Infrastructure Provision and Payment Evasion: Evidence from Energy Distribution in Brazil

Gabriel Dias Santamarina *

November 1, 2026

Updated regularly, please click here for the latest version.

Abstract

Electricity theft and outages remain widespread even in countries with near-universal electricity access. I study how regulators can improve welfare when consumers' payment decisions respond to service quality and firms' incentives to provide quality depend on whether consumers pay. I focus on Brazil, where roughly 15% of residential electricity consumption is unpaid. Using newly assembled administrative data on consumption, non-technical losses, outages, tariffs, and network infrastructure, I document a feedback loop between reliability and payment evasion. One additional annual interruption increases payment evasion by approximately 0.5 percentage points in lower-income areas, while a one-percentage-point increase in evasion reduces subsequent transformer replacement by approximately 3% of its mean. To quantify this mechanism, I estimate a structural model of household electricity demand and payment evasion together with firms' transformer replacement and outage restoration decisions. The model allows firms to value consumer surplus alongside profits, with observed operational decisions disciplining this weight. Counterfactuals show that eliminating payment evasion reduces outage frequency by 11.8% and raises paying households' welfare by R$5.0 per month, or 3.4% of the average bill. Allowing tariffs to vary across regions reduces payment evasion from 12.3% to 8.7% and increases consumer surplus by up to 10% in certain areas, holding the average tariff unchanged.

Keywords: Electricity Reliability; Payment Evasion; Infrastructure Regulation; Informality.

JEL Codes: L51, L94, Q41, O12, O17.

*Yale University, gabriel.santamarina@yale.edu. I gratefully acknowledge the support of the Yale Economic Growth Center. I thank Katja Seim, Mushfiq Mobarak, Rohini Pande, and Steven Berry for invaluable and frequent advice. I also thank Nick Ryan, Jingyuan Wang, Ken Gillingham, Mayara Felix, Andre Trindade, Marcelo Sant'Anna, Jingyi Cui, Ashley Langer, Carolina Tojal, Meichen Chen as well as seminar participants at Yale Industrial Organization and Development Workshops and Harvard-Yale Development Meeting for helpful comments and suggestions. Vivian Figer, Davi Resende and Henrique Lovisi gave great institutional insights. I thank ANEEL for promptly answering my multiple inquiries regarding the data. Sarah Lepkowitz provided excellent research assistance. All errors are my own.

1 Introduction

Electricity access does not guarantee reliable electricity service. In many developing countries, connection to the grid is nearly universal, yet households and firms continue to experience frequent and prolonged outages. Electricity distributors also face substantial payment evasion, which limits the revenue they collect from the services they provide. Reliability is generally studied as a question of infrastructure investment and regulation, and payment evasion as a question of enforcement and informality. Yet the two are closely related. The quality of electricity service can shape consumers' willingness to pay, while payment behavior can influence firms' incentives to maintain and improve the distribution network.

This paper studies this relationship in Brazil, where electricity access is widespread but outages remain frequent and unevenly distributed across locations (G1 (2024)). At the same time, approximately 15% of residential electricity consumption is not paid for (ANEEL (2026)). An important source of these losses is unmetered consumption: some households connect illegally to the grid or tamper with meters to avoid payment. I argue that reliability and payment evasion are jointly determined through a cycle of quality underprovision. First, poor reliability reduces the value of remaining in the formal payment system and increases the incentive to evade payment. Second, when a larger share of local electricity consumption goes unpaid, improvements in reliability generate less revenue for the distributor, weakening its incentives to replace aging equipment and restore service quickly. Third, lower investment and slower restoration lead to more frequent or prolonged interruptions, reinforcing the initial incentive to evade payment.

This cycle creates an externality within the local distribution network. When payment evasion changes the firm's reliability decisions, households that continue to pay also experience poorer service. Thus, two otherwise similar households can face different electricity reliability because of the payment behavior of other consumers connected to their networks. The interaction also complicates regulation: the returns to improving measured service performance depend on local payment behavior, even when consumer tariffs are regulated.

I investigate these mechanisms using administrative data from six Brazilian electricity distributors over 2021–2024, combined with weather records and Census demographics. The data include formal electricity consumption, non-technical losses, individual interruptions, residential tariffs, the age and location of distribution equipment, transformer replacement, and outage restoration. They reveal how payment behavior responds to reliability and how firms allocate investment and restoration effort across networks. Rather than inferring firm behavior only from realized outages or assuming that distributors maximize short-run profits, I use their actual operational and investment choices to estimate the costs and objectives underlying reliability provision.

I proceed in two parts. First, I provide empirical evidence on each direction of the feedback loop, using a distinct source of exogenous variation for each: extreme-wind events to identify the effect of interruptions on evasion, and generation-cost-induced tariff changes to identify the effect of evasion on investment. I report OLS estimates as benchmarks alongside the

instrumental-variables estimates, which are the main basis for interpreting these relationships causally.

To study the consumer response, I relate payment evasion to the number of interruptions experienced over the preceding twelve months. Establishing an irregular connection is a deliberate and potentially persistent decision, so accumulated service experience is more relevant than an isolated outage. I instrument interruptions with extreme-wind events, which disrupt local distribution infrastructure and provide variation in service quality conditional on local and time fixed effects and weather controls. The IV estimate implies that one additional interruption during the preceding year increases the share of unpaid electricity consumption by approximately 0.26 percentage points. As complementary evidence, the fixed-effects OLS association is larger in lower-income neighborhoods, reaching 0.38 percentage points; the subgroup estimate is not instrumented. These results suggest that unreliable service affects participation in the formal payment system, not only the value of electricity consumed.

I next examine whether payment evasion affects firms' provision of reliability. I focus on transformer replacement, a decision observed directly in the georeferenced equipment registry and linked to the future frequency of interruptions. To identify this response, I instrument payment evasion with variation in residential tariffs induced by changes in electricity generation costs. These cost changes are passed through to the prices households pay, affecting their incentives to evade payment, but do not directly change the distribution tariff retained by the utility or the funds available for investment. This distinction isolates a source of variation in payment behavior that does not mechanically alter distributors' investment resources and helps address endogeneity, including reverse causality from reliability decisions to evasion. The IV estimate is -0.00582 (standard error 0.00221), indicating that greater payment evasion reduces subsequent transformer replacement. The OLS estimates are also negative, including after controlling for the age distribution of installed equipment. I also examine the timing of these decisions: past evasion predicts current replacement, whereas future evasion does not predict earlier investment. These complementary patterns are consistent with distributors adjusting investment in response to payment behavior.

Transformer replacement links the estimated firm response to subsequent service quality. Older transformers are more likely to fail, so lower replacement rates increase expected interruptions as the equipment stock ages. I estimate the relationship between transformer age and failures in the structural analysis, alongside the effect of restoration effort on outage duration. Together with the two instrumental-variables estimates, these relationships trace the cycle from outages to evasion, investment, and future reliability.

The second part of the paper develops and estimates a structural model to quantify the consequences of this interaction and evaluate alternative policies. The reduced-form estimates identify the two directions of the feedback loop separately, but do not capture how households' payment decisions and firms' reliability choices adjust jointly in equilibrium. Quantifying the externality that nonpaying households impose on paying households therefore requires solving these responses together, including the firm's dynamic transformer-replacement deci-

sion. Welfare comparisons also require objects that the reduced-form estimates do not recover: households' valuation of reliability, the private costs of evasion, firms' replacement and restoration costs, and the weight they place on consumer surplus. Finally, the policies I consider—eliminating payment evasion and allowing region-specific tariffs—are not observed in the data. Regular residential tariffs are uniform within each distributor by regulation, so the effects of these policies must be computed using the estimated model.

On the demand side, households choose whether to participate in the formal electricity market or evade payment, and then choose electricity consumption conditional on that decision. Payment status depends on the regulated tariff and expected reliability, while consumption responds to realized service quality. The model allows tastes for electricity to be correlated with the decision to evade payment. Because individual consumption and installed capacity are not observed for households outside the formal system, I combine registered household counts, formal consumption, Census dwelling counts, and aggregate non-technical losses to recover both the number of evading households and the electricity they consume. This framework estimates how payment behavior responds to prices and reliability and recovers the expected private costs of evasion, which are central to evaluating enforcement and pricing policies.

On the supply side, distributors provide reliability through two decisions: replacing transformers to reduce the frequency of equipment-related interruptions, and allocating restoration resources to shorten outages. These margins have different cost structures and timing, but both depend on the value the firm places on reliability. The demand estimates recover the revenue implications of improving service and the associated gains to consumers. I then use observed replacement and restoration choices to estimate investment and operating costs and to discipline the firm's objective. In particular, I allow distributors to place weight on consumer surplus in addition to profits, capturing the regulatory, reputational, and other incentives that may matter for a regulated monopoly. The estimated weight on consumer surplus averages 0.23, relative to a normalized weight of one on profit. Observed decisions thus discipline firm conduct beyond static profit maximization.

I first use the estimated model to measure the reliability externality imposed by payment evasion on households that continue to pay. I consider a counterfactual in which the option to evade payment is eliminated, hold regulated tariffs fixed, and allow firms to re-optimize replacement and restoration decisions. I then compare the welfare of households already participating in the formal market under the observed and counterfactual levels of reliability. Eliminating evasion reduces outage frequency by 11.8% and raises the welfare of paying households by an average of R$5.0 per month, equivalent to 3.4% of the average electricity bill. The externality is substantially larger in poorer areas: for households in the bottom income quartile, the gain is R$21 per month, compared with R$1 for those in the top quartile. The corresponding decline in outage frequency is 27% in the bottom quartile, compared with 3% in the top. Payment evasion therefore imposes a particularly large cost on paying consumers in areas where nonpayment is most prevalent.

Finally, I use the model to study tariff regulation when payment evasion and reliability

respond endogenously to prices. I compare observed uniform tariffs with welfare-maximizing uniform and region-specific tariffs, allowing firms and households to adjust in each counterfactual equilibrium without changing enforcement. Regional pricing can account for differences in local payment incentives that a uniform tariff does not. Relative to current regulation, optimal regional tariffs reduce average payment evasion from 12.3% to 8.7%, lower the outage-frequency index from 100 to 95, and increase aggregate consumer surplus by 2.4% while leaving the average tariff unchanged. The benefits are concentrated in lower-income areas, where consumer surplus rises by 9.4%; the corresponding change in the highest-income quartile is -1%. These results illustrate how tariff design can improve reliability and welfare by changing payment behavior, even without expanding enforcement.

The paper contributes to the literature on electricity reliability by identifying formal payment as a channel through which service quality matters. This response creates an equilibrium spillover: payment evasion changes the reliability experienced by other consumers. It also contributes to the study of regulated firms by recovering their objectives from investment and operational decisions, and to the study of informality by quantifying its local welfare costs and the effects of tariff policies that account for endogenous payment behavior.

1.1 Related Literature

This paper relates to several strands of the literature. First, it contributes to work on the quality of electricity provision in low- and middle-income countries (LMICs). The closest paper is McRae (2015), which studies electricity distribution in Colombia and shows that firms have weak incentives to reduce outages in neighborhoods where government subsidies cover unpaid electricity consumption. In that setting, the additional revenue generated by reducing outages may be too small to justify the investment costs, leading to persistent underprovision of reliability. My paper differs from McRae (2015) in several ways. First, I use rich administrative data on both the demand and supply sides of the market, which allows me to estimate an equilibrium model in which household payment decisions and firms' reliability choices are jointly determined. Because payment evasion is modeled as a household choice, the framework allows consumer behavior to respond to firms' quality decisions and firms' investment incentives to respond to consumer nonpayment. Second, detailed operational data from distribution companies allow me to study the costs of reliability provision directly, incorporating outage-resolution effort and regional heterogeneity rather than relying on uniform cost measures provided by the regulator. Finally, the Brazilian setting allows me to study reliability provision as a recurring operational choice rather than as a one-time fixed investment decision.

Other related papers include Lipscomb et al. (2013), Lee et al. (2019), Lee et al. (2020), and Berkouwer et al. (2024), which study the effects of electrification projects and improvements in electricity service quality in developing countries. Another closely related paper is Cisse (2025), which studies reliability projects in Senegal to estimate willingness to pay for reductions in outage duration. I contribute to this literature by studying a margin through which reliability matters: households' decisions to pay for electricity. This margin generates

equilibrium spillovers beyond the direct effect of outages. By affecting payment behavior, reliability changes firms' incentives to invest in service quality, and therefore shapes the reliability experienced by other consumers on the same local distribution network. Closest to this margin, Dzansi et al. (2018) show that households in Accra exposed to more rolling blackouts accumulate larger unpaid balances, and argue that the resulting loss of revenue can create a revenue trap for utilities. I study a different margin of nonpayment, unmetered consumption rather than arrears on billed consumption, in a setting with near-universal access; I trace the reverse direction through the firm's transformer-replacement decisions; and I embed both directions in an equilibrium model.

