q_{-i}) + 1$. The difference in each element's harmonic term is

$$H_{n_e+1}(\rho) - H_{n_e}(\rho) = \sum_{k=1}^{n_e+1} k^{-\rho} - \sum_{k=1}^{n_e} k^{-\rho} = (n_e + 1)^{-\rho}. \tag{41}$$

As each term in the harmonic sum up to n_e cancels, the marginal change in the potential for each element is the share i pays for each e, $\alpha w_e (n_e + 1)^{-\rho}$. Similarly, for the fixed cost, an additional customer splitting it gives $H_{|Q_{-i}|+1}(1) - H_{|Q_{-i}|}(1) = (|Q_{-i}| + 1)^{-1}$, which is i's share of the fixed cost if they join the coalition. Scaling both the fixed and engineering charges by the marginal utility of money λ and initial supply-side incentives $m(a_t^g; h)$ gives the second term of Φ. For a marginal coalition member, the difference in potential is the difference in coalition-conditional ex-ante choice values because

$$\begin{aligned}
&\Phi_t(Q_{-i} \cup \{i\}; W) - \Phi_t(Q_{-i}; W) \\
&= B_i(\mathcal{E}_t) - W_i(\mathcal{E}_t) - \lambda m(a_t^g; h) \left[\mathbf{1}\{P_{it}^{\text{path}} > 0\} \frac{\chi}{|Q_{-i}| + 1} + \kappa_i(Q_{-i} \cup \{i\}; \rho, \mathcal{E}_t) \right] \\
&= v_i^1(Q_{-i}, \mathcal{E}_t) - v_i^0(Q_{-i}, \mathcal{E}_t),
\end{aligned}$$

where the last equality follows from Equation 37 and Equation 38. □

Shapley cost shares are the unique efficient rule under which the participation game admits an exact potential (Monderer and Shapley 1996; Hart and Mas-Colell 1989). The critical difference between Proposition 4 and Sections 3 and 6 of Monderer and Shapley (1996) is that I incorporate ρ, which measures the degree of strategic complements. In particular, if $\rho \neq 1$, Proposition 4 shows that the potential function

survives despite inefficiencies or scale economies. The component-wise split assumed by Shapley values allows me to estimate the Gibbs coalition distribution by maximum likelihood.

Having an exact potential also provides additional properties of the inner game that are helpful or intuitive. First, the potential tracks how additional coalition members change aggregate payoffs. Whenever a member deviates from a strategy, the potential only tracks their payoff change. This will be very useful in the next section for linking the outer game to the inner one. Second, every local maximum of the potential is a pure-strategy Nash equilibrium, so the sequence of value shocks within the meeting selects between potential maxima, in both senses of the term, because an unbounded Type I extreme value can shock the equilibrium out of its basin of attraction. Therefore, the Gibbs distribution describes the coalition distribution that results from the random sequence of shocks. Third, multiple local maxima in the potential function are separate multiple equilibria. This allows me to integrate over the distribution of the multiple equilbria given the private shocks instead of choosing between the least upper bound and greatest lower bound when using alternative methods involving Tarksi's fixed point algorithm or similar strategies (Jia 2008; Ackerberg and Gowrisankaran 2006).

8.5. Equilibrium Between the Inner and Outer Games

Customers choose to stop in the outer problem and participate in the inner problem when they choose to adopt. They consider their observation of the tentative coalition, the physical state, and also how their choices affect coalition formation. At a revision opportunity, a customer observes the tentative coalition and then chooses

$$y_{it\ell} = \arg\max_{y\in\{0,1\}} \left\{ v_i^y(Q_{-i}, \mathcal{E}_t) + \varepsilon_{it\ell}(y) \right\}.$$

They make this choice using three objects, all of which can be expressed in terms of the potential function.

Proposition 5 (CoCCPs and coalition distributions)**.** Assume a fixed neighborhood $\bar{Q}_t$, physical infrastructure state $\mathcal{E}_t$, and conjectured waiting value W. Under the meeting timing and stationarity assumption in subsection 8.1 and the exact potential of the game in Proposition 4, the following statements hold for every $i \in \bar{Q}_t$ and $Q_{-i} \subseteq \bar{Q}_t \setminus \{i\}$.

(i) Hotz–Miller inversion and Proposition 4 imply that potential differences, choice-value differences, and CoCCP log odds are equal,

$$\Phi_t(Q_{-i} \cup \{i\}; W) - \Phi_t(Q_{-i}; W) = v_i^1(Q_{-i}, \mathcal{E}_t) - v_i^0(Q_{-i}, \mathcal{E}_t) = \sigma_\epsilon \log \frac{\sigma_i(Q_{-i}, \mathcal{E}_t)}{1 - \sigma_i(Q_{-i}, \mathcal{E}_t)}. \tag{42}$$

(ii) Therefore, the coalition-conditional choice probability (CoCCP) can be expressed as a logit in potential differences.

$$\sigma_i(Q_{-i}, \mathcal{E}_t) = \Pr(y_{it\ell} = 1 \mid Q^\ell_{-i,t} = Q_{-i}, \mathcal{E}_t) = \frac{\exp\{\Delta_i \Phi_t(Q_{-i}; W)/\sigma_\epsilon\}}{1 + \exp\{\Delta_i \Phi_t(Q_{-i}; W)/\sigma_\epsilon\}}. \tag{43}$$

(iii) The overall coalition distribution is a Gibbs state (Blume 1993):

$$\mu(Q \mid \bar{Q}_t, \mathcal{E}_t) = \frac{\exp\{\Phi_t(Q; W)/\sigma_\epsilon\}}{\sum_{\widetilde{Q} \subseteq \bar{Q}_t} \exp\{\Phi_t(\widetilde{Q}; W)/\sigma_\epsilon\}}, \qquad Q \subseteq \bar{Q}_t. \tag{44}$$

(iv) The neighborhood-coalition distribution facing i is also stationary:

$$\begin{aligned}\eta_i(Q_{-i} \mid \mathcal{E}_t) &= \mu(Q_{-i} \mid \bar{Q}_t, \mathcal{E}_t) + \mu(Q_{-i} \cup \{i\} \mid \bar{Q}_t, \mathcal{E}_t) \\ &= \frac{\exp\{\Phi_t(Q_{-i}; W)/\sigma_\epsilon\} + \exp\{\Phi_t(Q_{-i} \cup \{i\}; W)/\sigma_\epsilon\}}{\sum_{\widetilde{Q} \subseteq \bar{Q}_t} \exp\{\Phi_t(\widetilde{Q}; W)/\sigma_\epsilon\}}, \qquad Q_{-i} \subseteq \bar{Q}_t \setminus \{i\}.\end{aligned} \tag{45}$$

Proof. The full proof appears in subsection B.6; I provide a brief summary here. Step *(i)* follows from Hotz–Miller inversion applied to coalition-conditional ex-ante choice values. Step *(ii)* rearranges to isolate σ. The Gibbs state in *(iii)* adapts the proof of Theorem 5.1 in Blume (1993). The neighbor-coalition distribution in *(iv)*, which excludes i, sums the Gibbs probabilities of Q_{-i} and $Q_{-i} \cup \{i\}$. □

Proposition 5 strings together all the layers of the outer and inner games. Potential differences from unilateral deviations at the revision-opportunity level determine CoCCPs conditional on the configuration customers face. Potential levels determine the Gibbs state over each coalition, and removing i from the distribution gives the neighbor coalition distribution that i faces just before they receive a revision opportunity. These objects allow me to move smoothly from within-meeting behavior to ex-ante choice values. Additionally, the proposition shows that the CoCCPs and the joint and marginal coalition distributions for i can all be written as normalized potential weights. The logit is the special case in which only the two configurations that differ in i's membership are feasible.

What remains is to construct analogues of the baseline dynamic-game objects that condition on the observed coalition. At a revision opportunity, the coalition-conditional inclusive value, the inner term of the Bellman equation, can be expressed as

$$\begin{aligned}\mathcal{V}_i(Q_{-i}, \mathcal{E}_t) &\equiv \mathbb{E}_\varepsilon \left[\max_{y \in \{0,1\}} \{v_i^y(Q_{-i}, \mathcal{E}_t) + \varepsilon_i(y)\} \right] \\ &= \sigma_\epsilon \log \left[\exp\left\{ \frac{v_i^0(Q_{-i}, \mathcal{E}_t)}{\sigma_\epsilon} \right\} + \exp\left\{ \frac{v_i^1(Q_{-i}, \mathcal{E}_t)}{\sigma_\epsilon} \right\} \right] + \sigma_\epsilon \gamma \\ &= v_i^0(Q_{-i}, \mathcal{E}_t) - \sigma_\epsilon \log\left[1 - \sigma_i(Q_{-i}, \mathcal{E}_t)\right] + \sigma_\epsilon \gamma \\ &= W_i(\mathcal{E}_t) + \sigma_\epsilon \log \left[1 + \exp\left\{ \frac{\Delta_i \Phi_t(Q_{-i}; W)}{\sigma_\epsilon} \right\} \right] + \sigma_\epsilon \gamma.\end{aligned} \tag{46}$$

where γ is Euler's constant. The first equality applies the expected-maximum formula for two independent Type-I extreme-value shocks. The second uses $1 - \sigma_i(Q_{-i}, \mathcal{E}_t) = \frac{e^{v_i^0/\sigma_\epsilon}}{e^{v_i^0/\sigma_\epsilon} + e^{v_i^1/\sigma_\epsilon}}$, the probability of waiting. The final equality substitutes the value of waiting and rewrites the sum using potential differences.