The paper also speaks to the industrial organization literature on electricity demand estimation and payment behavior (Dubin and McFadden (1984), Reiss and White (2005), Ito (2014), Jessoe and Rapson (2014), Jack and Smith (2020)). The closest paper in this strand is Resende et al. (2025), which studies how electricity prices affect payment evasion in Rio de Janeiro. I build on this insight but depart from it by treating payment evasion as a margin shaped not only by prices, but also by the quality of electricity service. This distinction is important because reliability changes the value of remaining connected to the formal system, while payment evasion changes the revenues and incentives of distribution companies. The paper therefore links demand estimation to a broader equilibrium problem in which household payment decisions and firms' quality choices are jointly determined. The paper also connects to the literature on the regulation of natural monopolies (Gagnepain and Ivaldi (2002), Fowlie (2010), Lim and Yurukoglu (2018), Abito (2020), Gowrisankaran et al. (2024), dos Santos and Guidetti (2025)). Much of this literature studies settings in which consumers participate in formal markets and enforcement is relatively strong. By contrast, I study a regulated infrastructure market in which a substantial share of consumers opts out of the formal payment system and enforcement is limited. In this environment, standard regulatory instruments may operate differently: prices affect not only consumption but also formal participation, and performance incentives affect firms in ways that depend on the local prevalence of nonpayment. The paper further relates to the empirical literature on regulation and firm conduct (Timmins (2002), Seim and Waldfogel (2013), Wilson (2025)) by using detailed operational data to discipline the objective functions of regulated firms. Rather than imposing profit maximization, I use data on equipment replacement, outage restoration, and local network conditions to estimate a weight on consumer surplus in the firm's objective, which summarizes how firms' observed decisions depart from profit maximization.

The paper also contributes to two literatures in development economics. First, it relates to the literature on poverty traps (Banerjee and Newman (1993), Balboni et al. (2022), Corbae and Glover (2025)). Much of this literature studies self-reinforcing mechanisms that prevent individuals from escaping a low-income equilibrium. The feedback loop I document is a candidate mechanism of this kind, operating through infrastructure provision: households in poor, high-evasion areas pay for less of the electricity they consume, and firms respond by investing less in the network that serves them. I do not show that this interaction produces multiple equi-

libria or persistence after temporary shocks; whether it does is a question the estimated model can address. Second, the paper speaks to the literature on tax morale and on the reciprocal relationship between citizens' payments and what the state provides (Alm et al. (1992), Luttmer and Singhal (2014), Dwenger et al. (2016)). Szabó and Ujhelyi (2015) show that a randomized water education campaign among low-income South African households substantially reduced nonpayment for water, and that the effect did not operate through improved knowledge. Weigel (2020) shows that citizens in the Democratic Republic of Congo respond to the state's effort to tax them by engaging more with the state. In the setting I study, households may stop paying for a state-regulated service when its quality deteriorates, which parallels the argument that compliance with payment obligations depends on what the state, or a state-regulated provider, delivers in return.

2 Setting and Data

2.1 Setting

Electricity distribution in Brazil is a regulated natural monopoly. Each municipality is served by a single distribution company, and consumer tariffs are regulated by the National Electric Energy Agency (ANEEL, *Agência Nacional de Energia Elétrica*). Prices are revised in periodic tariff-review cycles and readjusted annually. Final consumer prices combine the distribution tariff with pass-through charges related to generation costs, taxes, and energy tariff flags that vary with system-wide hydrological conditions. Within each distribution company, the regular residential tariff is linear and uniform across consumers, except for low-income households enrolled in the federal Social Electricity Tariff program, who face subsidized increasing-block tariffs.

Access to electricity in Brazil is nearly universal, but reliability varies substantially across locations. Most interruptions originate in the distribution network, due to transformer failures, deteriorating poles and wires, or weather events such as strong winds and falling trees. Reliability is regulated locally: ANEEL sets limits on the duration and frequency of interruptions (DEC and FEC, the Brazilian equivalents of SAIDI and SAIFI) within regulatory areas called consumer-unit groups, or CONJs. Poor performance relative to these limits can affect the allowed return incorporated into consumer tariffs, and registered consumers whose individual continuity indicators exceed regulatory limits receive compensation on their bills. This regulatory structure matters for the model because firms choose local reliability through maintenance, transformer replacement, capacity upgrades, and outage response, while prices are largely determined by regulation.

Payment evasion is widespread. Roughly 15% of residential electricity consumption in Brazil is not paid for (ANEEL, 2026). In this paper, payment evasion refers to unmetered electricity consumption, which occurs when households connect illegally to the grid by bypassing meters or tamper with their meters to avoid measurement. I exclude bill default by registered consumers, since billed consumption enters the distributor's energy balance regardless of whether the bill is

ultimately paid. Consumers caught evading payment may be required to repay the electricity consumed, face monetary penalties, or, in cases of illegal connections, face criminal sanctions (ConJur, 2019). Illegal connections may also damage the distribution network and increase the risk of outages.

2.2 Data and Variable Construction

The main unit of analysis is a CONJ-month, at which I measure payment evasion, reliability, tariffs, demographics, and formal participation. Feeders are distribution circuits that carry electricity from substations to connected consumers. They enter the analysis primarily in the construction of nontechnical losses and the estimation of the reliability technology, where within-feeder comparisons absorb local network conditions. I combine six data sources: the Distribution Utility Geographic Database (BDGD), additional administrative records from ANEEL, the 2022 national address register (CNEFE), the Brazilian Demographic Census, weather data from Copernicus, and geographic data from Topodata.

The primary source is the *Base de Dados Geográfica da Distribuidora* (BDGD), a georeferenced administrative database maintained by distribution companies and standardized by ANEEL. At the consumer-unit level, it reports monthly electricity consumption, tariff class, installed electrical load, and measures of service interruptions. Each consumer unit is linked to its CONJ, medium-voltage feeder, and serving distribution transformer. These records provide the main variables for estimating household electricity demand. The BDGD also contains detailed information on distribution-network assets, including transformer capacity, energization dates, substations, and distribution lines. I use these records to characterize the installed capital stock and identify transformer installations, replacements, and capacity upgrades, which are central to the supply-side model.

I construct nontechnical losses from the energy balances reported in the BDGD. These losses are defined as energy injected into the distribution network minus billed consumption and estimated technical losses. They are computed at the feeder level and aggregated to CONJs. Nontechnical losses serve as the measure of payment evasion throughout the paper. This measure includes unmetered consumption by households without registered connections and by registered consumers who tamper with their meters. It may also capture unmetered non-residential consumption and errors in estimated technical losses. Figure A4 illustrates the substantial spatial heterogeneity in non-technical losses across feeders, including within the same distribution company. Appendix B.1 describes the construction and cleaning of the non-technical-loss measure, including the treatment of negative energy balances, reporting irregularities, and outliers. Table A2 reports the resulting sample by distributor.

I complement the BDGD with ANEEL's administrative records of individual interruption events, which report their causes and durations, the number of registered consumers affected, and the associated CONJ and substation. I use these records to construct the frequency and duration of interruptions at the CONJ-month level. ANEEL also reports realized DEC and FEC indicators and their regulatory limits.

To measure formal participation, I combine registered residential consumer units from the BDGD with the 2022 *Cadastro Nacional de Endereços para Fins Estatísticos* (CNEFE), a georeferenced address register compiled during the Brazilian Census. Unlike registered consumer units, CNEFE includes dwellings without formal electricity connections, allowing me to estimate the total number of households in each CONJ. I assign dwellings to CONJs using the nearest registered electricity meter rather than regulatory polygons, some of which contain outdated boundaries. The formal participation share is constructed as the ratio of registered residential consumer units to the total number of dwellings in each CONJ. Appendix B.4 provides additional details on the consumer-unit sample restrictions, the geographic assignment of dwellings to CONJs, and the construction of the formal participation measure.

I also incorporate socioeconomic characteristics from the Brazilian Demographic Census, including household income, income dispersion, race, and access to basic services. These variables are originally reported at the census-tract level and aggregated to CONJs using registered consumer units as weights.

Finally, I merge these administrative records with weather and geographic data. From the Copernicus Climate Data Store, I obtain high-resolution measures of near-surface wind speed, temperature, and precipitation, which I use to construct local weather conditions and extreme-weather indicators. Geographic characteristics, including elevation, slope, and terrain ruggedness, come from Topodata, based on Shuttle Radar Topography Mission data. Figure A7 illustrates the geographic variation in elevation and terrain across the states covered by the distribution companies in the sample. Appendix B.5 provides further details on the construction of Census covariates, regulated residential tariffs, weather measures, and topographic characteristics.

3 Brazil's Quality Underprovision Cycle

This section provides evidence on the two mechanisms underlying the feedback between electricity reliability and payment evasion. I first study whether exposure to service interruptions induces households to obtain electricity outside the formal payment system. I then examine whether payment evasion affects utilities' subsequent investment in reliability, measured by the replacement of distribution transformers. In each case, I use an instrumental-variables strategy to isolate variation in the endogenous decision of interest: extreme-wind events for service interruptions and generation-cost-induced residential tariff changes for payment evasion. Together, these analyses establish the empirical basis for the equilibrium model developed below.

3.1 Service Quality Affects Payment Evasion

The first mechanism operates through households' incentives to participate in the formal electricity market. Repeated outages reduce the value of the service for which households pay and may consequently increase the attractiveness of irregular connections. Because establishing such a connection requires a deliberate change in payment behavior, I focus on accumulated exposure to interruptions rather than on the effect of an isolated outage. Figure A8 illustrates

the relationship between household income and electricity reliability, showing average outage duration in Panel (a) and monthly interruption frequency in Panel (b), both by income decile. Together, the two panels document how the duration and frequency of service interruptions vary across income groups.

Identification. The principal empirical challenge is that realized reliability is endogenous to local economic conditions and the operation of the distribution network. For example, changes in local infrastructure or payment evasion may jointly influence the frequency of outages and the share of electricity consumed outside the formal system. To identify the effect of reliability on evasion, I exploit interruptions induced by extreme-wind events. Sustained high winds damage distribution lines and other equipment, increasing the incidence of service interruptions. Conditional on local and time fixed effects and temperature controls, the timing of these weather events provides plausibly exogenous variation in households' accumulated exposure to outages.

The instrument is the number of extreme-wind days during the preceding twelve months, constructed over the same horizon as the interruption measure. Its relevance follows from the physical effect of high winds on the distribution network. The exclusion restriction is that, conditional on the included controls and fixed effects, past extreme-wind exposure affects current payment evasion through the resulting interruptions, rather than through an independent change in households' underlying propensity to obtain electricity informally. This variation therefore distinguishes the effect of service quality on payment behavior from the reverse effect of evasion on network performance.

Let FEC_{jt} denote interruptions per consumer unit in CONJ j during month t. I estimate

$$ntl_{jt} = \beta \, \overline{FEC}_{jt} + \alpha_X' X_{jt} + \alpha_j + \tau_t + \epsilon_{jt}, \qquad \overline{FEC}_{jt} \equiv \sum_{k=1}^{12} FEC_{j,t-k}, \tag{1}$$

where ntl_{jt} is the share of electricity delivered to CONJ j classified as non-technical losses, X_{jt} contains local temperature controls, and α_j and τ_t are CONJ and month fixed effects. Non-technical losses capture unmetered electricity consumption, rather than overdue payments on metered bills. I instrument $\overline{FEC}_{jt}$ with the corresponding twelve-month count of extreme-wind days. The coefficient β measures the effect of one additional interruption over the preceding year on the share of electricity consumed without payment. I also estimate the relationship separately for CONJs with mean household income below the sample median, using the 2022 Census.

Results. Table 1 reports the estimates. The IV coefficient in column (2) is 0.00263 (standard error 0.00141): one additional interruption over the previous twelve months raises the share of non-technical losses by approximately 0.26 percentage points. Panel B documents the relevance of the weather instrument. An additional extreme-wind day increases accumulated interruptions by 0.30 (standard error 0.06), and the first-stage F-statistic is 24. The fixed-effects OLS estimate in column (1) is 0.00285, corresponding to a 0.29 percentage-point increase, close to

the IV estimate.

The response is larger in lower-income areas. Column (3) reports an OLS coefficient of 0.00382 for CONJs with below-median mean household income, implying that one additional annual interruption is associated with a 0.38 percentage-point increase in payment evasion. The evidence thus indicates that accumulated service disruptions affect households' participation in the formal electricity market, particularly in economically disadvantaged areas.

Table 1: Accumulated Service Interruptions and Payment Evasion

	All CONJs		Below-Median Income
	OLS	IV	OLS
Panel A. Payment Evasion			
Interruptions, past 12 months	0.00285**	0.00263*	0.00382**
	(0.00138)	(0.00141)	(0.00163)
Panel B. First Stage: Interruptions, past 12 months			
Extreme-wind days, past 12 months	–	0.30	–
		(0.06)	
First-Stage F-Statistic	–	24	–
Observations	21,603	21,603	10,611
CONJ FE	✓	✓	✓
Month FE	✓	✓	✓
Local Controls	✓	✓	✓

Notes: The dependent variable is the share of electricity delivered to the CONJ classified as a non-technical loss. The main regressor is the total number of interruptions lasting at least three minutes during the preceding twelve months. Column (2) instruments interruptions using the number of days with sustained wind speeds of at least 64 km/h (40 mph) for at least one hour, accumulated over the same twelve-month window. Column (3) restricts the sample to CONJs with below-median mean household income in the 2022 Census. All specifications include CONJ and month fixed effects and local controls. Standard errors are two-way clustered by CONJ and month. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

As complementary evidence, I estimate a long-differences specification relating changes in interruption frequency to changes in non-technical losses between 2021 and 2024. Including distributor fixed effects, this exercise compares CONJs within the same concession area. One additional annual interruption is associated with a 0.36 percentage-point increase in the share of unpaid electricity ($p = 0.04$). The long-differences results reinforce the relationship between persistent service disruptions and payment evasion. As an additional identification strategy, Table A1 reports estimates that instrument average outage duration using the mean and standard deviation of local elevation. The IV estimates are also positive and statistically significant, providing complementary evidence that poorer electricity reliability increases payment evasion.