The law of total probability expresses the CCP as a function of the CoCCP and the neighbor-coalition distribution:

$$\Pr(i \in Y_t \mid \mathcal{E}_t) = \int_{\mathcal{Q}_{-i}(\mathcal{E}_t)} \sigma_i(Q_{-i}, \mathcal{E}_t) d\eta_i(Q_{-i} \mid \mathcal{E}_t). \tag{47}$$

This marginal probability provides the denominator for Bayes' rule, which gives the *action-conditioned coalition distributions*:

$$\omega_i^1(Q_{-i} \mid \mathcal{E}_t) \equiv \Pr(Q_{-i} \mid i \in Y_t, \mathcal{E}_t) = \frac{\eta_i(Q_{-i} \mid \mathcal{E}_t)\sigma_i(Q_{-i}, \mathcal{E}_t)}{p_i(\mathcal{E}_t)} = \frac{\mu(Q_{-i} \cup \{i\} \mid \bar{Q}_t, \mathcal{E}_t)}{p_i(\mathcal{E}_t)}, \tag{48}$$

$$\begin{aligned}\omega_i^0(Q_{-i} \mid \mathcal{E}_t) \equiv \Pr(Q_{-i} \mid i \notin Y_t, \mathcal{E}_t) &= \frac{\eta_i(Q_{-i} \mid \mathcal{E}_t)[1 - \sigma_i(Q_{-i}, \mathcal{E}_t)]}{1 - p_i(\mathcal{E}_t)} = \frac{\mu(Q_{-i} \mid \bar{Q}_t, \mathcal{E}_t)}{1 - p_i(\mathcal{E}_t)} \\ &= \frac{\exp\{\Phi_t(Q_{-i}; W)/\sigma_\epsilon\}}{\sum_{\widetilde{Q}_{-i} \subseteq \bar{Q}_t \setminus \{i\}} \exp\{\Phi_t(\widetilde{Q}_{-i}; W)/\sigma_\epsilon\}}.\end{aligned} \tag{49}$$

These expressions describe how the probability of neighbor coalitions depends on what i does, allowing i to understand how pivotal they are to neighbor coalition formation. The latter, waiting-conditioned coalition distribution is the distribution needed to form one-step-ahead continuation values, and Equation 49 shows that it can be expressed in terms of potentials.

The distributions over coalition configurations are specific to neighborhoods $\bar{Q}_t \in \bar{\mathcal{Q}}_t$, which partition the set of remaining unconnected customers R_t. Conditional on the physical state, the adoption distribution factors into neighborhood-specific distributions because connection charges never cross over into separate neighborhoods, but one can recover the full joint adoption distribution as their product

$$\Pr(Y_t = Y \mid \mathcal{E}_t) = \prod_{\bar{Q}_t \in \bar{\mathcal{Q}}_t} \mu(Y \cap \bar{Q}_t \mid \bar{Q}_t, \mathcal{E}_t). \tag{50}$$

One potential objection is that social-network peer effects may prevent independence across neighborhoods. However, the expression above is an ex-ante probability conditional only on the state. Social-network peer exposure $\bar{y}_{it}^D$ is predetermined within periods, so it never introduces a cross-neighborhood term into the potential. Instead, only dynamics of social-network peer effects create cross-neighborhood interactions, and they are captured by W. An equilibrium can now be defined:

Definition 5 (Equilibrium)**.** Given payoff parameters θ, an *equilibrium* is a triple $(\boldsymbol{\sigma}, \mu, W)$ such that, for each neighborhood $\bar{Q}$ and physical state $\mathcal{E}_t$:

(i) In the inner game, $\boldsymbol{\sigma}$ and μ are the CoCCPs and coalition distribution characterized by Proposition 5, evaluated at W.

(ii) In the outer game, $\boldsymbol{\sigma}$ and μ induce the state values in Equation 33 and the adoption distribution in Equation 50; W is the waiting value implied by the triple through Equation 36.

Definition 5 defines a fixed point between the strategy profile governing revision opportunities in the inner game and the waiting value in the outer game. The main link between the two is the Gibbs state of coalition

formation, μ, which aggregates outcomes of the inner game into a stationary distribution and permits an equilibrium between the two games.

To summarize, nesting the participation game within the dynamic game is possible for three reasons. First, realized shocks all occur within the inner meeting, so all expectations of future states are pushed into the outer game. Second, dynamic state valuations only enter the inner meeting as an option value, so within the inner game, customers who expect neighbors to adopt without them do not join the coalition. Third, terminal actions mean that continuation values are always one step ahead, so λ, the marginal utility of money, affects only monetary payoffs in the current charge and the option value. These three facts allow me to nest a self-enforcing participation game of strategic complements within a dynamic game of strategic substitutes.

9 — Structural Estimation and Identification

I estimate the equilibrium fixed point in three steps: an initial conditional-choice-probability (CCP) projection, Gibbs maximum likelihood, and potential-game nested pseudo-likelihood.

9.1. Model Parameters

In this baseline run, I fix the weekly discount factor to $\delta = 0.9^{1/52}$. Future versions of the paper will incorporate a grid of discount factor assumptions. These are not typically identified in dynamic models. Other parameters that are fixed rather than estimated are the cost pass-through rate α, the social distance-decay exponent ζ, and the engineering cost structure C_i.

The main structural parameters of the model are a benefit vector θ_π, the marginal utility of money λ, the sharing exponent ρ, a fixed cost of expanding the grid χ, and the rate of initial supply-side incentive decline h. Summarize these in a single structural vector.

$$\theta \equiv (\theta_\pi, \lambda, \rho, \chi, h).$$

I use units of 1,000 RWF (approximately USD 1), so χ is measured in thousands of RWF and λ converts this to utility. The customer characteristics used are the same as in the reduced form, using building and demand characteristics, 2019 nightlights, terrain, road access, and, as described above, non-grid network exposure.

9.2. Initialization

While the full physical state of the grid is a graph object, I simplify the state space only in an initial step to form initial guesses of CoCCPs. I summarize it and then walk through the terms.

$$s_{it} = \left(P_{it}^{\text{path}}, \kappa_{it}^{\bar{Q}}, |Y_t^*|, \bar{y}_{it}^{D}, g_{\max,it}^{\kappa}, g_{\sum,it}^{\kappa}, \varrho_t, \text{age}_{\mathcal{E}_t,it}, x_i^{\text{terr}}, x_i^{\text{own}} \right). \tag{51}$$

First, I follow the reduced form and use two measures for the upper and lower bound costs in P_{it}^{path} and $\kappa_{it}^{\bar{Q}}$. Second, the number of connected peers is counted in $|Y_t^*|$, as well as the inverse-distance weighted

peer statistics at the beginning of the state $\bar{y}_{it}^{D}$. Third, I take the path-based externality matrix and, for each row customer i, take the row maximum and sum, $g^{\kappa}_{\max,it}$ and $g^{\kappa}_{\sum,it}$, as well as the externality matrix's spectral radius ϱ_t. The spectral radius summarizes the strength of cycles in the externality matrix, which can be used as a summary statistic for the degree of public goods (Elliott and Golub 2019). To account for war-of-attrition incentives, I use $\text{age}_{\mathcal{E}_t,it}$, which measures the number of periods since each customer last received an externality. Finally, x_i^{terr} controls for elevation, the local slope, the distance to the nearest road, and indicators (0–5) for the quality of the nearest road, and x_i^{own} are customer characteristics for log building area and height, a missing google building indicator, a non-residential indicator, and log 2019 nightlights. These covariates are assumed to be on the demand side, not on the supply side, and so unlike x_i^{terr}, I will use them to reconstruct willingness-to-pay. This reduction simplifies the state space for an initial spline estimate of CoCCPs.

$$\widehat{P}_i(1 \mid s_{it}) \approx P_i(1 \mid \mathcal{E}_t). \tag{52}$$

While the above procedure would be enough to estimate CCPs if they were independent, here adoptions are correlated because of cooperation: even though the average baseline adoption hazard is around 0.5 percentage points, the realized adoption is current and joint.

The Gibbs log-likelihood of the second stage is

$$\ell_G^{(k)}\left(\theta; W^{(k)}\right) \equiv \sum_g \sum_t \sum_{\bar{Q}_{gt} \in \bar{\mathcal{Q}}_{gt}} \left[\frac{\Phi_{gt}\left(Y_{gt} \cap \bar{Q}_{gt}; W^{(k)}, \theta\right)}{\sigma_\epsilon} - \log \sum_{\widetilde{Q} \subseteq \bar{Q}_{gt}} \exp\left\{ \frac{\Phi_{gt}\left(\widetilde{Q}; W^{(k)}, \theta\right)}{\sigma_\epsilon} \right\} \right]. \tag{53}$$

it estimates the log Gibbs state for each realized coalition over the partition function for each alternative coalition that could form, for each distribution network g and week t within each neighborhood, conditioning on the one-step-ahead continuation value.

The second procedure that helps the sequential optimization start is to initialize the loop with some strategic complements. I initialize this with a myopic Gibbs log-likelihood, estimating the first equation of the following:

$$\widehat{\theta}^{M} = \arg\max_{\theta} \ell_G(\theta; W \equiv 0), \qquad \widehat{\theta}^{(k+1)} = \arg\max_{\theta} \ell_G\left(\theta; W^{(k)}\right). \tag{54}$$

Combining the first equation of Equation 54 and Equation 52 allows me to initialize $(\mu^{(0)}, \boldsymbol{\sigma}^{(0)}, \eta^{(0)})$ at an initial state without assuming zero strategic complements, which pushes ρ to infinity in the first sequential estimation loop. The second equation of Equation 54 is part of the sequential pipeline. With $(\mu^{(0)}, \boldsymbol{\sigma}^{(0)}, \eta^{(0)})$, the initial one-step-ahead waiting value is constructed if each customer i waits, so that the resulting adopters $Y \subseteq R_t \setminus \{i\}$ moves the state to $\mathcal{E}_t' = T(Y, \mathcal{E}_t)$, and so, given the normalization $\sigma_\epsilon = 1$, the iteration-k one-step-ahead continuation value, following Equation 36, is

$$\begin{aligned} W_i^{(k)}(\theta, \mathcal{E}_t) = \delta \int_{Y \subseteq R_t \setminus \{i\}} \int_{\mathcal{Q}_{-i}(\mathcal{E}_t')} & \left[\gamma + v_i^0(Q'_{-i}, \mathcal{E}_t') - \log\left(1 - \sigma_i^{(k)}(Q'_{-i}, \mathcal{E}_t')\right) \right] \\ & d\eta_i^{(k)}(Q'_{-i} \mid \mathcal{E}_t')\, d\Pr{}^{(k)}(Y \mid y_{it} = 0, \mathcal{E}_t). \end{aligned} \tag{55}$$

This allows me to start sequential estimation.

9.3. Sequential Potential-Gibbs Nested Pseudo-Likelihood Estimation

Starting from $W^{(0)}$, the estimator repeats the following five-arrow loop:

$$W^{(k)} \xrightarrow{(1)} \widehat{\theta}^{(k+1)} \xrightarrow{(2)} \Phi^{(k+1)} \xrightarrow{(3)} \mu^{(k+1)} \xrightarrow{(4)} \left(\boldsymbol{\sigma}^{(k+1)}, \eta^{(k+1)}\right) \xrightarrow{(5)} W^{(k+1)} \cdots . \tag{56}$$

This sequence adapts Nested Pseudo-Likelihoods (NPL) in Aguirregabiria and Mira (2007) to my model. NPL follows an expectation-maximization sequence in which a conjectured variable is used to form expectations, structural parameters maximize an objective function conditional on those expectations, and the resulting parameters are used to update expectations in a loop. *Pseudo* refers to evaluating the likelihood at a conjectured variable rather than the true variable. I describe each step in turn:

- Step (1) holds $W^{(k)}$ fixed and maximizes the Gibbs log-likelihood in Equation 53. Within each posited θ, θ determines a coalition potential $\Phi(\theta, W^{(k)})$. Normalizing the potential by the partition function gives the Gibbs state and its likelihood, and then maximization chooses the best structural vector that maximizes the potential of observing the realized coalitions $\widehat{\theta}^{(k+1)} = \arg\max_\theta \ell_G\big(\Phi(\theta, W^{(k)})\big)$.

- Step (2) takes the fitted potential at the maximized structural parameters and saves it

$$\widehat{\Phi}^{(k+1)} = \Phi\left(\widehat{\theta}^{(k+1)}, W^{(k)}\right),$$

- Step (3) enumerates each combination of adopters in the neighborhood, and constructs Gibbs states for each one following Equation (44)

$$\mu^{(k+1)}(Q \mid \mathcal{E}_t) = \frac{\exp\{\Phi^{(k+1)}(Q \mid \mathcal{E}_t)/\sigma_\epsilon\}}{\sum_{\widetilde{Q}} \exp\{\Phi^{(k+1)}(\widetilde{Q} \mid \mathcal{E}_t)/\sigma_\epsilon\}}.$$

- Step (4) updates strategy profiles (policy functions) and the probability of observing the neighborhood coalition excluding i, following Equation 29 and Equation (45):

$$\begin{aligned} \sigma_i^{(k+1)}(B, \mathcal{E}_t) &= \frac{\mu^{(k+1)}(B \cup \{i\} \mid \mathcal{E}_t)}{\mu^{(k+1)}(B \mid \mathcal{E}_t) + \mu^{(k+1)}(B \cup \{i\} \mid \mathcal{E}_t)}, \\ \eta_i^{(k+1)}(B \mid \mathcal{E}_t) &= \mu^{(k+1)}(B \mid \mathcal{E}_t) + \mu^{(k+1)}(B \cup \{i\} \mid \mathcal{E}_t). \end{aligned} \tag{57}$$

- Step (5) constructs the next continuation value by using the expectation in Equation 55:

$$W^{(k+1)} = \mathcal{W}\left(\boldsymbol{\sigma}^{(k+1)}, \eta^{(k+1)}, \mu^{(k+1)}\right), \tag{58}$$

  This integrates over sets of adopters and neighbor coalitions and is deterministic. Then the one-step-ahead continuation value $W^{(k+1)}$ becomes the input to Step (1) in the next sequential loop.

The loop continues until a fixed point in W has been found.

9.4. Identification

Identification comes from several contrasts that hold all other variables fixed. First, suppose that two otherwise identical coalitions differ only in whether customer i joins. Then their relative probabilities identify the model's incentive for i to join. Under the Gibbs state, and via Hotz–Miller inversion, the log probability ratio of coalition formation is exactly the change in the potential of the coalition (see Equation (B.5)). Because the game has an exact potential, the change in potential equals i's net payoff from connecting, given the same observed neighbors.

Second, suppose there are two different observed neighbor coalitions Q_{-i} and Q'_{-i} that give the same externality to i by making the closest pole to i the same if i waits. As all other state variables cancel, the difference between i's incentives to join each coalition is

$$\begin{aligned}&\Delta_i\Phi_t(Q_{-i};W)-\Delta_i\Phi_t(Q'_{-i};W)\\&=-\lambda m(a_t^g;h)\left[\kappa_i(Q_{-i}\cup\{i\};\rho,\mathcal{E}_t)-\kappa_i(Q'_{-i}\cup\{i\};\rho,\mathcal{E}_t)+\mathbb{1}\{P_{it}^{\text{path}}>0\}\chi\left(\frac{1}{|Q_{-i}|+1}-\frac{1}{|Q'_{-i}|+1}\right)\right].\end{aligned} \tag{59}$$

Therefore, within-state variation in coalition size identifies the connection-charge parameters. With these two broad comparisons in mind, narrower comparisons identify the specific parameters.

Third, consider two otherwise identical customers that require two line extensions of different lengths, both past the fixed cost. The only thing that changes across these two comparisons is the connection charge, so comparing them identifies the marginal utility of money via

$$-\lambda\left[\kappa_i(Q;\rho,\mathcal{E}_t)-\kappa_i(Q';\rho,\mathcal{E}_t)\right].$$

Fourth, compare a single customer under a service drop and under one pole. The difference between them, after accounting for the pole's engineering cost, identifies the fixed cost χ. Fifth, consider the same infrastructure engineering cost, but shared by differently sized coalitions. The differences in cost identify coalition economies or costs ρ. Finally, differences in the transformer age in the initial supply placement identify the decay in supply-side funding h.

The hardest parameters to identify are λ and χ, because both make connection costs reduce adoption. They would be impossible to identify if each coalition induced only one pole at a time. Among anchored coalitions, the charge per activated pole has a coefficient of variation (sd/mean) of 1.37, so there is variation in the relative sizes of connection charges and fixed costs. Networks with variation in charge per activated pole are necessary to identify these two. Last, connecting is a terminal action, so waiting values have finite dependence. The estimation fixed point requires the payoff parameters, policy law, and constructed waiting values to all be consistent and reproduce each other, so the convergence of one-step-ahead continuation values W is essential.

An instrument is not required because of the above arguments. In the reduced form, the instrument removes reflection and simultaneity; in the structural model, these are the main objects of interest. The Gibbs coalition distribution gives the joint distribution of neighbors' adoption implied by the valuation shocks: it uses the correlation between i's and its neighbors' choices to estimate cost-sharing incentives.

From this point of view, the model is complete, and the potential therefore gives the full state evolution.

The main assumptions are first that conditional on the observed physical state and covariates, the only unobservables that affect payoffs are Type I extreme-value revision shocks.[56] The second main assumption is that the observed physical infrastructure state is sufficient to measure incentives. As it is observed physically as infrastructure network expansions, it is reasonable to assume common knowledge of the state space.

9.5. Initial Estimates

Dimensionality is still a problem, so I chose eight transformers stratified by size and start time. I then used the second to the fifth week after each transformer was first placed, censoring neighborhoods with more than eight adopters. This run included 2,534 neighborhoods, 3,104 household-week observations, and 126 adoption events. Using an annual discount factor of 0.90, it converged to a log-likelihood of -431.64. I report the coefficients for each parameter in Table 12.

Table 12 — Structural Estimates

Parameter	Estimate	Clustered SE
Marginal value of money λ (utils per 10^3 RWF)	0.00161	0.00226
Fixed Costs $\chi = \psi/\lambda$	4,874	8,093
Coordination benefits ρ	0.66	0.45
Supply side discount taper h	0.572	0.313
Benefit: constant	264.9	67.7
Benefit: log area	6.52	78.8
Benefit: log nighttime lights at t_0	-40.7	633
Benefit: log building height	-21.3	122
Benefit: elevation	0.00749	0.213
Benefit: log slope	50.9	504
Benefit: log road distance	1.36	73.3
Benefit: building characteristics missing	281.7	140.3
Benefit: nonresidential building	57.4	417
Benefit: $\mathrm{DY}_{\mathrm{invdist}}$ (grid-peer exposure)	37.1	45.6

Notes. Small Sample NPL results. Results are clustered at the neighborhood level.

The table above shows that the method can converge, though it is slow and computationally costly. However, it produces reasonable estimates, even on a small sample.

56. A future version may include latent types following Arcidiacono and Miller (2011).

10 — Discussion and Conclusion

This paper shows that marginal-cost connection pricing in infrastructure connections creates a dynamic externality. Whoever connects first pays a high price to provision the infrastructure they need, but their neighbors can then connect to that same infrastructure at a low price. The externality is double-edged. On one hand, the positive externality delays adoption and is an efficiency cost. On the other hand, high-value customers value it more, so they select into provisioning infrastructure that benefits their low-value neighbors, an equity benefit.

Evidence from 177 electricity distribution networks in Rwanda supports the proposed mechanisms. A 10% anticipated reduction in straight-line distance to the grid over the next eight weeks lowers the weekly adoption hazard by at least 0.910 percentage points, an effect equal to 2.4 times the unconditional mean weekly hazard. A contemporaneous externality of the same size raises adoption by 0.655 percentage points in OLS; the IV estimate is similar but imprecise. A placebo test of an alternative supply-side explanation for the estimated delay produced no comparable response, so the only plausible mechanism appears to be strategic behavior. The probability of joint rather than sequential adoption rises from 32% to as much as 81% as more infrastructure is shared and strategic complements increase in strength. Consistent with selection, customers with larger buildings connect earlier and pay more: relative to a mechanical zero under a flat fee, the elasticity of realized connection charges with respect to building area is 0.04 in the utility data. Similar elasticities with respect to willingness-to-pay proxies range from 0.10 to 0.24 in the World Bank surveys.