3.2 Payment Evasion Affects Distribution Quality

The second mechanism concerns the effect of payment evasion on the firm's incentives to provide reliability. When a larger share of electricity delivered to a local network is not paid for, improvements in service quality generate a smaller return in collected revenue. I study this mechanism through a direct firm decision: replacing distribution transformers. The admin-

istrative records identify the installed stock, the age of each transformer, and the equipment replaced in each year. Appendix B.3 details how I identify physical equipment replacements across annual asset snapshots, distinguishing replacements from changes in equipment identifiers. Unlike realized outages, replacement decisions measure an investment margin chosen by the distributor and directly linked to the future frequency of equipment-related interruptions.

Identification: residential tariff pass-through. I identify the effect of payment evasion on transformer replacement using changes in residential electricity tariffs generated by variation in electricity generation costs. In Brazil, the regulated price paid by households includes a component that passes through the cost of electricity generation. Changes in this component alter the tariff faced by residential consumers and hence their incentive to evade payment: a higher tariff raises the private gain from obtaining electricity without paying for it.

The key institutional distinction is between the *price paid by households* and the *revenue retained by the distributor*. The generation-cost component is passed through to consumers; it is not an increase in the distribution tariff retained by the utility. Thus, these shocks do not directly increase the distributor's regulated distribution margin or provide additional funds for transformer investment. They shift payment behavior without mechanically changing the financial return to replacing a transformer. The identification strategy therefore follows the chain

$$\text{Generation costs} \longrightarrow \text{Residential tariffs} \longrightarrow \text{Payment evasion} \longrightarrow \text{Replacement}$$

The instrument addresses the endogeneity of payment evasion, including reverse causality from investment and reliability to theft and unobserved local conditions affecting both variables. Relevance requires that generation-cost-induced residential tariff changes shift payment evasion; the first-stage estimates below document this relationship. The exclusion restriction is that, conditional on the fixed effects and controls, the pass-through shocks affect transformer replacement through their effect on payment evasion, rather than through an independent change in distributors' retained revenues, investment resources, or incentives to replace equipment. In particular, the identifying variation comes from the generation-cost component of the consumer tariff, not from changes in the remuneration received by distributors for their distribution services. Because consumer tariffs can also affect metered electricity demand, the exclusion restriction additionally requires that tariff-induced changes in formal consumption do not independently alter replacement incentives at a given level of evasion. Under these assumptions, the IV coefficient isolates the response of investment to payment evasion.

I estimate the following relationship at the CONJ-year level:

$$R_{jt} = \beta\, ntl_{j,t-1} + \gamma_1 \bar{a}_{j,t-1} + \gamma_2 \sigma^a_{j,t-1} + \alpha'_X X_{jt} + \alpha_j + \tau_t + \epsilon_{jt}, \tag{2}$$

where R_{jt} is the number of transformers replaced during year t as a share of the installed stock in CONJ j; $ntl_{j,t-1}$ is the preceding year's non-technical-loss share; and $\bar{a}_{j,t-1}$ and $\sigma^a_{j,t-1}$ are the

lagged mean and standard deviation of transformer age. Depending on the specification, I include CONJ and year fixed effects and controls for local temperature. The age controls account for replacement needs determined by the condition of installed capital. I instrument $ntl_{j,t-1}$ with generation-cost-induced variation in residential tariffs over the corresponding period.

Results. Table 2 presents the results. I begin with the IV specification in column (5), which uses variation in payment evasion induced by residential tariff pass-through. Panel B reports a positive first-stage coefficient of 0.09 (standard error 0.02), with a first-stage F-statistic of 12.98. This relationship is consistent with higher consumer prices increasing the incentive to evade payment. The second-stage coefficient in Panel A is −0.00582 (standard error 0.00221), indicating that greater payment evasion reduces the subsequent rate of transformer replacement.

Columns (1)–(4) provide complementary OLS estimates. The coefficients are negative throughout, ranging from −0.00367 to −0.00482, and remain statistically significant after controlling for additional moments of the transformer-age distribution and including CONJ and year fixed effects. The preferred fixed-effects estimate in column (4), −0.00482 (standard error 0.00222), is similar to the IV estimate. Taken together, the estimates show that areas with higher payment evasion receive fewer transformer replacements, even after accounting for the age and composition of installed equipment.

Table 2: Payment Evasion and Transformer Replacement

Dep. Variable: Replacement Rate	(1) OLS	(2) OLS	(3) OLS	(4) OLS	(5) IV
Panel A. Replacement Rate					
Nontechnical Losses, $t-1$	−0.00396**	−0.00367**	−0.00379**	−0.00482**	−0.00582***
	(0.00170)	(0.00168)	(0.00171)	(0.00222)	(0.00221)
Panel B. First Stage: Nontechnical Losses, $t-1$					
Generation-Cost-Induced Price Variation	–	–	–	–	0.09
					(0.02)
First-Stage F-Statistic	–	–	–	–	12.98
Observations	1,687	1,687	1,685	1,687	1,687
Lagged Stock Mean Age	✓	✓	✓	✓	✓
Lagged Stock Age Moments		✓	✓	✓	✓
CONJ FE			✓	✓	✓
Year FE				✓	✓

Notes: The dependent variable is the number of distribution transformers replaced during the year as a share of the installed stock. Panel A reports OLS estimates in columns (1)–(4) and the IV estimate in column (5). Panel B reports the first stage for column (5), which instruments lagged non-technical losses with generation-cost-induced changes in regulated residential tariffs. All specifications control for local temperature; transformer-age controls are measured at the end of the previous year. Standard errors are clustered by CONJ. The F-statistic tests the excluded instrument. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Timing and complementary evidence. The instrumental-variables strategy provides the principal source of identification. I further use the timing of observed evasion and replacement

decisions to examine whether the data exhibit the pattern implied by a firm response to nonpayment. The main specification relates current investment to *lagged* payment evasion, so that the payment behavior precedes the observed replacement decision. This timing also corresponds to a firm's ability to use information about previous losses when planning maintenance and replacement.

I complement this specification with a placebo test based on *future* payment evasion. If investment decisions respond to previously observed losses, future evasion should not predict earlier transformer replacement. Table 3 reports the comparison. The coefficient on future evasion is 0.00303 (standard error 0.00190) and is not statistically significant, whereas the coefficient on lagged evasion is -0.00379 (standard error 0.00171) and is statistically significant. The absence of a corresponding association with future evasion is consistent with the timing of the mechanism: firms adjust investment after observing payment losses. This placebo is supporting evidence for the interpretation of the main IV result, rather than a substitute for the exclusion restriction.

Table 3: Timing of Payment Evasion and Transformer Replacement

	Lead	Lag
Nontechnical Losses	0.00303	-0.00379^{**}
	(0.00190)	(0.00171)
Observations	1,685	1,685

Notes: The dependent variable is the transformer replacement rate. The columns report separate specifications relating current replacement to a lead or a lag of non-technical losses. Both specifications control for transformer-age characteristics, CONJ fixed effects, and time effects. Standard errors are clustered by CONJ. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

I also estimate a long-differences specification comparing changes in payment evasion and transformer replacement between the beginning and end of the sample. With distributor fixed effects, the identifying comparisons are across CONJs in the same concession. Locations experiencing greater increases in non-technical losses also experience larger declines in replacement, with an estimated coefficient of -0.0176 ($|t| = 4.12$). This longer-horizon pattern is consistent with the panel and IV estimates.

Overall, the estimates provide evidence for both sides of the quality-underprovision cycle. Extreme-wind-induced interruptions increase payment evasion on average, and the association is stronger in lower-income areas. Generation-cost-induced variation in payment evasion, in turn, reveals a negative effect on subsequent transformer replacement. The timing placebo and additional long-horizon evidence support the interpretation that utilities respond to observed payment losses by reducing reliability-related investment. The structural model developed in the following sections combines these behavioral responses, links investment and restoration choices to realized reliability, and quantifies their equilibrium and welfare implications.

4 Demand Model

To understand the determinants of electricity consumption and households' decisions to engage in electricity payment evasion, I develop a two-stage demand model. In the first stage, households choose whether to steal electricity, taking an action $a \in \{0, 1\}$, where $a = 1$ denotes theft and $a = 0$ denotes participation in the formal market. In the second stage, conditional on this decision, households choose their level of electricity consumption.

Households are indexed by i, regions by j, and time by t; a region is a CONJ. Electricity reliability is denoted by r_{jt} and captures the effective availability of electricity, the fraction of hours in month t in which supply is available, $r_{jt} = 1 - F_{jt}d_{jt}/T$, as defined in equation (13) of Section 5. The two decisions are taken on different information. The regime decision is deliberate and persistent, so households take it on expected reliability, which I measure as availability over the preceding twelve months,

$$\bar{r}_{jt} \;\equiv\; \frac{1}{12}\sum_{k=1}^{12} r_{j,t-k}. \tag{3}$$

Consumption is chosen after reliability is realized and responds to r_{jt}. The model is solved by backward induction: I first characterize the conditional demand and the resulting indirect utility under each regime, and then study the regime-choice problem.

Second Stage

Conditional on the choice of a, households choose electricity consumption q^a_{it}. Preferences are quasi-linear in a numéraire good, so welfare is money-metric. Reliability enters multiplicatively on the marginal value of consumption, so that electricity and service availability are complements: an additional kilowatt-hour is worth more to a household whose supply is dependable. Let X_{ij} denote household and local demographics and c_i the household's installed capacity. The per-period utility derived from consuming q units of electricity is

$$u(q;\, r_{jt}, c_i, X_{ij}, \chi_{jt}, \nu_{it}) \;=\; A_{it}\cdot r_{jt}^{\gamma}\cdot\frac{q^{\,1-1/\eta} - q_0^{\,1-1/\eta}}{1-1/\eta}, \qquad \log A_{it} \;=\; a_0 + a_c \log c_i + a_X' X_{ij} + \chi_{jt} + \nu_{it}, \tag{4}$$

where $\eta > 0$ is the price elasticity of demand, $\gamma > 0$ governs the elasticity of the utility shifter with respect to reliability, $q_0 > 0$ is a reference level of consumption, and A_{it} is a taste shifter with an observable component, which depends on installed capacity and demographics, a market-level unobservable χ_{jt}, and a household-level unobservable ν_{it}.

The market-level component χ_{jt} captures local demand conditions, such as economic activity or heat beyond the temperature controls, that the firm may observe when choosing reliability and that may covary with tariffs. It is the reason tariffs and reliability are instrumented. The household component ν_{it} is independent of tariffs, reliability, the market-level unobservables, and the instruments, conditional on c_i and X_{ij}.

The reference level q_0 fixes the level of utility and plays no role in demand or in regime

choices. Without it, for $\eta < 1$, the empirically relevant case, the term $q^{1-1/\eta}/(1-1/\eta)$ is negative, so a higher value of $A_{it}r_{jt}^{\gamma}$ would lower utility: demand would rise with reliability while welfare fell. With q_0 below the consumption levels in the data, utility is positive and increasing in reliability. I interpret q_0 as a subsistence level of consumption, below which service has no value.

Formal market. A household participating in the formal market faces the linear tariff p_{ut} and solves $\max_q \; u(q;\, r_{jt}, c_i, X_{ij}, \chi_{jt}, \nu_{it}) - p_{ut}\, q$. The first-order condition yields the conditional demand

$$q_{it}^{F} \;=\; \left(\frac{A_{it}\, r_{jt}^{\gamma}}{p_{ut}}\right)^{\eta}, \qquad \log q_{it}^{F} \;=\; \eta\, a_0 + \eta\, a_c \log c_i + \eta\, a_X' X_{ij} - \eta \log p_{ut} + \eta\gamma \log r_{jt} + \eta\, \chi_{jt} + \eta\, \nu_{it}. \tag{5}$$

Equation (5) is the formal-market demand equation that disciplines the estimation of the demand-side parameters; it does not depend on q_0. Note that the elasticity of consumption with respect to reliability is the product $\eta\gamma$ rather than γ alone: reliability raises the marginal value of consumption, and the household's response to that shift is mediated by its willingness to substitute. The associated per-period indirect utility is

$$v^{F}(p_{ut}, r;\, c_i, X_{ij}, \chi_{jt}, \nu_{it}) \;=\; \frac{\eta}{1-\eta}\, A_{it}\, r^{\gamma}\, q_0^{\,1-1/\eta} \;-\; \frac{\left(A_{it}\, r^{\gamma}\right)^{\eta} p_{ut}^{\,1-\eta}}{1-\eta}, \tag{6}$$

which increases with reliability whenever formal consumption exceeds the reference level:

$$\frac{\partial v^{F}}{\partial r} \;=\; \frac{\gamma}{r}\,\frac{\eta}{1-\eta}\, A_{it}\, r^{\gamma}\left[q_0^{\,1-1/\eta} - \left(q_{it}^{F}\right)^{1-1/\eta}\right] \;>\; 0 \quad\Longleftrightarrow\quad q_{it}^{F} > q_0.$$

Payment evasion. Households that evade payment do not face the formal tariff—the marginal price of consumption is zero—but they cannot consume without bound. Consumption is capped at a technical capacity $\bar{q}_i$, which depends on the household's installed capacity and on the physical characteristics of the illegal connection. I treat $\bar{q}_i$ as a household-specific object, jointly determined by installed electrical capacity, which I observe in the data, and local demographics. Recovering this object is not straightforward. Because observed prices are not close to zero, extrapolating formal demand to a zero price would be uninformative. I therefore discipline $\bar{q}_i$ using three pieces of information: the number of households in each area that do not hold a registered consumer unit, total unmetered electricity consumption, and the observed distribution of installed capacity among consumers in the formal market. More specifically, I model consumption under an illegal connection as a deterministic function of the household's installed capacity and local characteristics:

$$\log \bar{q}_i = \kappa_0 + \kappa_1 \log c_i + \kappa_X' X_j, \tag{7}$$

where c_i is the installed capacity of household i and X_j is a vector of local characteristics in region j. The parameters of (7) are estimated jointly with the other demand parameters

(Section 7). The empirical challenge is that I observe these household-level characteristics only for consumers in the formal market. I do not observe individual characteristics for households outside the formal payment system. To address this issue, I use a Bayes-rule approach similar to Resende et al. (2025).

Let $f(c_i)$ denote the population distribution of installed capacity. Then,

$$\mathbb{P}(c_i \mid a_i = 0) = \frac{\mathbb{P}(a_i = 0 \mid c_i) f(c_i)}{\mathbb{P}(a_i = 0)}. \tag{8}$$

The left-hand side, $\mathbb{P}(c_i \mid a_i = 0)$, is observed in the data because I observe installed capacity for consumers in the formal market. The aggregate probability $\mathbb{P}(a_i = 0)$ is also observed, as the ratio of registered consumer units to dwellings (Section B.4). The conditional probability $\mathbb{P}(a_i = 0 \mid c_i)$ is recovered from the demand model. The remaining object, $f(c_i)$, is the population distribution of installed capacity, including both households in and outside the formal payment system. Recovering this distribution is crucial for counterfactual simulations, because counterfactual payment decisions depend on the installed-capacity distribution of all households, not only those observed in the formal market. Identification and estimation are discussed in later sections.

Evasion carries an expected cost that is absent from the formal regime. A household with an illegal connection may be detected, disconnected, and required to pay for the electricity it consumed. This expected cost is not independent of local network conditions: a utility that maintains its network in a given area also has crews, metering infrastructure, and monitoring capacity in that area, so enforcement intensity and reliability move together. I therefore allow the expected cost of evasion to depend on local reliability, on the scale of the illegal connection, and on the tariff through a flexible term $g(\cdot)$, and the resulting indirect utility from evasion is

$$v^S(\bar{q}_i, r, p_{ut};\, c_i, X_{ij}, \chi_{jt}, \nu_{it}) \;=\; u\big(\bar{q}_i;\, r, c_i, X_{ij}, \chi_{jt}, \nu_{it}\big) \;+\; \underbrace{\phi_1\, r \;+\; \phi_2\, r^2 \;+\; \phi_3\, r\, \bar{q}_i \;-\; \phi_4\, p_{ut}\, \bar{q}_i}_{g(r,\, \bar{q}_i,\, p_{ut})}. \tag{9}$$

The last term is expected repayment: a detected household is required to pay for the electricity it consumed at the formal tariff, so $\phi_4 \in [0, 1]$ is the expected share of stolen consumption eventually paid for. The sign and curvature of the reliability terms are left to the data. A negative ϕ_1 indicates that better local service raises the expected cost of evasion, so that reliability deters theft; a positive ϕ_2 implies this deterrent effect attenuates as reliability improves, consistent with enforcement capacity saturating. The interaction term ϕ_3 captures that larger illegal loads are more easily detected in the utility's loss accounting and carry larger liabilities when detected.

The reference level q_0 cancels from the comparison between regimes:

$$v^S - v^F \;=\; \frac{\big(A_{it}\, r^{\,\gamma}\big)^{\eta}\, p_{ut}^{\,1-\eta}}{1-\eta} \;-\; \frac{\eta}{1-\eta}\, A_{it}\, r^{\gamma}\, \bar{q}_i^{\,1-1/\eta} \;+\; g(r, \bar{q}_i, p_{ut}), \tag{10}$$

so choices do not identify it. Its role is to fix the level of the value of reliability, and for η well

below one the term in q_0 dominates $\partial v^F/\partial r$: the welfare value of reliability is governed by q_0 rather than by the curvature parameters. I fix q_0 at 1 kWh, the first percentile of monthly consumption in the demand sample (Appendix B.4). The same argument implies a broader limit on what the model identifies: any reliability term common to both regimes, $H(r)$, would leave demand and regime choices unchanged while changing willingness to pay for reliability. The specification rules such terms out by assumption: reliability affects welfare only through the value of electricity consumption.

Reliability therefore enters the evasion regime through two opposing channels. It raises the direct utility of stolen consumption, since capped consumption is worth more when the lights stay on, and it raises the expected cost of being caught. Which channel dominates is an empirical question and is precisely what governs whether investment in the network attenuates or amplifies payment evasion. Because g depends on reliability, holding the enforcement technology fixed in counterfactuals means holding the function g fixed, not the expected cost of evasion, which moves with reliability. The model does not separate monitoring and enforcement from other channels through which reliability changes the attractiveness of evasion; g summarizes all of them. It also assumes that reliability affects the cost of evasion in the same way whatever its source, so that an interruption caused by a storm and one caused by an aging transformer shift g equally.

4.1 First Stage

Based on the indirect utilities derived above and on local characteristics, households decide whether to engage in electricity theft or to participate in the formal electricity market. Households evaluate each regime at expected reliability $\bar{r}_{jt}$ and the current tariff. Total utility in each regime is

$$\begin{aligned} U_{it}^F &= \lambda\, v^F(p_{ut}, \bar{r}_{jt};\, c_i, X_{ij}, \chi_{jt}, \nu_{it}) \;+\; \xi_{jt}^F \;+\; e_{it}^F, \\ U_{it}^S &= \lambda\, v^S(\bar{q}_i, \bar{r}_{jt}, p_{ut};\, c_i, X_{ij}, \chi_{jt}, \nu_{it}) \;+\; \xi_{jt}^S \;+\; e_{it}^S, \end{aligned} \tag{11}$$

where $\lambda > 0$ converts the money-metric indirect utility into the scale of the choice problem, ξ_{jt}^a are region–time unobservables capturing local determinants of evasion that are not observed by the econometrician — the prevalence of existing illegal connections, local enforcement intensity, or social norms regarding payment — and e_{it}^F, e_{it}^S are idiosyncratic regime-specific preference shocks. The household chooses $a = 1$ if and only if $U_{it}^S > U_{it}^F$.

The household-level unobservables in the model are the taste shock ν_{it} entering the utility shifter in (4), and the preference shocks e_{it}^F and e_{it}^S entering the regime utilities in (11). I assume they are jointly normally distributed:

$$\begin{pmatrix} \nu_{it} \\ e_{it}^F \\ e_{it}^S \end{pmatrix} \;\sim\; \mathcal{N}(\mathbf{0},\, \Sigma).$$

Because only the difference $\omega_{it} \equiv e^F_{it} - e^S_{it}$ affects behavior, the individual variances of e^F_{it} and e^S_{it} are not separately identified, and the model is written in terms of ω_{it}. Two parameters of Σ are therefore estimable: the standard deviation σ_ν of the taste shock, identified from the dispersion of formal-market consumption in (5) within markets, and the correlation $\rho \equiv \text{corr}(\nu_{it}, \omega_{it})$. This correlation captures selection on unobservables: households with unobservedly high taste for electricity may also differ systematically in their unobserved attachment to the formal market. A negative ρ indicates that high-consumption households are, all else equal, more inclined to evade. Selection on taste does not require $\rho \neq 0$: because ν_{it} enters the formal and capped-evasion payoffs differently, participation depends on ν_{it} even when $\rho = 0$, and ρ measures selection over and above this structural channel. Estimating this correlation is essential for the welfare and counterfactual exercises, which require evaluating outcomes under counterfactual regimes for households whose realized regime may differ.

Similarly, only the difference $\Delta\xi_{jt} \equiv \xi^S_{jt} - \xi^F_{jt}$ affects behavior. The level of a component common to both regimes is irrelevant and normalized to zero; the difference itself is not normalized but recovered market by market.

To fix the scale of the choice problem I normalize $\text{Var}(\omega_{it}) = 1$. Conditional on the taste shock, the probability of evasion then takes the probit form

$$\Pr\left(a_{it} = 1 \mid c_i, X_{ij}, p_{ut}, \bar{r}_{jt}, \chi_{jt}, \Delta\xi_{jt}, \nu_{it}\right) = \Phi\left(\frac{\lambda\left[v^S_{it} - v^F_{it}\right] + \Delta\xi_{jt} - \rho\,\nu_{it}/\sigma_\nu}{\sqrt{1-\rho^2}}\right), \tag{12}$$

where v^F_{it} and v^S_{it} abbreviate (6) and (9) evaluated at $\bar{r}_{jt}$. The unconditional choice probability requires integrating (12) over the distribution of ν_{it}, which enters the indirect utilities nonlinearly and asymmetrically across regimes: it shifts formal consumption on both the intensive and extensive margins, while under evasion consumption is capped and only the extensive margin responds. This integral has no closed form, so choice probabilities, consumer surplus, and counterfactual outcomes are computed by simulation. Equation (12), together with the formal-market demand equation (5) and the joint normality of the unobservables, fully characterizes the joint distribution of consumption and regime choice in the model.

Limitation. Measuring expected reliability over twelve months matches the evidence that households respond to accumulated experience, but the model has no switching cost: it treats each month's regime as a fresh choice given expected reliability. The estimated extensive-margin responses should therefore be read as long-run responses to persistent changes in prices and reliability. This is also the horizon of the counterfactuals, which compare stationary equilibria.

5 Supply Side

I do not assume that electricity distributors maximize static local profits. Although profit maximization is a natural benchmark, it is unlikely to fully describe reliability provision in this setting for three reasons. First, an important margin of reliability provision is dynamic. Transformer replacement is a capital decision whose benefits accrue over the life of the asset,

not only in the current period. A static model would miss the forward-looking incentives that shape whether firms replace aging equipment before it fails.

Second, regulated distributors operate in a political and institutional environment that extends beyond local revenues and costs. They interact with regulators, local governments, and civil-society actors, and both outage penalties and the classification of interruptions involve administrative processes.[1] These features can create incentives to provide reliability that are not captured by variable profits alone.

Third, the cost of reliability varies substantially across regions. Network conditions, equipment age, terrain, and load all affect how costly it is to prevent interruptions or restore service after an outage. As a result, the same observed level of reliability can reflect very different underlying costs across locations. For these reasons, I estimate a supply-side model that allows firms' reliability choices to reflect both profit incentives and additional motives associated with regulation, political constraints, reputation, and consumer welfare.

The model presented in the next sections addresses each point. The firm supplies reliability through two margins: a dynamic replacement decision for each transformer, which governs how often interruptions occur, and a static restoration decision, which governs how long they last. The firm's objective places a weight λ^f on consumer surplus alongside profit; this weight nests profit maximization at $\lambda^f = 0$ and summarizes the motives in the second point. Costs on both margins depend on observed network and regional characteristics. Rather than imposing an objective, I let observed replacement and restoration decisions determine λ^f.

Together with the household demand model, these supply decisions determine the feedback between electricity reliability and payment evasion. Appendix C develops a simple model illustrating how this interaction can amplify shocks to reliability through households' payment decisions and distributors' investment incentives.

5.1 Firm's Problem

The unit of analysis is the CONJ, indexed by j, the geography at which the regulator measures and sanctions reliability. CONJ j contains N_j consumer units and is served by a set of distribution transformers $k \in j$; transformer k serves the set $\mathcal{N}_k$, containing N_k consumer units. In each period t, the firm decides whether to replace each transformer and chooses the mean restoration time d_{jt} in each CONJ. Replacement decisions are taken first; the firm then chooses d_{jt}.[2]

Reliability. Let F_{jt} denote the number of interruptions per consumer unit in CONJ j during period t (the FEC index) and d_{jt} the mean duration of an interruption (the ratio DEC/FEC).

[1] Outages caused by "extreme weather events" are not included in the calculation of SAIDI and SAIFI. The classification of whether an outage was caused by an extreme weather event depends, in part, on information and declarations from local governments.

[2] The asset registry is annual, while outage records are monthly. I write both problems at a common period t; in estimation, the replacement payoff aggregates monthly payoffs within the year.

Reliability, the fraction of hours in which service is available, is

$$r_{jt} = 1 - \frac{F_{jt}\, d_{jt}}{T}, \tag{13}$$

where T is the number of hours in the period. Interruptions decompose into a component caused by transformers and a component arising from all other causes:

$$F_{jt} = \sum_{k \in j} \frac{N_k}{N_j} F^{\mathrm{tr}}\big(\tilde{x}_{kt}\big) + F^{\mathrm{net}}_{jt}, \qquad \frac{\partial F^{\mathrm{tr}}}{\partial x} > 0, \qquad \tilde{x}_{kt} = \begin{cases} 0 & \text{if } y_{kt} = 1, \\ x_{kt} & \text{otherwise,} \end{cases} \tag{14}$$

where x_{kt} is the age of transformer k at the start of period t, $y_{kt} = 1$ denotes replacement, and $F^{\mathrm{tr}}(x)$ is the number of interruptions per period that a transformer of age x causes to each consumer unit it serves. Because replacement decisions are taken before restoration effort, frequency is evaluated at post-replacement ages $\tilde{x}_{kt}$: a unit replaced in period t operates at age zero during that period. The weights N_k/N_j arise because F_{jt} averages over consumer units, and an interruption caused by transformer k reaches only the units it serves. The term F^{net}_{jt} collects interruptions caused by the medium- and low-voltage network, connection lines, and weather. It depends on local weather and on stolen consumption in the previous period, $Q^S_{j,t-1}$, since illegal connections overload low-voltage lines. Because it depends on past rather than current theft, the firm takes it as given within the period; the firm does not internalize that current reliability, through future theft, changes future non-transformer faults.