The results show a substantive policy trade-off. The selection effects are relatively small, while the delay is large but partly recouped after a neighbor provides infrastructure. Price discrimination may allow the utility to raise revenue from connections, and it may use that revenue to lower minimum connection charges. Equity effects are enhanced by proximal heterogeneity in willingness to pay, which allows high-value neighbors to cross-subsidize lower-value neighbors. Which policy is optimal depends on the government's weights on universal access, utility cost recovery, spatial efficiency, and the incidence of expansion costs relative to the size of strategic delay.

In my review of other countries' connection pricing policies, I found that other countries clearly use policies that address or mute some of the paper's strategic incentives. For example, Zambia predicts strategic incentives ex-ante, conditioning connection charges on population density, using density as an ex-ante proxy for how public or private an extension is likely to be (Zambia Electricity Regulation Board 2023). Ghana and Côte d'Ivoire use Coasian pseudo-ownership policies to internalize the externality: Consumers who pay for an extension upfront are later reimbursed when other customers connect to the same infrastructure.[57] Utilities clearly recognize the strategic incentives created by asking first movers to provision public utility infrastructure.

The paper highlights a broader tension between planning-based and market-based approaches to pricing infrastructure connections. A planning-based approach uses increasingly rich data to predict demand (Lee et al. 2021), asserting that better data will help target infrastructure more effectively toward customers who use it productively, increase the utility's revenue, or otherwise. A market-based approach instead

57. Ghana calls this a capital contribution policy, while Côte d'Ivoire calls it "droit de suite" (Compagnie Ivorienne d'Electricité 2026; Public Utilities Regulatory Commission Ghana 2024).

asks high-value customers to reveal their connection values and provision the grid to cross-subsidize their poorer neighbors by exposing customers to declining prices over time and seeing when they choose to take up electricity. It uses prices to get customers to reveal preferences rather than predict them with machine learning. As Mission 300 deploys $35 billion toward connecting 300 million people by 2030, connection policy will determines how effective financing will be by changing who pays, how long others wait, and where distribution networks expand in the last mile.

References

Ackerberg, Daniel A, and Gautam Gowrisankaran. 2006. "Quantifying equilibrium network externalities in the ACH banking industry." *The RAND journal of economics* 37 (3): 738–761. (Cited on pages 5, 86).

Aguirregabiria, Victor, Allan Collard-Wexler, and Stephen P Ryan. 2021. "Dynamic games in empirical industrial organization." 4 (1): 225–343. (Cited on pages 4, 78, 80).

Aguirregabiria, Victor, and Pedro Mira. 2007. "Sequential estimation of dynamic discrete games." *Econometrica* 75 (1): 1–53. (Cited on pages 78, 91).

Allouch, Nizar. 2015. "On the private provision of public goods on networks." *Journal of Economic Theory* 157:527–552. (Cited on page 3).

Andrews, Isaiah, James H. Stock, and Liyang Sun. 2019. "Weak Instruments in Instrumental Variables Regression: Theory and Practice." *Annual Review of Economics* 11:727–753. https://doi.org/10.1146/annurev-economics-080218-025643. (Cited on page 64).

Arcidiacono, Peter, and Robert A Miller. 2011. "Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity." *Econometrica* 79 (6): 1823–1867. (Cited on page 93).

Bala, Venkatesh, and Sanjeev Goyal. 2000. "A Noncooperative Model of Network Formation." *Econometrica* 68, no. 5 (September): 1181–1229. ISSN: 0012-9682, 1468-0262, accessed March 22, 2021. https://doi.org/10.1111/1468-0262.00155. (Cited on page 3).

Banerjee, Abhijit, Esther Duflo, and Nancy Qian. 2020. "On the Road: Access to Transportation Infrastructure and Economic Growth in China." *Journal of Development Economics* 145:102442. (Cited on page 4).

Beatty, Lauren. 2022. *How Do Natural Gas Pipeline Networks Affect Emissions From Drilling and Flaring*. Technical report. Working paper, University of Maryland. (Cited on page 5).

Bergstrom, Theodore, Lawrence Blume, and Hal Varian. 1986. "On the private provision of public goods." *Journal of public economics* 29 (1): 25–49. (Cited on page 3).

Bernard, Tanguy, and Maximo Torero. 2015. "Social Interaction Effects and Connection to Electricity: Experimental Evidence from Rural Ethiopia." *Economic Development and Cultural Change* 63 (3): 459–484. https://doi.org/10.1086/679746. (Cited on pages 4, 42, 54, 70).

Björkegren, Daniel. 2019. "The Adoption of Network Goods: Evidence from the Spread of Mobile Phones in Rwanda." *The Review of Economic Studies* 86 (3): 1033–1060. (Cited on page 5).

Blimpo, Moussa, and Malcolm Cosgrove-Davies. 2019. *Electricity Access in Sub-Saharan Africa: Uptake, Reliability, and Complementary Factors for Economic Impact.* World Bank Publications. (Cited on pages 18, 22, 104).

Blimpo, Moussa, Shaun McRae, and Jevgenijs Steinbuks. 2018. "Why Are Connection Charges so High? An Analysis of the Electricity Sector in Sub-Saharan Africa." *An Analysis of the Electricity Sector in Sub-Saharan Africa (April 16, 2018). World Bank Policy Research Working Paper,* no. 8407, (cited on pages 5, 21).

Bliss, Christopher, and Barry Nalebuff. 1984. "Dragon-slaying and ballroom dancing: The private supply of a public good." *Journal of public Economics* 25 (1-2): 1–12. (Cited on pages 3, 15, 16, 62).

Blume, Lawrence E. 1993. "The statistical mechanics of strategic interaction." *Games and economic behavior* 5 (3): 387–424. https://doi.org/10.1006/game.1993.1023. (Cited on pages 5, 50, 77, 80, 87, 109, 110).

Blume, Lawrence E, William A Brock, Steven N Durlauf, and Rajshri Jayaraman. 2015. "Linear social interactions models." *Journal of Political Economy* 123 (2): 444–496. (Cited on pages 4, 42, 53).

Borenstein, Severin. 2016. "The Economics of Fixed Cost Recovery by Utilities." *The Electricity Journal* 29, no. 7 (September): 5–12. ISSN: 1040-6190, accessed March 12, 2023. https://doi.org/10.1016/j.tej.2016.07.013. (Cited on page 19).

Bramoullé, Yann, Habiba Djebbari, and Bernard Fortin. 2009. "Identification of Peer Effects through Social Networks." *Journal of econometrics* 150 (1): 41–55. (Cited on pages 4, 54, 55, 59).

Burlig, Fiona, and Louis Preonas. 2022. "Out of the Darkness and into the Light? Development Effects of Rural Electrification," 42. (Cited on page 5).

Chakravorty, Ujjayant, Kyle Emerick, and Majah-Leah Ravago. 2016. "Lighting up the Last Mile: The Benefits and Costs of Extending Electricity to the Rural Poor." *Resources for the Future Discussion Paper,* 16–22. (Cited on pages 5, 18).

Dabalen, Andrew, Justice Tei Mensah, Aimable Nsabimana, and Alexandre Nshunguyinka. 2025. "Agricultural Windfalls and Electrification," (cited on page 25).

Dahl, Gordon B., Katrine V. Løken, and Magne Mogstad. 2014. "Peer Effects in Program Participation." *American Economic Review* 104, no. 7 (July): 2049–2074. ISSN: 0002-8282, accessed May 4, 2022. https://doi.org/10.1257/aer.104.7.2049. (Cited on page 14).

Dinkelman, Taryn. 2011. "The Effects of Rural Electrification on Employment: New Evidence from South Africa." *American Economic Review* 101, no. 7 (December): 3078–3108. ISSN: 0002-8282, accessed March 14, 2022. https://doi.org/10.1257/aer.101.7.3078. (Cited on page 5).

dos Santos, Carolina Tojal R, and Bruna Morais Guidetti. 2024. *Expansion of Piped Water and Sewer Networks: The Effects of Regulation.* Working Paper. (Cited on page 5).

Elliott, Matthew, and Benjamin Golub. 2019. "A Network Approach to Public Goods." *Journal of Political Economy* 127 (2): 730–776. https://doi.org/10.1086/701032. eprint: https://doi.org/10.1086/701032. (Cited on pages 3, 4, 47, 78, 90).

Elvidge, Christopher D, Mikhail Zhizhin, Tilottama Ghosh, Feng-Chi Hsu, and Jay Taneja. 2021. "Annual time series of global VIIRS nighttime lights derived from monthly averages: 2012 to 2019." *Remote Sensing* 13 (5): 922. (Cited on page 28).

Faber, Benjamin. 2014. "Trade Integration, Market Size, and Industrialization: Evidence from China's National Trunk Highway System." *Review of Economic Studies* 81 (3): 1046–1070. (Cited on pages 4, 39).

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. 2020. "Bartik instruments: What, when, why, and how." *American Economic Review* 110 (8): 2586–2624. (Cited on page 4).

Golumbeanu, Raluca, and Douglas F Barnes. 2013. "Connection Charges and Electricity Access in Sub-Saharan Africa." *World Bank Policy Research Working Paper,* no. 6511, https://doi.org/10.1596/1813-9450-6511. (Cited on pages 8, 21).

Griffith, Alan. 2022. "Name your friends, but only five? the importance of censoring in peer effects estimates using social network data." *Journal of Labor Economics* 40 (4): 779–805. (Cited on page 59).

Grogan, Louise, and Asha Sadanand. 2013. "Rural Electrification and Employment in Poor Countries: Evidence from Nicaragua." *World Development* 43:252–265. https://doi.org/10.1016/j.worlddev.2012.09.002. (Cited on page 5).

Hagberg, Aric, Pieter J Swart, and Daniel A Schult. 2007. *Exploring network structure, dynamics, and function using NetworkX.* Technical report. Los Alamos National Laboratory (LANL). (Cited on page 29).