Payoff. Let $D_j(r)$ denote formal (billed) consumption in CONJ j at reliability r, $Q^S_j(r)$ stolen consumption, $\mathcal{L}_j(r) = D_j(r) + Q^S_j(r)$ total load, and $CS_j(r)$ the expected money-metric surplus of all households in j, formal and evading. These objects come from the demand model, evaluated at the consumer tariff and aggregated over the households of j. Their derivatives with respect to r include both the intensive response of consumption and the movement of households across regimes. Because households choose their regime on reliability accumulated over the preceding year, a change in reliability affects regime shares gradually. I evaluate the derivatives at the steady state, in which a permanent change in r moves $\bar{r}$ one for one. This is the relevant derivative for replacement, whose effect on reliability persists over the life of the asset, and an approximation for restoration effort, whose effect is contemporaneous. The firm's per-period payoff from reliability in j is

$$\Pi_j(r) = \underbrace{p_t\, D_j(r) - w_t\, \mathcal{L}_j(r)}_{\text{variable profit}} + \lambda^f\, CS_j(r), \tag{15}$$

where p_t is the tariff the firm retains per billed kilowatt-hour, net of taxes and sector charges, and w_t is the cost of energy purchased per kilowatt-hour delivered, grossed up for technical losses. The tariff embeds the regulatory allowance for non-technical losses, fixed in advance, so an additional stolen kilowatt-hour costs the firm w_t and is not recovered. The parameter λ^f is

the weight distributor f places on consumer surplus relative to profit: $\lambda^f = 0$ corresponds to profit maximization and $\lambda^f = 1$ to a planner. It summarizes every motive to supply reliability beyond variable profit, including regulatory penalties, political pressure, and reputation, to the extent that these scale with the value of reliability to households. Appendix B.5 derives the consumer-surplus expressions for formal and evading households and shows how their marginal valuation of reliability enters the distributor's objective.

This is a substantive restriction. Regulatory incentives are tied to DEC and FEC relative to CONJ-level limits and to compensation owed to registered consumers whose individual indicators exceed their limits. Their marginal effect depends on the distance to the limit and on the number of registered consumers, not on the surplus of all households, including those who evade. To the extent that regulatory exposure varies across CONJs with evasion, λ^f partly captures it. Holding λ^f fixed in counterfactuals that change tariffs or remove evasion assumes that these incentives respond to the counterfactual as consumer surplus does.

Because each transformer serves a small share of its CONJ, the firm's problem separates into a static choice of restoration time in each CONJ and one replacement problem per transformer, in which CONJ-level conditions are taken as given.

5.1.1 Static Problem

In each period, the firm chooses how much restoration effort to allocate to each CONJ, trading off the value of shorter outages against the cost of crews, electricians, vehicles, and equipment:

$$\max_{d_{jt}} \quad \Pi_j\left(1 - \frac{F_{jt}\, d_{jt}}{T}\right) - \psi\big(d_{jt};\, \Xi_j, X_j, \upsilon_{jt}\big), \tag{16}$$

where ψ is the per-period cost of maintaining the crew capacity required to deliver mean restoration time d_{jt}:

$$\psi\big(d;\, \Xi_j, X_j, \upsilon_{jt}\big) = \psi_{0jt}\, d^{-\psi_1}, \qquad \log \psi_{0jt} = \theta_{\psi,0} + \theta'_{\psi,\Xi}\, \Xi_j + \theta'_{\psi,X}\, X_j + \upsilon_{jt}, \qquad \psi_1 > 0. \tag{17}$$

Faster restoration is costly, and increasingly so: $\psi_d < 0$ and $\psi_{dd} > 0$. The cost depends on the structure of the local network Ξ_j and on regional characteristics X_j, and υ_{jt} is a cost shock observed by the firm but not by the econometrician, such as crew availability or difficulty of access. The cost does not depend on the number of interruptions, which rules out congestion of crews when faults are numerous or clustered. The first-order condition equates the marginal benefit of faster restoration, weighted by the frequency of interruptions, to its marginal cost:

$$\frac{F_{jt}}{T}\, \Pi'_j(r_{jt}) = -\psi_d\big(d_{jt}\big) = \psi_1\, \psi_{0jt}\, d_{jt}^{-\psi_1 - 1}, \qquad \Pi'_j(r) = (p_t - w_t)\frac{\partial D_j}{\partial r} - w_t \frac{\partial Q_j^S}{\partial r} + \lambda^f \frac{\partial CS_j}{\partial r}. \tag{18}$$

I assume $\Pi'_j > 0$ and $\Pi''_j \leq 0$ over the relevant range. The first is necessary for an interior solution; the second ensures that (18) characterizes the optimum.

Two features of (18) are worth noting. First, interruption frequency F_{jt} scales the marginal benefit of effort: faster restoration is more valuable where outages are frequent, which links

this decision to the replacement problem below. Second, evasion lowers the profit component of the marginal benefit through two channels: evading households generate no revenue when better service raises their consumption, and the firm pays for the energy they draw. Consumer surplus CS_j includes households that evade, so where evasion is widespread, faster restoration generates social value that the firm internalizes only to the extent that $\lambda^f > 0$.

5.1.2 Dynamic Problem

Beyond restoration effort, the firm decides in each period whether to replace each transformer. Newer units are less likely to fail, so replacement lowers the frequency of interruptions rather than their duration. I model this decision as a dynamic discrete choice problem in the tradition of Rust (1987), in which the firm weighs a lower interruption rate and a lower risk of breakdown against the cost of replacement.

In each period the firm chooses $y_{kt} \in \{0, 1\}$, where $y_{kt} = 1$ denotes replacement. The state of unit k is its age $x_{kt} \in \{0, 1, \ldots, \bar{x}\}$, measured from the date of energization and observed in the asset records, together with the CONJ-level conditions $\Omega_{kt} \equiv \big(p_t,\, w_t,\, \xi_{jt},\, d_{jt},\, F_{-k,t}\big)$, where

$$F_{-k,t} \equiv \sum_{m \in j,\, m \neq k} \frac{N_m}{N_j}\, F^{\mathrm{tr}}\big(\tilde{x}_{mt}\big) + F^{\mathrm{net}}_{jt}$$

collects interruptions from all other sources. Because transformer k serves a small share of the CONJ, its replacement has a negligible effect on crew allocation and on the rest of the network, so the firm treats Ω_{kt} as unaffected by y_{kt}; I assume Ω_{kt} follows an exogenous first-order Markov process. The same approximation implies that the firm does not account for the effect of a single unit's replacement on the CONJ's future evasion: the steady-state regime response enters through Π'_j, but the gradual adjustment of households' expected reliability does not. Reliability in j as a function of the age of unit k is

$$r_j(x; \Omega_{kt}) = 1 - \left[\frac{N_k}{N_j}\, F^{\mathrm{tr}}(x) + F_{-k,t}\right] \frac{d_{jt}}{T}.$$

Age evolves according to

$$x_{k,t+1} = \begin{cases} 1 & \text{if } y_{kt} = 1, \\ 0 & \text{if } y_{kt} = 0 \text{ and the unit fails, with probability } h(x_{kt}), \\ \min\{x_{kt} + 1,\ \bar{x}\} & \text{otherwise.} \end{cases} \tag{19}$$

A replaced unit operates at age zero during period t. A unit that fails is replaced at the end of the period at cost RC_k and enters $t + 1$ at age zero. The failure hazard $h(x)$ is the probability that a unit of age x fails irreparably within the period, increasing in x. It is distinct from $F^{\mathrm{tr}}(x)$: most interruptions caused by a transformer, such as a blown fuse or an overload trip, do not destroy the unit. Observed replacements in the asset registry therefore combine chosen replacements and forced replacements after failure. Unless the two can be distinguished in the

data, the probability that a unit of age x is observed to be replaced is

$$\Pr\big(\text{observed replacement} \mid x, \Omega\big) \;=\; P(1 \mid x, \Omega) + \big[1 - P(1 \mid x, \Omega)\big]\, h(x), \tag{20}$$

where $P(1 \mid x, \Omega)$ is the probability of a chosen replacement, and the replacement likelihood uses this mixture.

The per-period payoff from each action is

$$u(x, y; \Omega) = \begin{cases} \Pi_j\big(r_j(0; \Omega)\big) - RC_k & \text{if } y = 1, \\ \Pi_j\big(r_j(x; \Omega)\big) - c(x; \theta_c) - h(x)\, RC_k & \text{if } y = 0, \end{cases} \tag{21}$$

where $RC_k = RC(\Xi_k, X_j)$ is the cost of installing a new unit, which depends on unit characteristics Ξ_k and regional characteristics X_j, and $c(x; \theta_c)$ is the cost of maintaining a unit of age x, increasing in x with $c(0; \theta_c) = 0$. Only the difference between the two payoffs affects the choice, so including the CONJ-level payoff Π_j in each unit's problem does not double count. Payment evasion enters through Π_j, with the same weight λ^f as in the static problem.

Replacement is forward-looking: the firm incurs a cost today in exchange for a stream of lower interruption and failure risk. Let $\varepsilon_{kt}(y)$ denote a cost shock observed by the firm but not by the econometrician, capturing transitory variation in crew availability, site access, or the scheduling of nearby work, distributed i.i.d. type-I extreme value with scale σ_ε. The value function satisfies

$$V(x, \Omega) = \mathbb{E}_\varepsilon \max_{y \in \{0,1\}} \Big\{ u(x, y; \Omega) + \varepsilon(y) + \delta\, \mathbb{E}\big[V(x', \Omega') \mid x, y, \Omega\big] \Big\}, \tag{22}$$

where δ is the discount factor and the expectation over x' follows (19).

To first order in N_k/N_j, the flow gain from replacing a unit of age x is

$$u(x, 1; \Omega) - u(x, 0; \Omega) \approx \Pi_j'(r_{jt})\, \frac{N_k}{N_j}\, \frac{d_{jt}}{T}\, \big[F^{\mathrm{tr}}(x) - F^{\mathrm{tr}}(0)\big] + c(x; \theta_c) - \big(1 - h(x)\big) RC_k. \tag{23}$$

The replacement cost enters net of $h(x)\, RC_k$, the expected cost of a forced replacement if the unit is kept. Because the reliability gain from replacing unit k is spread over the CONJ through the weight N_k/N_j and valued at the CONJ's marginal payoff Π_j', the replacement problem values reliability at the CONJ's mix of paying and evading households, not at the mix among the households unit k serves.

Duration and frequency as substitutes. Equations (13) and (14) place the two instruments in a common metric, expected hours without service, while keeping them distinct. Replacement lowers F^{tr}; restoration effort lowers d_{jt}; both raise r_{jt}. Because reliability is multiplicative in the two components,

$$\frac{\partial r_{jt}}{\partial d_{jt}} = -\frac{F_{jt}}{T}, \qquad \frac{\partial r_{jt}}{\partial F_{jt}} = -\frac{d_{jt}}{T}, \qquad \frac{\partial^2 r_{jt}}{\partial d_{jt}\, \partial F_{jt}} = -\frac{1}{T} \neq 0, \tag{24}$$

the return to each instrument depends on the level of the other. Comparing (18) with (23), both margins value reliability at the same marginal rate $\Pi_j'(r_{jt})$: restoration effort scales it by F_{jt}/T, and replacement by $(N_k/N_j)(d_{jt}/T)\big[F^{\text{tr}}(x) - F^{\text{tr}}(0)\big]$. Faster restoration is worth more where interruptions are frequent, and replacing a transformer is worth more where restoration is slow.

This shared valuation makes the two problems mutually disciplining rather than merely parallel. The instruments differ sharply in cost structure: replacement is a lumpy, irreversible outlay whose benefit accrues over the life of the asset, while restoration effort is a continuous flow expenditure with contemporaneous returns. They also differ in scope, since replacement acts only on F^{tr} while effort shortens all interruptions, including those in F_{jt}^{net} that the firm cannot prevent. Yet both are governed by the same weight λ^f: a firm that weighs household surplus at λ^f should do so whether it is deciding how quickly to dispatch a crew or whether to replace an aging unit. Estimating λ^f from both margins therefore yields an overidentifying restriction on the firm's objective.

The cross-partial in (24) also delivers a testable implication that requires no additional parameters. Under $\Pi_j' > 0$ and $\Pi_j'' \leq 0$, the first-order condition (18) implies

$$\frac{\partial d_{jt}^*}{\partial F_{jt}} < 0, \tag{25}$$

so that, holding the cost of restoration fixed, restoration is faster in CONJs where interruptions are more frequent. Three qualifications apply. The prediction is conditional on the cost of restoration: storms can impede access and occupy crews, raising ψ_{0jt} where interruptions are frequent, so the unconditional relationship need not be negative. The restoration cost does not depend on the number of interruptions; with crew congestion, the sign of (25) is ambiguous, and the check becomes a joint test of the payoff model and of that assumption. And measuring restoration time as the ratio DEC/FEC would create a mechanical negative relationship with FEC through measurement error in FEC. I therefore measure restoration time from event-level durations and test the prediction conditional on restoration-cost shifters, including weather in the month. Whether the data bear this out is an informative check on the joint framework.

6 Identification

This section discusses the sources of variation that identify the parameters of the model. The argument proceeds in two blocks. The demand-side parameters are identified from household consumption and regime choices, using variation in tariffs and in local service quality. The supply-side parameters are identified from the firm's observed operating decisions, taking the demand estimates as known. This sequencing is substantive rather than merely computational: the contribution of the supply-side analysis rests on bringing an independently estimated measure of household welfare to bear on the firm's problem, so that the firm's objective can be assessed against an external benchmark rather than rationalized by construction.

Throughout, I distinguish between parameters identified from within-regime behavior —

the intensive margin — and those identified from movements across regimes — the extensive margin. This distinction does substantial work, because it allows the direct effect of reliability on utility to be pinned down separately from its effect on the incentive to evade payment.