Haller, Hans. 2016. "Networks as Public Infrastructure: Externalities, Efficiency, and Implementation: Networks as Public Infrastructure." *Journal of Public Economic Theory* 18, no. 2 (April): 193–211. ISSN: 10973923, accessed November 2, 2022. https://doi.org/10.1111/jpet.12139. (Cited on page 3).

Hart, Sergiu, and Andreu Mas-Colell. 1989. "Potential, value, and consistency." *Econometrica: Journal of the Econometric Society,* 589–614. (Cited on pages 3, 10, 85).

———. 1996. "Bargaining and value." *Econometrica: Journal of the Econometric Society,* 357–380. (Cited on page 10).

Hotz, V Joseph, and Robert A Miller. 1993. "Conditional choice probabilities and the estimation of dynamic models." *The Review of Economic Studies* 60 (3): 497–529. (Cited on page 78).

Jackson, Matthew O, and Asher Wolinsky. 1996. "A strategic model of social and economic networks." *Journal of economic theory* 71 (1): 44–74. (Cited on page 3).

Jeuland, Marc, T. Robert Fetter, Yating Li, Subhrendu K. Pattanayak, Faraz Usmani, Randall A. Bluffstone, Carlos Chávez, et al. 2021. "Is Energy the Golden Thread? A Systematic Review of the Impacts of Modern and Traditional Energy Use in Low- and Middle-Income Countries." *Renewable and Sustainable Energy Reviews* 135 (January): 110406. ISSN: 13640321, accessed February 16, 2022. https://doi.org/10.1016/j.rser.2020.110406. (Cited on page 5).

Jia, Panle. 2008. "What happens when Wal-Mart comes to town: An empirical analysis of the discount retailing industry." *Econometrica* 76 (6): 1263–1316. (Cited on pages 5, 86).

Johnston, Sarah, Yifei Liu, and Chenyu Yang. 2023. "An empirical analysis of the interconnection queue," (cited on page 5).

Joskow, Paul L. 2007. "Regulation of natural monopoly." *Handbook of law and economics* 2:1227–1348. (Cited on page 19).

Khandker, Shahidur R, Douglas F Barnes, and Hussain A Samad. 2009. "Welfare Impacts of Rural Electrification: A Case Study from Bangladesh." *World Bank Policy Research Working Paper,* no. 4859, (cited on page 5).

———. 2013. "Welfare Impacts of Rural Electrification: A Panel Data Analysis from Vietnam." *Economic Development and Cultural Change* 61, no. 3 (April): 659–692. ISSN: 0013-0079, 1539-2988, accessed March 14, 2022. https://doi.org/10.1086/669262. (Cited on page 5).

Kuersteiner, Guido M, and Ingmar R Prucha. 2020. "Dynamic spatial panel models: Networks, common shocks, and sequential exogeneity." *Econometrica* 88 (5): 2109–2146. (Cited on page 54).

Lee, Kenneth, Eric Brewer, Carson Christiano, Francis Meyo, Edward Miguel, Matthew Podolsky, Javier Rosa, and Catherine Wolfram. 2016. "Electrification for "Under Grid" Households in Rural Kenya." *Development Engineering* 1 (June): 26–35. ISSN: 23527285, accessed May 9, 2022. https://doi.org/10.1016/j.deveng.2015.12.001. (Cited on page 18).

Lee, Kenneth, Edward Miguel, and Catherine Wolfram. 2020a. "Does Household Electrification Supercharge Economic Development?" *Journal of Economic Perspectives* 34, no. 1 (February): 122–144. ISSN: 0895-3309, accessed March 14, 2022. https://doi.org/10.1257/jep.34.1.122. (Cited on page 5).

———. 2020b. "Experimental Evidence on the Economics of Rural Electrification." *Journal of Political Economy,* https://doi.org/10.1086/705417. (Cited on pages 3, 5, 22).

Lee, Stephen J, Dhruv Suri, Priyanshi Somani, Christopher L Dean, Jason Pacheco, Robert Stoner, I Perez-Arriaga, JW Fisher III, and Jay Taneja. 2021. "How probabilistic electricity demand forecasts can expedite universal access to clean and reliable electricity." *Energy for Economic Growth,* (cited on pages 23, 94).

Lipscomb, Molly, A. Mushfiq Mobarak, and Tania Barham. 2013. "Development Effects of Electrification: Evidence from the Topographic Placement of Hydropower Plants in Brazil." *American Economic Journal: Applied Economics* 5, no. 2 (January): 200–231. ISSN: 1945-7782, 1945-7790, accessed March 14, 2022. https://doi.org/10.1257/app.5.2.200. (Cited on pages 4, 5, 39).

Littlechild, Stephen C, and GF Thompson. 1977. "Aircraft landing fees: a game theory approach." *The Bell Journal of Economics,* 186–204. (Cited on pages 41, 81).

Littlechild, Stephen C., and Guillermo Owen. 1973. "A simple expression for the Shapley value in a special case." *Management Science* 20 (3): 370–372. (Cited on pages 3, 41, 81).

Manski, Charles F. 1993. "Identification of Endogenous Social Effects: The Reflection Problem." *The Review of Economic Studies* 60 (3): 531–542. ISSN: 0034-6527, accessed March 4, 2023. https://doi.org/10.2307/2298123. JSTOR: 2298123. (Cited on pages 4, 45).

Márkus, Judit, Péter Miklós Pintér, and Anna Radványi. 2012. *The Shapley value for airport and irrigation games.* Technical report. IEHAS Discussion Papers. (Cited on pages 41, 81).

Masselus, Lise, Jörg Ankel-Peters, Gabriel Gonzalez Sutil, Vijay Modi, Joel Mugyenyi, Anicet Munyehirwe, Nathan Williams, and Maximiliane Sievert. 2024. *10 years after: Long-term adoption of electricity in rural Rwanda.* 1086. Ruhr Economic Papers. (Cited on page 5).

Mele, Angelo. 2017. "A structural model of dense network formation." *Econometrica* 85 (3): 825–850. (Cited on page 5).

Monderer, Dov, and Lloyd S Shapley. 1996. "Potential games." *Games and economic behavior* 14 (1): 124–143. (Cited on pages 5, 77, 84, 85).

OpenStreetMap Contributors. 2026. *OpenStreetMap.* Available online: https://download.geofabrik.de/africa/rwanda.html (accessed May 2026). (Cited on page 28).

Parshall, Lily, Dana Pillai, Shashank Mohan, Aly Sanoh, and Vijay Modi. 2009. "National electricity planning in settings with low pre-existing grid coverage: Development of a spatial model and case study of Kenya." *Energy Policy* 37 (6): 2395–2410. (Cited on page 23).

Rosenthal, Robert W. 1973. "A class of games possessing pure-strategy Nash equilibria." *International Journal of Game Theory* 2 (1): 65–67. (Cited on page 84).

Rossi, Jim. 2000. "Universal Service in Competitive Retail Electric Power Markets: Whither the Duty to Serve?" *Energy LJ* 21:27. (Cited on page 21).

Rud, Juan Pablo. 2012. "Electricity Provision and Industrial Development: Evidence from India." *Journal of Development Economics* 97, no. 2 (March): 352–367. ISSN: 03043878, accessed March 14, 2022. https://doi.org/10.1016/j.jdeveco.2011.06.010. (Cited on page 5).

Rwanda Energy Group (REG). 2017. *New connection Policy.* July 15, 2017. Accessed December 1, 2022. https://www.reg.rw/fileadmin/user_upload/RevisedNewConnectionPolicy__1_-75742.pdf. (Cited on pages 25, 26, 103).

Sadler, Evan. 2014. "Bounded Rationality and the Network Economist's Dilemma." *Available at SSRN 2501690,* (cited on page 48).

Sanderson, Eleanor, and Frank Windmeijer. 2016. "A weak instrument F-test in linear IV models with multiple endogenous variables." *Journal of econometrics* 190 (2): 212–221. (Cited on pages 59, 64).

Shapley, Lloyd S. 1953. "A value for n-person games," (cited on pages 3, 10, 80).

Simon, Herbert A. 1955. "A behavioral model of rational choice." *The quarterly journal of economics,* 99–118. (Cited on page 48).

Sirko, William, Emmanuel A Brempong, Juan T C Marcos, Akosua Annkah, Aida Korme, Mohamed A Hassen, Kishor Sapkota, Tov Shekel, Abdoulaye Diack, Shlomo Nevo, et al. 2023. "High-Resolution Building and Road Detection from Sentinel-2." *arXiv preprint arXiv:2310.11622,* (cited on page 28).

Trimble, Chris, Masami Kojima, Ines Perez Arroyo, and Farangis Mohammadzadeh. 2016. *Financial Viability of Electricity Sectors in Sub-Saharan Africa: Quasi-Fiscal Deficits and Hidden Costs.* Policy Research Working Paper 7788. Washington, DC: World Bank. https://openknowledge.worldbank.org/handle/10986/24869. (Cited on page 21).

UNDP. 2023. *UNDP Rwanda Data Hub.* Https://geohub.data.undp.org/data/00d5add9be37e465398b081683c3ec03 (accessed December 20, 2022). (Cited on page 28).

van de Walle, Dominique, Martin Ravallion, Vibhuti Mendiratta, and Gayatri Koolwal. 2013. *Long-Term Impacts of Household Electrification in Rural India.* Policy Research Working Papers. The World Bank, June. Accessed March 14, 2022. https://doi.org/10.1596/1813-9450-6527. (Cited on page 5).

Van de Walle, Dominique, Martin Ravallion, Vibhuti Mendiratta, and Gayatri Koolwal. 2017. "Long-Term Gains from Electrification in Rural India." *The World Bank Economic Review* 31 (2): 385–411. (Cited on page 5).

Virtanen, Pauli, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, et al. 2020. "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python." *Nature Methods* 17:261–272. https://doi.org/10.1038/s41592-019-0686-2. (Cited on page 28).

Walter, Torsten Figueiredo, and Niclas Moneke. 2024. "When Does Electrification Work? Evidence from Sub-Saharan Africa." *Working Paper,* https://niclasmoneke.com/wp-content/uploads/FigueiredoWalter_Moneke-When_Does_Electrification_Work-2024-full.pdf. (Cited on page 5).

Warr, Peter G. 1983. "The private provision of a public good is independent of the distribution of income." *Economics letters* 13 (2-3): 207–211. (Cited on page 3).

Zvoleff, Alex, Ayse Selin Kocaman, Woonghee Tim Huh, and Vijay Modi. 2009. "The impact of geography on energy infrastructure costs." *Energy Policy* 37 (10): 4066–4078. (Cited on page 19).

Utility and Policy Sources

Compagnie Ivorienne d'Electricité. 2026. "Droit de suite." Accessed June 21, 2026. https://www.cie.ci/particuliers/droit-de-suite. (Cited on page 94).

Ministry of Infrastructure (MININFRA) Rwanda. 2015. *Energy Sector Strategic Plan.* Accessed July 3, 2026. https://rise.esmap.org/sites/default/files/library/rwanda/Electricity%20Access/Rwanda_Energy%20sector%20strategic%20plan_2015.pdf. (Cited on pages 103, 104).

———. 2019. *Design of the National Electrification Plan in Rwanda: NESP Task-2 Final Report.* Accessed July 3, 2026. https://rise.esmap.org/sites/default/files/library/rwanda/Electricity%20Access/Rwanda_Design%20of%20the%20National%20Electrification%20Plan_2019.pdf. (Cited on page 104).

Ministry of Infrastructure (MININFRA) Rwanda, and ESRI. 2024. *Rwanda Ministry of Infrastructure Geoportal.* Accessed December 2, 2024. https://mininfra-geoinfo.maps.arcgis.com/home/gallery.html. (Cited on page 27).

Public Utilities Regulatory Commission Ghana. 2024. "PURC Capital Contribution Guidelines," May. Accessed June 21, 2026. https://www.purc.com.gh/attachment/156321-20240423030423.pdf. (Cited on page 94).

Rwanda Energy Group (REG). 2020. *Distribution Reticulation Standards, Version VII.* Accessed June 14, 2026. https://www.reg.rw/fileadmin/user_upload/Rev7_Distribution_Reticulation_Standards_Version_VII_FINAL_1_and_signed.pdf. (Cited on pages 25, 104, 105).

———. 2024a. *Rwanda Electricity Distribution Master Plan.* Accessed June 17, 2026. https://www.reg.rw/fileadmin/user_upload/Rwanda_Electricity_Distribution_Master_Plan_June_2024_Version__13_08_2024_clean_signed_gn.pdf. (Cited on page 104).

———. 2024b. *Rwanda Energy Group Annual Report 2023–24.* Accessed June 17, 2026. https://www.reg.rw/fileadmin/user_upload/REG_Annual_Report_2023-24FY.pdf. (Cited on pages 104, 105).

Rwanda Energy Group (REG), and ESRI. 2022. *Rwanda Survey Up: Transformers, LV Support, LV Lines, Customers, Customer Map.* Accessed May 17, 2022. https://euclgis.reg.rw/server/rest/services/Rwanda_Survey_up/MapServer. (Cited on page 27).

Woods, Todd. 2019. *Final Report – Volume 3a: New Customer Connection Policy and Connection Cost Methodology - Guidelines from the Electricity Regulatory Authority.* Technical report. Accessed September 1, 2025. https://www.era.go.ug/download/final-report-volume-3a-new-customer-connection-policy-and-connection-cost-methodology-guidelines-from-the-electricity-regulatory-authority/. (Cited on pages 23, 25).

World Bank. 2015. *Project Appraisal Document: Rwanda Electricity Sector Strengthening Project (P150634).* November 19, 2015. Accessed June 22, 2026. https://documents.worldbank.org/curated/en/426911468179342511/pdf/PAD1301-P150634-IDA-R2015-0295-1-Box393256B-OUO-9.pdf. (Cited on pages 104, 105).

———. 2018a. *Multi-Tier Framework for Measuring Energy Access: Rwanda.* Accessed May 31, 2026. https://energydata.info/dataset/rwanda-multi-tier-framework-mtf-survey-2018. (Cited on pages 29, 72, 74, 104).

———. 2018b. *Rwanda Electricity Access Scale-up and Sector Wide Approach (SWAp) Development Project (P111567).* September 30, 2018. Accessed June 22, 2026. https://documents.worldbank.org/curated/en/790591539716903028/pdf/Implementation-Completion-and-Results-Report-ICR-Document-10102018.pdf. (Cited on pages 104, 105).

———. 2020. *Rwanda: Energy Access and Quality Improvement Project (P172594).* August 25, 2020. Accessed June 22, 2026. https://documents.worldbank.org/curated/en/819241600653622828/pdf/Rwanda-Energy-Access-and-Quality-Improvement-Project.pdf. (Cited on pages 104, 105).

———. 2022. *Rwanda — Multi-Tier Framework Energy Survey.* Accessed August 16, 2026. https://microdata.worldbank.org/index.php/catalog/6429. (Cited on pages 29, 72, 74, 104).

———. 2023. *Implementation Completion and Results Report: Rwanda Electricity Sector Strengthening Project (P150634).* April 20, 2023. Accessed June 22, 2026. https://documents.worldbank.org/curated/en/099042423180540738/pdf/BOSIB0256b108a0e80aca401a846b141890.pdf. (Cited on pages 104, 105).

World Bank Independent Evaluation Group. 2019. *ICR Review: Rwanda Electricity Access Scale-up and Sector-Wide Approach Development Project (P111567).* March 20, 2019. Accessed June 22, 2026. https://documents.worldbank.org/curated/en/102241553114033231/pdf/Rwanda-Rwanda-Electricity-Access-Scale-up-Proj.pdf. (Cited on pages 35, 105).

Zambia Electricity Regulation Board. 2023. "Connection Charges Consultation Paper." Accessed June 21, 2026. https://www.erb.org.zm/wp-content/uploads/ConnectionChargesConsultationPaper.pdf.pdf. (Cited on pages 23, 94).

A — Rwanda's Connection Costs and Charges

I report sources and figures each of the costs used in the paper in Table A1 and then discuss them.

A.1. Connection fees 56,000RWF

The online schedule charges a flat 56,000 RWF connection fee, and it is documented in the connection policy (Rwanda Energy Group (REG) 2017). This number is also reported in the EARP Connection Policy schedule reproduced in the 2015 sector strategy (Ministry of Infrastructure (MININFRA) Rwanda 2015, Annex Table 24).

Table A1 — Documentary evidence on connection costs, charges, and subsidies

Object	Value	Basis	Source	Location
Connection charges				
Standard connection fee	56,000 RWF	posted	Ministry of Infrastructure (MININFRA) Rwanda (2015)	Annex T.24
Per-pole charge	US$200	discussion, bunching	Blimpo and Cosgrove-Davies (2019) and World Bank (2018a)	microdata
Installed cost per km				
LV (ABC 3×35 / 3×70)	US$15,000 / 17,931	planning catalogue	Ministry of Infrastructure (MININFRA) Rwanda (2019)	T.13
LV rehabilitation	∼US$50,000	project tables	Rwanda Energy Group (REG) (2024a)	T.54, p. 88
LV in mixed extensions	US$17–21k		Rwanda Energy Group (REG) (2024a)	T.52–53
Cost per connection				
EARP actual	US$411	completion	World Bank (2018b)	p.56, p. 21
ESSP actual	∼US$533	completion	World Bank (2023)	p. 43
EAQIP appraisal	US$603–758	planned	World Bank (2020)	p.41, p. 35
FY2023/24 implied	∼US$857	derived	Rwanda Energy Group (REG) (2024b)	outputs
Implied subsidy				
Lower bound subsidy	∼60–66% of total	derived	—	this appendix
Customer share of total	∼1/10	stated	World Bank (2015)	p.22, p. 9

Note. 1,000 RWF ≈ US$1 over the sample period. Per-km-to-per-pole conversions use 20.4 poles per km (median observed span 49 m over 239,829 real LV spans, consistent with the codified 44–55 m standard, Rwanda Energy Group (REG) 2020, Table 18). "Installed" costs include materials, labor, transport, design, margins.

A.2. Pole and Line connection charges

The paper assumes customers are charged US$2 per meter of line beyond a 37 m service drop and US$200 per pole. The $200 per pole figure was shared with me by Rwanda's Ministry of Infrastructure in 2022. However, I did not find this exact number in utility or Ministry of Infrastructure documents. I nevertheless find evidence broadly consistent with this number in the following sources:

- Figure 3.5b in Blimpo and Cosgrove-Davies (2019), which indicates USD 189 for poles and wiring in Rwanda alongside other distance-based connection charges; and
- a World Bank survey in which consumers reported the price they paid to connect to the grid, indicating bunching at USD 50 intervals as well as shared connections (World Bank 2018a, 2022);

In the World Bank Multi-Tier Framework survey, household-reported connection payments bunch at US$100, US$150, and US$200, round splits of a US$200-per-pole charge among coalition members (World Bank 2018a, 2022).

A.3. Pole and Line Supply Costs and Subsidies

The total *installed* cost of LV extensions is US$15,000–17,931 per km of LV aerial bundled cable in Rwanda's National Electrification Plan's network catalogue (Ministry of Infrastructure (MININFRA) Rwanda 2019, Tables 13–15). The same document shows the supply cost of an MV/LV transformers is between US$3,600 (15 kVA) to US$28,455 (400 kVA). More recent records from the Electricity Distribution Master Plan indicate a total cost of $17,000–21,000 per km of LV (Rwanda Energy Group (REG) 2024a, Tables 52–54). A World Bank document on the Electricity Sector Strengthening Project reports 1,295.4 km of LV and 849.8 km of MV built,at a blended cost of roughly US$20,750 per km (World Bank 2023, results framework, p. 43). At

the observed 44-meter average and 55-meter reticulation standard, these crude estimates correspond to a raw supply costs of $660 - -1155$ per pole-span (Rwanda Energy Group (REG) 2020, Table 18).