Demand

The demand-side parameter vector is

$$\theta^D = \big(\underbrace{a_0,\ a_c,\ a_X}_{\text{taste shifters}},\quad \underbrace{\eta,\ \gamma}_{\text{utility curvature}},\quad \underbrace{\sigma_\nu,\ \rho}_{\text{unobservables}},\quad \underbrace{\phi_1,\ \phi_2,\ \phi_3,\ \phi_4}_{\text{cost of evasion}},\quad \underbrace{\kappa_0,\ \kappa_1,\ \kappa_X}_{\text{capacity of the connection}},\quad \underbrace{m_0,\ m_X,\ s}_{\text{capacity distribution}} \big),$$

together with the market-level unobservables $\{\Delta\xi_{jt}, \chi_{jt}\}$, which are recovered by inversion. The parameters governing the technical capacity $\bar{q}_i$ in (7) and the population capacity distribution are estimated jointly with the rest of θ^D, not in a preliminary step (Section 7). The reference consumption level q_0 is not identified and is fixed (Section 4).

Normalizations. Two normalizations are required and neither is restrictive. First, because only the difference $\omega_{it} \equiv e^F_{it} - e^S_{it}$ enters the regime comparison, the individual variances of the two regime-specific shocks are not separately identified, and I set $\text{Var}(\omega_{it}) = 1$ to fix the scale of the discrete choice. Adding any common component to both shocks leaves every equation of the model unchanged, so no economic content is lost. Second, only the difference $\Delta\xi_{jt} = \xi^S_{jt} - \xi^F_{jt}$ enters the regime comparison, so the level of the component common to both regime utilities is normalized to zero. The difference $\Delta\xi_{jt}$ is not normalized: it shifts the local formal share and is recovered separately for each market.

Curvature of preferences. The parameters η and γ are identified from the formal-market demand equation (5), which is linear in logs with coefficients $-\eta$ on the log tariff and $\eta\gamma$ on log reliability. Both regressors are endogenous with respect to the market-level taste shock χ_{jt}. The tariff reflects the firm's regulated cost recovery, which itself depends on network conditions, and reliability is chosen by the firm in ways that may correlate with unobserved local demand.

Tariff variation is instrumented using the national tariff flag interacted with utility-specific passthrough, together with the annual tariff readjustments of each concession. The flag is set centrally on the basis of aggregate hydrological conditions in the principal generating basins and applies simultaneously to all distributors. Flags are pass-through charges that fund generation costs, and I treat them as changing the price households pay but not the tariff the distributor retains per billed kilowatt-hour. The exclusion restriction is that generation-cost shocks affect household electricity decisions only through the tariff, and not through local service quality or local taste for electricity. This is credible because generation and distribution are separately regulated activities: a change in the national cost of generation does not alter the condition of a local feeder.

The two margins draw on different tariff variation. Monthly flag changes identify the intensive-margin response η. They are a weak source of variation for the extensive margin, because a household is unlikely to set up or remove an illegal connection in response to a one-month change in the flag. The extensive-margin price response is identified instead from persistent tariff changes: the water-scarcity flag, in force from September 2021 to April 2022, and the annual readjustments, which vary by company and year. This mirrors Resende et al. (2025), who identify the extensive margin from a permanent tariff increase.

Reliability is instrumented using local weather. Storms, high winds, and heavy precipitation cause interruptions that are unrelated to household preferences, and their timing is plausibly exogenous conditional on the seasonal patterns absorbed by fixed effects. Realized reliability in the consumption equation is instrumented with storm-related weather in the month, and expected reliability in the regime choice with the same measures accumulated over the preceding twelve months, as in Section 3.1. The threat to this restriction is that weather also enters demand directly, since temperature drives cooling and heating load. I therefore restrict the excluded instruments to storm-related measures and control flexibly for temperature, so that the identifying variation comes from precipitation and wind events rather than from thermal conditions.

A second concern is the joint relevance of the two instruments. The tariff flag responds to reservoir levels, which are themselves a function of accumulated rainfall. Identification does not require the instruments to be independent of one another; it requires valid exclusion restrictions and a first stage of full rank after the controls. Correlation between the instruments short of a rank failure weakens identification without eliminating it. Two features limit the correlation: the flag is determined by hydrological conditions in the generating basins of the Southeast and Center-West, whereas local reliability responds to storm events in the service territory of each distributor, and the sample includes utilities located outside those basins. I report first-stage statistics that account for the joint endogeneity of the two regressors.

The exclusion restriction also has a dynamic dimension. A tariff change can alter evasion and, through the firm's response, subsequent reliability. Because reliability enters the model directly and is instrumented, this pathway is part of the model rather than a violation, provided the relevant reliability measure is included: the regime choice depends on reliability accumulated over the preceding twelve months, which absorbs the effect of past tariff-induced changes in quality.

Dispersion and selection. The standard deviation σ_ν of the household taste shock is identified from the dispersion of log consumption among formal households within markets, conditional on capacity. Because formal consumption is observed only after selection, the relevant object is the model-implied variance of log consumption among formal households, which depends on σ_ν and on the truncation induced by selection; the estimator matches this conditional variance directly. Because the coefficient on the taste shock in the log demand equation is η, the scale of ν is recovered once η is known.

The correlation $\rho = \text{corr}(\nu_{it}, \omega_{it})$ is identified from selection. Consumption in the formal market is observed only for households that chose the formal regime, so the conditional mean of log consumption differs from its unconditional counterpart by an amount that depends on ρ. Identification comes from how the selected distribution of consumption shifts as the instruments move households across the participation margin. When the tariff falls or reliability improves and the formal market expands, the composition of formal households changes. The direction of that change reflects two channels: the structural channel, since ν enters the formal and capped-evasion payoffs differently, and the correlation ρ. If $\rho < 0$, households with unobservedly high taste for electricity are also those least attached to the formal market, so the correlation channel draws entrants from the upper part of the ν distribution and raises average formal consumption; $\rho > 0$ works in the opposite direction. The estimator computes both channels, so what identifies ρ is not the level of consumption but how its selected distribution co-moves with the size of the formal market as the instruments shift it.

This argument makes clear that ρ places a more demanding requirement on the instruments than η and γ do. It is not enough that the instruments shift tariffs and reliability; they must shift a sufficient number of households across the regime boundary to trace out the composition effect. An instrument may be strong for the intensive margin and weak for the extensive one. Identification of ρ also leans on the joint normality of the unobservables. Variation that shifts participation without shifting formal consumption directly helps: the recovered $\Delta\xi_{jt}$, the evasion-cost terms, and the capacity ceiling $\bar{q}_i$ all enter the regime choice but not formal demand. Even so, the selection correction is identified partly by functional form, and I report estimates of the remaining parameters with ρ fixed at zero.

The cost of evading payment. The parameters ϕ_1 through ϕ_4 govern the term $g(r_{jt}, \bar{q}_i, p_{ut})$ capturing the expected cost of an illegal connection. The identification of the reliability terms rests on a separation between the two channels through which reliability enters the evasion regime. Reliability raises the direct utility of stolen consumption through $u(\bar{q}_i; r_{jt})$, and this channel is fully determined by η and γ, which are identified from the intensive margin of formal consumption. Any response of evasion rates to reliability beyond what these parameters predict must therefore be attributed to $g(\cdot)$. In this sense the intensive margin disciplines the direct utility channel and frees the extensive margin to identify the enforcement channel.

The reliability coefficients are then separated by the shape of that residual response. The linear term ϕ_1 is identified from the average slope of evasion with respect to instrumented reliability, and the quadratic term ϕ_2 from curvature in that relationship. Identification of ϕ_2 requires genuine dispersion in reliability. Where availability is confined to a narrow band, the linear and quadratic terms are collinear and the curvature parameter is not separately identified; I therefore report specifications with and without the quadratic term. The interaction ϕ_3 cannot be identified from a household-level gradient, because neither household-level evasion nor the capacity of evading households is observed. It is identified indirectly: from how the reliability gradient of market evasion shares varies across markets with different capacity distributions,

and from how reliability shifts the selected capacity distribution of registered households in moment (28). The repayment share ϕ_4 is identified from how the price gradient of evasion varies with the ratio of evasion capacity to formal consumption: by (10), a tariff increase raises the value of evading relative to paying by $q_{it}^F - \phi_4 \bar{q}_i$ per unit of price.

Market-level unobservables. The terms $\Delta\xi_{jt}$ and χ_{jt} are not estimated as free parameters. For any candidate θ^D, $\Delta\xi_{jt}$ is recovered so that the model matches the observed formal share, the ratio of registered residential units to dwellings (Section B.4), and χ_{jt} so that it matches mean log consumption among formal households. Non-technical losses enter separately, through an energy moment. The identifying content of the recovered unobservables lies in their orthogonality to the instruments: the requirement that unobserved local determinants of evasion — the prevalence of existing illegal connections, local enforcement intensity, prevailing norms regarding payment — and of local demand are uncorrelated with the tariff instruments and with local storm events provides the market-level moment conditions.

Supply

The supply-side parameter vector is

$$\theta^S = \big(\underbrace{\theta_F,\ h}_{\text{failure technology}},\ \underbrace{RC,\ \theta_c,\ \sigma_\varepsilon}_{\text{replacement}},\ \underbrace{\theta_\psi,\ \psi_1}_{\text{repair effort}},\ \underbrace{\lambda^f}_{\text{weight on surplus}} \big),$$

where λ^f is permitted to vary across the six distributors. Throughout, the demand-side objects $D_j(\cdot)$, $Q_j^S(\cdot)$, and $CS_j(\cdot)$ and their derivatives are treated as known.

Normalizations. Because variable profit and replacement cost are both denominated in currency, the coefficient on profit in the firm's payoff is set to unity. This is what gives λ^f its interpretation as the value of a unit of consumer surplus relative to a unit of profit, and it is also what makes the scale of the unobserved cost shocks, σ_ε, separately identifiable rather than absorbed into the remaining parameters. The discount factor is not identified from choice data alone, a standard result in the dynamic discrete choice literature, and is fixed.

Because the welfare value of reliability depends on the reference level q_0, so does λ^f. The replacement and restoration decisions identify the product $\lambda^f \, \partial CS_j / \partial r$, the value the firm attaches to reliability through households' surplus. A higher value of reliability to households, implied by a lower q_0, implies a lower estimated weight.

Failure technology. The age profile of transformer-caused interruptions, θ_F, is identified from comparisons among transformers of different ages serving the same feeder in the same month. Conditioning at this level absorbs weather, vegetation conditions, upstream faults, and any feeder-level shock common to the units it contains, so that the remaining variation is driven by differences in the age of the equipment itself.

Three features of this step deserve emphasis. First, θ_F is estimated outside the firm's choice

problem, from realized interruption counts rather than from replacement behavior. This is precisely what distinguishes the present framework from a standard replacement model, in which every consequence of equipment age is folded into a single cost function and no separate weight on service quality can be recovered. Its separation from replacement behavior rests on an assumption, however: conditional on age and capacity, units within a feeder do not differ in condition or in loading in ways the firm observes. If the firm replaces units in persistently poor condition early, the surviving stock at older ages is positively selected and the estimated age profile is too flat; if new or upgraded units are placed on unusually heavily loaded branches, the profile at young ages is too high. Second, because equipment ages deterministically, age and installation cohort are collinear in the surviving stock, and θ_F therefore confounds physical degradation with differences in the quality of successive vintages. The firm faces the same composite when deciding whether to replace a unit, so this is harmless for the replacement decision in the data. It matters for counterfactuals, in which units installed today age through future states: using the estimated profile there assumes that new vintages age as the current stock did. Third, the failure hazard $h(x)$ is distinct from $F^{\text{tr}}(x)$ and is not delivered by the interruption-count regression. It requires records of destructive failures; where these are not available, $h(x)$ is parameterized and identified jointly with the replacement model through the mixture (20), which relies on functional form.

The weight on consumer surplus. The central parameter is λ^f, and its identification is the crux of the supply-side argument. Replacing a transformer or restoring service faster raises reliability, and the firm values the gain at $\Pi_j'(r)$ in (18), the sum of a profit component, $(p-w)\,\partial D_j/\partial r - w\,\partial Q_j^S/\partial r$, and a surplus component, $\lambda^f\,\partial CS_j/\partial r$. If all households paid, the two components would move nearly in proportion across CONJs, and any observed behavior could be rationalized by an appropriate cost function with λ^f set to zero.

What breaks this proportionality is payment evasion. Evading households contribute to $\partial CS_j/\partial r$ but not to billed revenue, and their consumption rises with reliability at cost w to the firm. At the same time, better reliability draws some households into the formal regime, which raises $\partial D_j/\partial r$. The demand model computes the net effect; the relevant fact is that the ratio of the surplus component to the profit component varies across CONJs with evasion, holding physical load and equipment characteristics fixed. A firm with $\lambda^f = 0$ responds only to the profit component; a firm with $\lambda^f > 0$ replaces and restores more than the profit component justifies where that ratio is high. Because the replacement payoff (23) values reliability at the CONJ's marginal payoff, identification comes from variation across CONJs, not from variation across transformers within a CONJ.

Costs of replacement. Given the failure technology and the demand-side objects, the cost parameters are identified from the age profile of observed replacement. The fixed cost RC governs the overall frequency with which units are replaced, and is therefore identified from the level of the replacement hazard; the maintenance cost parameters θ_c govern how sharply the

incentive to replace rises with age, and are identified from the slope of that hazard. Because the reliability channel is pinned down independently by θ_F, the residual age dependence in replacement behavior is attributed to costs.