Conservatively assuming the lower end of material costs, $200-per-pole and $2-per-meter connection charge is approximately a 30% subsidy on materials, perhaps in addition to a 100% subsidy on unobserved labor, service, and transport costs. This 60–66% subsidy or one-third pass-through, is primarily financed by multilateral development banks and other donors (World Bank 2023, 2020). This rate is conservative and perhaps reflective of the time period of the study; earlier World Bank records indicated that customers contributed only about one-tenth of the total connection cost (World Bank 2015). I assume a two-thirds subsidy rate, equivalently a one-third pass-through of supply-side costs throughout the study. I document additional sources behind the cost figures in section 3 and the subsidy calibration in the structural model below. Within the sample period, 1,000 RWF $\approx$ US$1.

A.4. Average Costs per connection

Per-connection totals are the best-documented object because they show up in World Bank reports. The EARP completion report puts the realized average cost per connection at US$411 against an appraisal baseline of US$1,000 (World Bank 2018b; World Bank Independent Evaluation Group 2019, p.56, p. 21). A more recent project appraised the average supply cost of a connection at US$603–758 per connection (World Bank 2020, p.41, p. 35). The Electricity Sector Strengthening Project realized roughly US$533 per connection (World Bank 2023), and its appraisal stated that the customer's US$100 contribution covered about one tenth of the total (World Bank 2015, p.22, p. 9): it appears that different policies were used prior to 2017 that may not have been distance based. Utility financial reports show that in fiscal years 2023/2024, supply costs are about US$857 per connection when solving for them from the reported US$50,000 per km cost, 2,357.7 km of LV, and 137,502 connections added in that year (Rwanda Energy Group (REG) 2024b).

B — Theoretical Proofs

B.1. Derivation of the Shapley Cost Share

Specifically, the Shapley cost shares average each customer's marginal contribution across possible coalitions. Given surplus functions,

$$\begin{aligned} w(\emptyset) &= 0, \\ w(\{i\}) &= V_i - c_i, \\ w(\{j\}) &= V_j - c_j, \\ w(\{i,j\}) &= V_i + V_j - (1+\rho)C, \end{aligned}$$

Shapley payoffs to customer i are

$$\begin{aligned}\phi_i &= \frac{1}{2}w(\{i\}) + \frac{1}{2}\left[w(\{i,j\}) - w(\{j\})\right] \\ &= \frac{1}{2}(V_i - c_i) + \frac{1}{2}\left[(V_i + V_j - (1+\rho)C) - (V_j - c_j)\right] \\ &= V_i - \frac{1}{2}\left[(1+\rho)C + c_i - c_j\right],\end{aligned}$$

and substituting into the cooperation payoff gives Shapley cost shares κ_i.

B.2. Proof of Lemma 1

The strategy profile *(Wait$_i$, Adopt$_j$)* is a pure-strategy Nash equilibrium if and only if

$$V_j - c_j \geq 0 \qquad \text{and} \qquad \beta(V_i - c_i') \geq V_i - \kappa_i.$$

Customer j is willing to pay the unilateral connection charge, and customer i prefers connecting second to adopting jointly.

As i's net surplus is assumed to be higher $V_i - c_i \geq V_j - c_j$, both customers are willing to pay the unilateral connection charge $V_i - c_i \geq V_j - c_j \geq 0$. As the Shapley cost shares from Equation 2 will imply that $\kappa_i - \kappa_j = c_i - c_j$, $(V_i - \kappa_i) - (V_j - \kappa_j) = (V_i - c_i) - (V_j - c_j)$, so the j's free-riding condition implies

$$\beta(V_j - c_j') - (V_j - \kappa_j) \geq \beta(V_i - c_i') - (V_i - \kappa_i) \qquad \geq 0 \tag{B.1}$$

$$\beta(V_j - c_j') \geq V_j - \kappa_j, \tag{B.2}$$

So j prefers to connect second when i adopts. As i is willing to adopt unilaterally and j prefers to wait, *(Adopt$_i$, Wait$_j$)* is also a pure-strategy Nash equilibrium. Therefore, if i has a greater net unilateral surplus $V_i - c_i \geq V_j - c_j$, the parameter region supporting *(Wait$_i$, Adopt$_j$)* is contained in the region supporting *(Adopt$_i$, Wait$_j$)*, so *(Wait$_i$, Adopt$_j$)* will never be a unique Nash equilibrium. □

B.3. Proof of Lemma 2

For any $\sigma_j \in [0,1]$, the net gain from adoption is strictly increasing in the customer's connection value $\frac{\partial \Delta_i}{\partial V_i} = \sigma_j(1-\beta) + (1-\sigma_j)\left[1 - \beta W'(V_i)\right] > 0$. This implies a cutoff rule where the customer is indifferent to adopting or not when $\Delta_i(V^*(\sigma_j);\sigma_j) = 0$. The neighbor's adoption probability σ_j determines this cutoff. The cutoff need not lie in $[\underline{V}, \overline{V}]$, so there are corner cases in which every type adopts or every type waits. The net gain $\Delta_i(V_i;\sigma_j)$ is continuously differentiable in both parameters and strictly increasing in V_i, so the Implicit Function Theorem gives Equation 8. As just indicated, the denominator is positive. Under strategic substitutes, the numerator is negative $\frac{\partial \Delta_i}{\partial \sigma_j} = \frac{\partial^2 \pi_i}{\partial \sigma_i \partial \sigma_j} < 0$, so the overall term is positive and the cutoff increases with σ_j. Under strategic complements the cutoff decreases with σ_j. □

B.4. *Proof of Proposition 2*

I first establish existence. Without loss of generality, suppose that neighbor j follows a cutoff rule with threshold v, adopting when $V_j \geq v$. Because V_j is an independent draw from the common distribution F, customer i assigns probability $\sigma_j = \Pr[s_j(V_j) = \text{Adopt}] = 1 - F(v)$ to the event that j adopts. Let $\widetilde{V}^*(\sigma_j) \in [\underline{V}, \overline{V}]$ denote customer i's best-response cutoff from Lemma 2 as a function of the neighbor's mixed strategy. The best-response cutoff as a function of the neighbor's cutoff is therefore $\text{BR}(v) = \widetilde{V}^*(1 - F(v))$. Because F and $\widetilde{V}^*(\sigma_j)$ are continuous, BR is a continuous map from $[\underline{V}, \overline{V}]$ into itself. Let $g(v) = \text{BR}(v) - v$. Since $\text{BR}(v) \in [\underline{V}, \overline{V}]$, $g(\underline{V}) \geq 0$ and $g(\overline{V}) \leq 0$. The intermediate value theorem therefore gives some $V^* \in [\underline{V}, \overline{V}]$ such that $g(V^*) = 0$, and $V^* = \text{BR}(V^*)$. As the cutoff rule with threshold V^* is a best response to itself, it is a symmetric cutoff equilibrium.

I now show uniqueness under strategic substitutes. By Lemma 2, strategic substitutes imply that $\widetilde{V}^*(\sigma_j)$ is weakly increasing in σ_j. Because $1 - F(v)$ is weakly decreasing in v, $\text{BR}(v) = \widetilde{V}^*(1 - F(v))$ is weakly decreasing in v. Hence, for any $v_2 > v_1$, $g(v_2) - g(v_1) = \text{BR}(v_2) - \text{BR}(v_1) - (v_2 - v_1) < 0$. Therefore, g is strictly decreasing and can have at most one zero. As existence shows that it has at least one zero, there is exactly one fixed point and so the symmetric cutoff equilibrium is unique under strategic substitutes. Under strategic complements, BR may not necessarily be decreasing, and so there may be multiple fixed points and equilibria.

Finally, consider behavior in any symmetric cutoff equilibrium with threshold V^*. By Lemma 2, a player adopts if $V_k > V^*$, waits if $V_k < V^*$, and is indifferent if $V_k = V^*$. Consequently, whenever two players' values lie on opposite sides of the cutoff, only the higher-value player adopts first. Further, if a type adopts, then every higher type also adopts, so the adoption decision is weakly increasing in connection value. High-connection-value players therefore select into earlier adoption at the first-mover charge, while low-connection-value players select into waiting for the second-mover charge. □

B.5. *Proof of Proposition 3*

Holding $W(V)$ fixed, a waiting customer can choose between continuing the baseline war of attrition or holding out until T^k. The marginal payoff from increasing their own adoption probability becomes

$$\Delta_i^{T^k}(V; \sigma_j) = \sigma_j \big[(V - \kappa) - \beta(V - c')\big] + (1 - \sigma_j) \left[(V - c) - \max\left\{\beta W(V),\ \beta^{T^k - t}(V - c^k)\right\}\right]. \quad \text{(B.3)}$$

The unilateral net benefit of adoption (the second term) is the minimum of $V - c - \beta W(V)$ and $V - c - \beta^{T^k - t}(V - c^k)$. Both terms are strictly increasing in V because $\beta W'(V) < 1$ and $\beta^{T^k - t} < 1$. The minimum is continuous and strictly increasing, including at the point where the maximizing continuation payoff switches branches, but it will not necessarily be differentiable there. Therefore, $\Delta_i^{T^k}$ is continuous and strictly increasing in V, and the best response in the foreseeable-event equilibrium remains a cutoff rule, as in Lemma 2.

1. *Baseline Adoption Unaffected.* When the baseline continuation value is weakly preferable to waiting for the foreseeable event at the baseline cutoff, $\beta W(V^*) \geq \beta^{T^k - t}(V^* - c^k)$, the foreseeable option value does not change the payoff. Therefore, the net gains across the wars are the same, $\Delta_i^{T^k}(V^*; \sigma^*) =$

$\Delta_i(V^*;\sigma^*) = 0$, where the last equality comes from indifference, so the baseline symmetric cutoff equilibrium V^* is unaffected. Uniqueness from Proposition 2 then implies $V^*_{t,T^k} = V^*$.

2. *Delay into the externality.* The baseline cutoff strategy is no longer optimal when the foreseeable network extension raises the payoff from waiting $\Delta_i^{T^k}(V^*;\sigma^*) < 0$. As $\Delta_i^{T^k}$ is strictly increasing in V, the cutoff best response to the baseline equilibrium strategy is above V^*. Let $r \in \{0, T^k\}$ correspond to the baseline and foreseeable-event wars, respectively, so that $\mathrm{BR}_0(v)$ and $\mathrm{BR}_{T^k}(v)$ denote the cutoff best responses to a neighbor cutoff, and $g_r(v) = \mathrm{BR}_r(v) - v$. As in Proposition 2, strategic substitutes imply that each g_r is strictly decreasing and has a unique zero, so each best-response function has a unique fixed point. The higher payoff from waiting implies a higher cutoff strategy $\mathrm{BR}_{T^k}(v) \geq \mathrm{BR}_0(v)$, so the fixed point of the foreseeable-event war is farther to the right and implies a higher cutoff, $g_{T^k}(V^*) \geq g_0(V^*) = 0 \implies V^*_{t,T^k} \geq V^*$. As F is increasing, foreseeable event adoption probabilities are lower than baseline, so $\sigma_{t,T^k} = 1 - F(V^*_{t,T^k}) \leq 1 - F(V^*) = \sigma^*$.

3. *Frozen adoption hazards (optional phase).* If the other customer waits, holding out for customer k weakly dominates unilateral adoption because $\beta^{T^k - t}(V - c^k) \geq V - c$ in the maximum of Equation B.3. If the other customer adopts, $\beta(V - c') > V - \kappa$ makes free-riding preferable to joint adoption. As waiting dominates adoption, both players wait and current adoption is zero.

4. *Joint adoption with k at time T^k.* At the event, customer i can either join k's connection request and receive the contemporaneous payoff $V_i - \kappa^k$, or wait one more period and receive the discounted post-event payoff $\beta(V_i - c^k)$. The benefit of joining is $(V_i - \kappa^k) - \beta(V_i - c^k) = (1-\beta)V_i - \kappa^k + \beta c^k$, which is strictly increasing in V_i. Therefore, customer i joins k if and only if

$$V_i \geq V^{\text{joint}}_{T^k} \equiv \frac{\kappa^k - \beta c^k}{1 - \beta}.$$

   Only customers who have not yet adopted at T^k can join. Let $V^{\text{surv}}_{T^k} \equiv \min_{s<T^k} V^*_{s,T^k}$ denote the lowest cutoff faced before the event. A customer reaches T^k if and only if $V_i < V^{\text{surv}}_{T^k}$. The time T^k joint-adoption probability among the remaining customers is consequently $\left(F(V^{\text{surv}}_{T^k}) - F(V^{\text{joint}}_{T^k})\right)/F(V^{\text{surv}}_{T^k})$ when $V^{\text{joint}}_{T^k} < V^{\text{surv}}_{T^k}$, and zero otherwise.

5. *Realization of free-riding after time T^k.* Consider a customer who reached time T^k because waiting for k's foreseeable extension providing externality $c - c^k \geq 0$, but who did not jointly adopt with k. This implies $V_i - \kappa^k < \beta(V_i - c^k)$, or equivalently $V_i < V^{\text{joint}}_{T^k}$. Their option value of waiting for the event was realized. They now observe the post-event charge c^k and the continuation value $W^k(V_i)$ of the war that resumes after the event, and adopt immediately if $V_i - c^k \geq \beta W^k(V_i)$. Customers who harvest and delay adoption in order to receive neighbor externalities immediately and sequentially adopt after the event.

6. *A higher baseline after the event.* If neither i nor j joined k's coalition nor did they free-ride, they still observe a lower connection charge c^k. Holding $W(V)$, c', and κ fixed, replace the unilateral charge c in Equation 7 by c^k and denote the resulting marginal payoff from own adoption by $\Delta_i^{T^k+1}$. Then for

every V and σ_j,

$$\Delta_i^{T^k+1}(V;\sigma_j) - \Delta_i(V;\sigma_j) = (1-\sigma_j)(c-c^k) \geq 0,$$

so the post-event cutoff best response is weakly below the baseline best response. As above, let BR_k be the post-event best response and let $g_k(v) = \mathrm{BR}_k(v) - v$. Under strategic substitutes, g_k is strictly decreasing and has a unique zero. At the baseline cutoff V^*, $g_k(V^*)$ is nonpositive:

$$g_k(V^*) = \mathrm{BR}_k(V^*) - V^* \leq \mathrm{BR}_0(V^*) - V^* = 0,$$

so $V^{*,k} \leq V^*$. As F is increasing, baseline adoption hazards are higher after the event. $\sigma^{*,k} = 1 - F(V^{*,k}) \geq 1 - F(V^*) = \sigma^*$. In the resumed war, the remaining customers adopt exactly when $V_i \geq V^{*,k}$. Therefore, the sequential adopters of the previous phase are those with $V^{*,k} \leq V_i < V_{T^k}^{\text{joint}}$.

□

B.6. Proof of Proposition 5

(i) and (ii). The Type one EV1 shocks at each revision opportunity gives the logit

$$\sigma_i(Q_{-i}, \mathcal{E}_t) = \frac{\exp\{v_i^1(Q_{-i}, \mathcal{E}_t)/\sigma_\epsilon\}}{\exp\{v_i^1(Q_{-i}, \mathcal{E}_t)/\sigma_\epsilon\} + \exp\{v_i^0(Q_{-i}, \mathcal{E}_t)/\sigma_\epsilon\}}.$$

Proposition 4 gives $v_i^1(Q_{-i}, \mathcal{E}_t) - v_i^0(Q_{-i}, \mathcal{E}_t) = \Delta_i \Phi_t(Q_{-i}; W)$. Taking log odds gives (i). Dividing the numerator and denominator by $\exp\{v_i^0(Q_{-i}, \mathcal{E}_t)/\sigma_\epsilon\}$ gives (ii).
(iii). This adapts Blume (1993, Theorem 5.1 and Appendix 1) to a finite coalition $2^{\bar{Q}_t}$. Let

$$Z_t = \sum_{\widetilde{Q} \subseteq \bar{Q}_t} \exp\{\Phi_t(\widetilde{Q}; W)/\sigma_\epsilon\},$$

denote the partition function and define a probability distribution over coalitions/configurations Q through the Gibbs state of its potential

$$\widetilde{\mu}(Q \mid \bar{Q}_t, \mathcal{E}_t) = \frac{\exp\{\Phi_t(Q; W)/\sigma_\epsilon\}}{Z_t}, \qquad Q \subseteq \bar{Q}_t. \tag{B.4}$$

What I call CoCCPs are what Blume (1993) calls selection probabilities of the stochastic revision process. If coalition configurations are distributed according to Equation B.4, then conditioning on the joint distribution of neighbor membership to the coalition gives the CoCCPs. Spefcifically, the conditional probabilities of $\widetilde{\mu}$ are exactly the selection probabilities of Blume's stochastic revision process, what I call CoCCPs. For any $i \in \bar{Q}_t$ and $Q_{-i} \subseteq \bar{Q}_t \setminus \{i\}$,

$$
\begin{aligned}
\Pr_{\widetilde{\mu}}\left(i \in Q \mid Q_{-i}, \bar{Q}_t, \mathcal{E}_t\right) &= \frac{\widetilde{\mu}(Q_{-i} \cup \{i\} \mid \bar{Q}_t, \mathcal{E}_t)}{\widetilde{\mu}(Q_{-i} \mid \bar{Q}_t, \mathcal{E}_t) + \widetilde{\mu}(Q_{-i} \cup \{i\} \mid \bar{Q}_t, \mathcal{E}_t)} \\
&= \frac{\exp\{\Phi_t(Q_{-i} \cup \{i\}; W)/\sigma_\epsilon\}}{\exp\{\Phi_t(Q_{-i}; W)/\sigma_\epsilon\} + \exp\{\Phi_t(Q_{-i} \cup \{i\}; W)/\sigma_\epsilon\}} \\
&= \frac{\exp\{\Delta_i \Phi_t(Q_{-i}; W)/\sigma_\epsilon\}}{1 + \exp\{\Delta_i \Phi_t(Q_{-i}; W)/\sigma_\epsilon\}} \\
&= \sigma_i(Q_{-i}, \mathcal{E}_t). \qquad \text{(B.5)}
\end{aligned}
$$

The first equality expresses i joining the coalition over its potential choices. The second substitutes in potentials. The third equality divides by $\exp\{\Phi_t(Q_{-i}; W)/\sigma_\epsilon\}$. And the last equality follows from item 43. Therefore, for every customer and every tentative coalition, the conditional probability distribution under $\widetilde{\mu}$ is the selection probability distribution of the stochastic revision process. As $\widetilde{\mu}$ is a reversible distribution via Blume (1993, Appendix 1), and every reversible distribution is stationary. Because the configuration space is finite, every customer has a positive revision rate, and every CoCCP assigns strictly positive probability to both actions, the revision process is irreducible as the configuration space cannot be subdivided into two or more closed groups. Therefore, it has a unique stationary distribution. A low mean waiting time $r \to 0$ also means that the stationary distribution is achieved rapidly. Therefore it must be the case that $\mu = \widetilde{\mu}$, and

$$
\mu(Q \mid \bar{Q}_t, \mathcal{E}_t) = \frac{\exp\{\Phi_t(Q; W)/\sigma_\epsilon\}}{\sum_{\widetilde{Q} \subseteq \bar{Q}_t} \exp\{\Phi_t(\widetilde{Q}; W)/\sigma_\epsilon\}}, \qquad Q \subseteq \bar{Q}_t. \qquad \text{(B.6)}
$$

(iv). A neighbor coalition of Q_{-i} exists when the total coalition is either Q_{-i} and $Q_{-i} \cup \{i\}$, so summing their probabilities gives (45).