It follows that θ_c is a residual object and absorbs any misspecification of the reliability mapping or of the vintage composition of the equipment stock. A large estimated maintenance cost should therefore not be read as direct evidence about the expense of servicing old units. Nor is λ^f insulated from such misspecification: because the reliability gain $F^{\text{tr}}(x) - F^{\text{tr}}(0)$ multiplies the CONJ-specific marginal payoff, including its λ^f component, an error in the age profile, or an age-dependent cost that varies with evasion, would affect the cross-CONJ contrasts that identify λ^f. The separation between the two parameters rests on θ_c being common across CONJs with different levels of evasion, conditional on the cost shifters.

The dispersion of the unobserved cost shocks, σ_ε, is identified from the extent to which replacement decisions are imperfectly predicted by the observed state. Because benefits and costs vary with unit capacity and CONJ conditions, even a deterministic policy would replace some units and retain others of the same age. The overlap that identifies σ_ε is therefore the overlap between replaced and retained units conditional on the full observed state — age, capacity, and Ω — not the overlap of their unconditional age distributions.

The cost of repair effort. The parameters θ_ψ and ψ_1 governing the cost of maintaining crew capacity are identified from the first-order condition of the static problem. The level parameter θ_ψ is identified from the average restoration time observed across regions, and the curvature parameter ψ_1 from how restoration time responds to variation in the marginal benefit of speed — that is, from the fact that regions differing in interruption frequency and in the composition of the households they serve are restored at systematically different speeds.

Restoration time is a choice variable and is therefore correlated with unobserved determinants of the cost of speed, such as crew availability or the difficulty of physical access. Identification requires an instrument that shifts the marginal benefit of faster restoration without shifting its cost. I use the national tariff flag interacted with baseline exposure to evasion. Tariff movements change the price households pay and, through regime choices, the composition of paying and non-paying households, which changes the ratio of profit to surplus at stake in faster restoration; they also change $\partial D_j/\partial r$ directly, since formal consumption falls with the price. Two qualifications apply. Because regime choices respond to persistent rather than monthly price changes, the composition channel operates through persistent flag episodes, and monthly flags act mainly through the direct channel. And because flags respond to hydrological conditions that may coincide with local weather, which affects restoration costs, I control for local weather in the cost shifters; the instrument remains valid only if flag-related weather does not change crew costs differentially across CONJs with different baseline evasion.

A distinct concern is that observed restoration times reflect the composition of faults as well as the effort devoted to resolving them, since some equipment failures require longer repairs than others regardless of staffing. I therefore construct restoration times from event-level durations

within cause category and rely on within-category comparisons, so that the identifying variation reflects how quickly the firm resolves a given type of fault rather than which types of fault it experiences.

7 Estimation Method

Estimation proceeds in two blocks. The demand-side parameters are estimated jointly by simulated method of moments. The supply-side parameters are then estimated taking the demand estimates as known. The sequencing across blocks is substantive rather than computational: household preferences are primitives with respect to the firm's problem, and the contribution of the supply-side analysis rests on confronting the firm's observed decisions with an independently estimated measure of what reliability is worth to the households it serves. Within each block, all parameters are estimated simultaneously.

Two features of the setting shape the design of the demand estimator and are worth stating at the outset. First, formal-market consumption is observed only for households that select into the formal regime, so instrumental-variables estimation of the demand equation in isolation is inconsistent whenever selection on the taste shock is present, including through ρ. Because ρ is itself a parameter of interest, the instrumental-variables logic must enter through the choice of moment conditions rather than as a separate estimation step. Second, and more fundamentally, installed capacity is recorded only for registered consumer units. The distribution of capacity in the population is therefore not observed, and it cannot be treated as data: the counterfactuals require integrating over households that are not currently formal, and precisely which households those are changes with the policy under evaluation. The estimator must recover that distribution rather than approximate it by the distribution among the formal.

Demand

Parameters and unobserved distributions

The demand block estimates

$$\theta^D = \big(\underbrace{a_0,\, a_c,\, a_X,\, \eta,\, \gamma}_{\text{preferences}},\ \underbrace{\sigma_\nu,\, \rho}_{\text{unobservables and scale}},\ \underbrace{\phi_1, \phi_2, \phi_3, \phi_4}_{\text{cost of evasion}},\ \underbrace{\kappa_0, \kappa_1, \kappa_X}_{\text{capacity of the connection}},\ \underbrace{m_0, m_X, s}_{\text{capacity distribution}} \big),$$

together with the market-level unobservables $\{\Delta\xi_{jt}, \chi_{jt}\}$, which are recovered by inversion rather than estimated as free parameters. The reference level q_0 is fixed.

The last two groups replace what would otherwise be a preliminary regression. The capacity available under an illegal connection is given by (7). That equation describes a technical ceiling anchored to the physical connection, not consumption evaluated at a counterfactual price: κ_1 is the elasticity of the ceiling with respect to capacity and $\exp(\kappa_0)$ is a utilization rate. Because $\bar{q}_i$ enters the indirect utility of evasion and therefore the regime choice, estimating it as a first stage and holding it fixed would ignore the dependence of the estimating equations on κ; it is estimated inside the criterion instead.

The population distribution of installed capacity is specified as

$$\log c \mid j \;\sim\; \mathcal{N}\big(\mu_j,\, s^2\big), \qquad \mu_j \;=\; m_0 + m_X' X_j, \tag{26}$$

with X_j constructed from Census variables available at the census-tract level and apportioned to the CONJ by the dasymetric weighting described in Section B.4. Parameterizing μ_j rather than allowing free region effects is necessary for identification with the moments I use: with free region-specific locations, both μ_j and $\Delta\xi_{jt}$ would shift the local formal share, and the pooled shape moments below could not separate them. Regional distributions of registered capacity could in principle separate the two; the parameterization trades that flexibility for moments that remain informative where few formal households are observed. Tying μ_j to observable characteristics lets $\Delta\xi_{jt}$ be identified from the level of the formal share and μ_j from the shape of the capacity distribution observed among formal households.

Recovering the population distribution of capacity

Installed capacity is observed only conditional on formal participation, so the empirical distribution in the consumer-unit records is the population distribution tilted by the probability of being formal:

$$f_j^F(c) \;=\; \frac{\Pr\big(\text{formal} \mid c, X_j; \theta^D\big)\; f_j(c)}{\Pr\big(\text{formal} \mid X_j; \theta^D\big)}. \tag{27}$$

Capacity enters the regime comparison through two channels. It raises the value of both regimes through the taste shifter, though not in the same proportion, since the formal value scales partly with A_{it}^{η} and the evasion value with A_{it}. And it relaxes the constraint $\bar{q}_i$, which binds only under evasion. Both channels tilt the distribution of capacity among formal households.

Equation (27) could in principle be inverted nonparametrically, by reweighting each observed formal household by the reciprocal of its model-implied participation probability. That approach fails in exactly the region that matters. Where participation probabilities are low, few or no formal households are observed, and no reweighting of an empty cell can generate the missing mass. Since it is precisely the households with low participation probability that account for non-technical losses, the nonparametric inversion is least reliable where the model most requires it. The parametric specification (26) avoids this by extrapolating on functional form in the sparse region, and the resulting distribution is disciplined by an aggregate energy moment rather than by the sparse cells themselves.

Accordingly, the estimator matches the first two moments of the model-implied selected distribution to their observed counterparts, with the first moment interacted with the local characteristics that enter μ_j. Writing $W_j = (1, X_j')'$ and $E_j^{\theta}[\cdot] \equiv \int(\cdot)\, \Pr\big(\text{formal} \mid c, X_j\big)\, dF_j(c) \big/ \int \Pr\big(\text{formal} \mid c, X_j\big)\, dF_j(c)$ for the model-implied mean among formal households,

$$g_3(\theta^D) \;=\; \frac{1}{M} \sum_{j,t} \begin{pmatrix} W_j \Big(\widehat{\mathbb{E}}\big[\log c \mid \text{formal}, j\big] - E_j^{\theta}\big[\log c\big] \Big) \\ \widehat{\mathbb{E}}\big[(\log c)^2 \mid \text{formal}, j\big] - E_j^{\theta}\big[(\log c)^2\big] \end{pmatrix}, \tag{28}$$

where the empirical terms are computed directly from the consumer-unit records. The first block, interacted with $(1, X_j')$, identifies (m_0, m_X); the second identifies s.

Market-level unobservables

Because the formal share and mean formal consumption are observed in every market, with the number of households outside the formal market measured as the difference between dwellings and registered consumer units, the market-level unobservables are recovered by inversion rather than estimated. For any candidate θ^D, $(\Delta\xi_{jt}, \chi_{jt})$ solve

$$\begin{aligned} \underbrace{\iint \Pr\big(\text{formal} \mid c, \nu, X_j; \theta^D, \Delta\xi_{jt}, \chi_{jt}\big)\, d\Phi(\nu/\sigma_\nu)\, dF_j(c)}_{s_{jt}^{\text{sim}}(\theta^D, \Delta\xi_{jt}, \chi_{jt})} &= \underbrace{\frac{N_{jt}^{\text{UCBT}}}{N_j^{\text{CNEFE}}}}_{s_{jt}^{\text{obs}}}, \\ m_{jt}^{\text{sim}}\big(\theta^D, \Delta\xi_{jt}, \chi_{jt}\big) &= \overline{\log q}_{jt}^{F}, \end{aligned} \tag{29}$$

where N_{jt}^{UCBT} is the number of registered residential units, N_j^{CNEFE} the number of dwellings (Section B.4), $\overline{\log q}_{jt}^{F}$ mean log consumption among formal households, and m_{jt}^{sim} its model counterpart. For given χ_{jt}, the predicted formal share is strictly monotone in $\Delta\xi_{jt}$; for given $\Delta\xi_{jt}$, mean formal log consumption increases with χ_{jt}. I solve the two conditions by nested one-dimensional updates, separately for each region and period, and verify convergence numerically. Note that the population distribution F_j appears inside (29): the inversion is over the whole population of households, not the observed formal sample, which is what makes the recovered $\Delta\xi_{jt}$ interpretable as a local shifter of evasion rather than a composite of that shifter and the local capacity composition.

Moment conditions

Four sets of moments identify θ^D.

The first uses within-market variation in formal consumption. Let $m_{it}(\theta^D) \equiv \widehat{\mathbb{E}}\big[\log q \mid \text{formal}, c_i, X_j, p_{ut}, r_{jt}, \bar{r}_{jt}, \Delta\xi_{jt}, \chi_{jt}; \theta^D\big]$ denote the model-implied mean of log consumption for a household with capacity c_i, conditional on formal participation, and $V_{it}(\theta^D)$ the corresponding variance. The moments are

$$g_1(\theta^D) = \frac{1}{N_F} \sum_{i \in F} \begin{pmatrix} \log c_i \big[\log q_{it} - m_{it}(\theta^D)\big] \\ \big[\log q_{it} - m_{it}(\theta^D)\big]^2 - V_{it}(\theta^D) \end{pmatrix}. \tag{30}$$

The conditional moments are computed by simulation as participation-weighted averages across draws of the taste shock; for the mean,

$$\widehat{\mathbb{E}}\big[\log q \mid \text{formal}, \cdot\big] = \frac{\sum_s \Pr\big(\text{formal} \mid \nu^{(s)}, \cdot\big)\, \log q\big(\nu^{(s)}, \cdot\big)}{\sum_s \Pr\big(\text{formal} \mid \nu^{(s)}, \cdot\big)},$$

evaluated at the household's own observed capacity c_i. Weighting by participation integrates the taste shock over its distribution conditional on formal participation, which is valid because ν_{it} is independent of tariffs, reliability, and the market-level unobservables. Because χ_{jt} is inverted

to match mean formal log consumption in each market, the residual has mean zero within each market by construction, and (30) uses only within-market variation: the first element identifies a_c from how formal consumption rises with capacity, net of selection on capacity, and the second identifies σ_ν from the selected dispersion of consumption.

The second imposes orthogonality between the instruments and the two recovered market-level unobservables:

$$g_2(\theta^D) \;=\; \frac{1}{M}\sum_{j,t} Z_{jt} \otimes \begin{pmatrix} \chi_{jt}(\theta^D) \\ \Delta\xi_{jt}(\theta^D) \end{pmatrix}, \tag{31}$$

where Z_{jt} contains the excluded instruments — the tariff instruments and storm-related weather, contemporaneous and accumulated over twelve months — and the exogenous covariates X_j. Orthogonality of χ_{jt} identifies the parameters of formal demand, (a_0, a_X, η, γ), in the manner of an instrumental-variables regression of market mean log consumption corrected for selection; orthogonality of $\Delta\xi_{jt}$ identifies the parameters of the regime choice. The selection correlation ρ enters both: if it is misspecified, the selection correction in the market mean is wrong, and $\chi_{jt}(\theta^D)$ absorbs the error, which then co-moves with the instruments that shift participation.

The third is the shape moment (28).

The fourth matches measured non-technical losses, and is what identifies the capacity of illegal connections:

$$g_4(\theta^D) \;=\; \frac{1}{M}\sum_{j,t} Z_{jt}\left[\mathrm{NTL}_{jt} \;-\; N_j^{\mathrm{CNEFE}} \iint \Pr\left(\text{steal} \mid c, \nu, X_j\right) \bar{q}(c)\, d\Phi(\nu/\sigma_\nu)\, dF_j(c)\right]. \tag{32}$$

Moments (29) and (32) use distinct information and identify distinct objects. The inversion matches the *number* of households consuming outside the formal market; the energy moment matches the *quantity* they consume. Their ratio is the average consumption of an evading household, which is the empirical counterpart of $\mathbb{E}[\bar{q} \mid \text{steal}, j]$ and therefore what pins down κ. Note that the evasion probabilities weighting $\bar{q}(c)$ in (32) come from the model itself, so the estimator does not require the auxiliary assumption that evading households are a random draw from the local population; the selection into evasion that the model predicts is the selection used in constructing the moment.

The estimator solves

$$\hat{\theta}^D \;=\; \arg\min_{\theta^D}\; g(\theta^D)'\, W\, g(\theta^D), \qquad g \;=\; \left(g_1', g_2', g_3', g_4'\right)',$$

using the two-step weighting matrix, with the first step weighted by the inverse of the instrument cross-product.

Computation

Two features keep the problem tractable. Because the regime-specific shocks enter the choice comparison additively, their difference ω_{it} is integrated analytically and only the taste shock requires simulation: conditional on a draw of ν_{it}, the participation probability takes the probit form in (12). And because $\Delta\xi_{jt}$ and χ_{jt} are inverted rather than estimated, the dimension of

the optimization does not grow with the number of regions.

For a candidate θ^D, the objective is evaluated as follows.

1. Construct F_j from (26) and draw capacities $\{c_j^{(m)}\}$ for the population integrals; the household moments use observed c_i instead.

2. Form $\bar{q}$ from (7) and evaluate v^F and v^S at each simulation draw, at expected reliability $\bar{r}_{jt}$ for the regime choice and realized reliability r_{jt} for consumption.

3. Solve the inversion (29) for $(\Delta\xi_{jt}, \chi_{jt})$ in each region and period.

4. Form g_1 through g_4 and evaluate the criterion.

Simulation uses Halton sequences held fixed across parameter evaluations, so that the objective is continuous in θ^D; the draws are redrawn across bootstrap replications, as described below. Because the number of draws enters the bias of the simulated moments at order $1/S$, I verify that the estimates are stable when the number of draws is doubled.

As a diagnostic, the estimated F_j implies a distribution of capacity among evading households that can be compared with the distribution among formal households in the same region. The model does not fix the sign of the difference: higher capacity relaxes the evasion ceiling, which favors evasion, but it also raises the value of both regimes through the taste shifter, and the reliability-dependent evasion cost interacts with capacity through ϕ_3. I report the comparison as a description of the recovered selection on capacity rather than as a test of the model.

Supply

The supply-side parameters are estimated taking $\hat{\theta}^D$ as given. The demand model is simulated at the observed reliability of each household to obtain $D_j(r)$, $Q_j^S(r)$, $CS_j(r)$, and their derivatives, which are aggregated to the CONJ level; the number of consumer units served by each transformer, N_k, comes from the observed assignment of consumer units to serving equipment. These objects enter the firm's problem as data. Estimation then proceeds in two steps.

First step: reliability technology

The age profile of transformer-caused interruptions is estimated by Poisson regression of transformer-attributed interruption counts on age bins, with feeder–month fixed effects and controls for equipment capacity. Conditioning at the feeder–month level absorbs weather, vegetation conditions, and upstream faults, so the identifying variation comes from comparisons among differently-aged transformers operating under common local conditions. This step uses realized interruptions rather than replacement decisions, which is what allows the reliability channel to be separated from the cost channel in the second step, under the assumption on unit condition stated in Section 6. The failure hazard entering the transition rule is estimated separately, from destructive failures; if these cannot be distinguished from other replacements, $h(x)$ is parameterized and estimated jointly with the replacement model through (20). Restoration times

are computed from event-level durations within cause category, so that the effort measure reflects how quickly the firm resolves a given type of fault rather than which types of fault it experiences.

Second step: replacement and effort

Two considerations shape the design of the second step.

First, replacement is a renewal action: after a replacement the unit enters the next period at age one, whatever its age before, and the flow payoff from replacing does not depend on that age. This delivers one-period finite dependence in the sense of Arcidiacono and Miller (2011). The continuation values of the two branches do not cancel directly, because keeping the unit leads to age zero after a failure or to age $x+1$ otherwise, while replacing leads to age one. They cancel after one further replacement. Let $v(x, y; \Omega)$ denote the choice-specific value, net of the shock. Using the representation $V(x', \Omega') = v(x', 1; \Omega') - \sigma_\varepsilon \ln P(1 \mid x', \Omega') + \sigma_\varepsilon \gamma_E$, where γ_E is Euler's constant, and the fact that neither $u(x', 1; \Omega')$ nor the continuation after a replacement depends on x', the choice-specific value difference is

$$\begin{aligned} v(x,1;\Omega) - v(x,0;\Omega) \;&=\; u(x,1;\Omega) - u(x,0;\Omega) \\ &+\; \delta\,\sigma_\varepsilon\,\mathbb{E}_{\Omega'|\Omega}\Big[\, h(x)\ln P(1 \mid 0,\Omega') + \big(1-h(x)\big)\ln P\big(1 \mid \min\{x+1,\bar{x}\},\Omega'\big) - \ln P(1 \mid 1,\Omega')\,\Big]. \end{aligned} \tag{33}$$

The correction term depends only on replacement probabilities at the ages the unit can reach next period, including the young ages zero and one, and on the transition of Ω. Following Hotz and Miller (1993), conditional choice probabilities are estimated in a preliminary flexible logit of the replacement indicator on age bins interacted with the CONJ-level state, smoothing across CONJs at young ages, where replacements are rare. When forced replacements cannot be separated from chosen ones, this preliminary step is fit to the mixture (20), so that it recovers the probability of a chosen replacement. Because Ω follows an exogenous Markov process, the expectation over Ω' is replaced by its sample counterpart, the observed state in the following year. With δ and $h(\cdot)$ given, the correction enters the second step, after division by σ_ε, as a constructed offset containing no unknown parameters, and the dynamic program is never solved. When $h(\cdot)$ is instead parameterized and estimated jointly, the offset depends on its parameters but on none of $\big(RC, \theta_c, \lambda^f, \sigma_\varepsilon\big)$. Conditional on the offset, the probability of a chosen replacement is a binary logit whose index is linear in $1/\sigma_\varepsilon$, which multiplies variable profit, and in $\big(RC, \theta_c, \lambda^f\big)/\sigma_\varepsilon$; the replacement moments use this probability through the mixture (20) when forced replacements cannot be separated.

Second, the weight on consumer surplus appears in both the replacement decision and the repair-effort decision, so the two margins are estimated jointly under the restriction that this weight is common:

$$\hat{\theta}^S \;=\; \arg\min_{\theta^S} \begin{pmatrix} g^{\text{repl}}(\theta^S) \\ g^{\text{eff}}(\theta^S) \end{pmatrix}' W \begin{pmatrix} g^{\text{repl}}(\theta^S) \\ g^{\text{eff}}(\theta^S) \end{pmatrix}.$$

The effort moments are those implied by the first-order condition (18), in which restoration time is instrumented using the national tariff flag interacted with baseline exposure to evasion, subject to the qualifications in Section 6. Because the two margins differ in cost structure, in timing, and in the set of interruptions they affect, the common weight is a substantive restriction rather than an accounting identity, and the resulting overidentification test assesses directly whether a single weight rationalizes behavior on both.

The parameters RC, θ_c, σ_ε, θ_ψ, and λ^f are permitted to vary across the six distributors. Regional cost heterogeneity enters parametrically through observable characteristics rather than as free conjunto-level effects, which with two transition periods would introduce an incidental-parameters problem and would absorb the cross-sectional variation on which λ^f depends. The discount factor is not identified from choice data and is fixed; results are reported for a range of values.

Inference

Standard errors must account for four sources of variability: sampling error; simulation error; the propagation of parameters estimated in earlier steps, including the demand estimates and the reliability technology entering the supply block; and dependence within geographic units and within months. I compute standard errors by a block bootstrap that resamples CONJs with replacement and re-executes the entire procedure, demand and supply, at each replication. Resampling CONJs rather than feeders preserves the dependence among the feeders, transformers, and households of the same CONJ, and keeps together the monthly restoration moments and the annual replacement moments of each CONJ. Simulation draws are redrawn at each replication, since holding them fixed would treat simulation error as deterministic and understate the resulting variance. Each replication is initialized at the point estimates to reduce computational cost.

The CONJ-level bootstrap preserves serial dependence within CONJs but does not account for shocks common across CONJs in the same month. The tariff instrument is national in origin, so the effective number of independent price shocks is governed by the number of distinct tariff episodes rather than by the number of CONJs, and persistent flags and twelve-month reliability measures reduce the effective time-series variation further. I therefore also report a two-way bootstrap that resamples CONJs and months independently, and I report the number of distinct flag regimes and tariff readjustments in the sample. Analytical standard errors clustered along both the geographic and the temporal dimension, following Cameron et al. (2011), are reported for the final-stage moments as a complement: they capture common monthly shocks but treat earlier-stage estimates as known.

8 Results

This section presents the estimates of the household demand and payment-evasion model and of the distributors' reliability-provision problem. I first discuss the household estimates and then turn to the parameters governing transformer replacement, outage restoration, and the

weight firms place on consumer surplus.

8.1 Household Demand and Payment Evasion

Table 4: Demand Model Parameter Estimates

Parameter	Notation	Estimate	Std. Error
Panel A. Formal Electricity Demand			
Price elasticity	η	0.19	(0.03)
Reliability elasticity of formal demand	$\eta\gamma$	3.20	(0.90)
Standard deviation of electricity taste	σ_ν	0.55	(0.06)
Panel B. Consumption Capacity under Payment Evasion			
Installed-capacity elasticity of $\bar{q}_i$	κ_1	1.00	(0.10)
Panel C. Expected Cost of Payment Evasion			
Reliability	ϕ_1	-3.50	(1.90)
Reliability squared	ϕ_2	0.00	(0.50)
Reliability × illegal-consumption capacity	ϕ_3	0.00	(0.01)
Expected repayment share	ϕ_4	0.20	(0.08)
Panel D. Selection into the Formal Market			
Correlation between electricity taste and regime preference	ρ	-0.35	(0.09)
Variance of regime-choice shock, Var(ω_{it})		1 (normalized)	
Reference consumption, q_0		Fixed (see notes)	
Market-level unobservables $\Delta\xi_{jt}$, χ_{jt}		Recovered by inversion	

Notes: This table reports estimates of the structural parameters of the household demand and payment-evasion model. The price elasticity of electricity consumption conditional on formal participation is $-\eta$, while the elasticity with respect to reliability is $\eta\gamma$. The parameter κ_1 governs how the technical consumption ceiling under payment evasion varies with installed electrical capacity. The expected-cost function is

$$g(r_{jt}, \bar{q}_i, p_{ut}) = \phi_1 r_{jt} + \phi_2 r_{jt}^2 + \phi_3 r_{jt}\bar{q}_i - \phi_4 p_{ut}\bar{q}_i.$$

The parameter ρ captures correlation between the household taste shock for electricity and the unobserved component of the formal-participation decision. The variance of $\omega_{it} \equiv e_{it}^F - e_{it}^S$ is normalized to one. The market-level unobservables $\Delta\xi_{jt}$ and χ_{jt} are recovered by inversion. The reference consumption level q_0 is fixed, and welfare results are reported for alternative values. Additional intercepts, demographic shifters, and parameters governing the population distribution of installed capacity are omitted from the table for brevity.

Table 4 shows that electricity consumption among formal households is relatively price inelastic. The estimated conditional price elasticity is 0.19 (s.e. 0.03). Thus, a 10 percent increase in the regulated tariff reduces electricity consumption among households that remain formal by approximately 1.9 percent. This estimate is close to existing estimates of residential electricity demand in Brazil and, more importantly for the analysis here, implies that the intensive-margin response to prices is relatively limited. Changes in tariffs can nevertheless have larger effects on total billed demand because they also affect households' decision to remain in the formal market.

Reliability also affects electricity consumption directly. The estimated elasticity of formal consumption with respect to reliability is 3.20 (s.e. 0.90). Because reliability is measured as the fraction of hours in which electricity service is available, economically relevant changes in r_{jt}

are small in absolute terms. For example, a 0.1 percent increase in availability raises predicted formal consumption by approximately 0.32 percent. The estimates therefore imply that outages affect not only households' participation decisions but also their electricity use conditional on continuing to pay.

Installed electrical capacity is also strongly related to the amount a household can consume outside the formal market. The estimate $\kappa_1 = 1.00$ (s.e. 0.10) implies that a 10 percent increase in installed capacity raises the technical consumption ceiling $\bar{q}_i$ by approximately 10 percent. This result is consistent with the interpretation of $\bar{q}_i$ as a physical constraint on consumption under an illegal connection rather than as formal electricity demand evaluated at a zero price.

The estimates also provide evidence on the costs households associate with payment evasion. The linear reliability coefficient is -3.50 (s.e. 1.90), while the quadratic term 0.00 and the interaction with illegal-connection capacity 0.00 are close to zero. The point estimates therefore suggest that the relationship between local reliability and the attractiveness of payment evasion is approximately linear over the range observed in the data, with little evidence of substantial curvature or heterogeneity by connection size. The expected-repayment parameter is estimated at 0.20 (s.e. 0.08), implying that households behave as if approximately 20 percent of the tariff value of stolen electricity is eventually repaid in expectation.

Finally, the selection parameter is negative: $\rho = -0.35$ (s.e. 0.09). Thus, conditional on observed household and local characteristics, households with a higher unobserved taste for electricity are more likely to select into payment evasion. This selection is important for the counterfactual analysis because the households induced to move between the formal and informal regimes are not a random draw from the population.

8.2 Reliability Provision

The replacement estimates imply broadly similar capital costs across distributors. The average estimated transformer replacement cost is approximately R$ 24.8 thousand. The scale of the unobserved replacement-cost shock averages R$ 8.6 thousand, indicating substantial residual heterogeneity in replacement decisions even after conditioning on the observed state of the transformer and local network.

The restoration estimates display considerably more heterogeneity across distributors. For example, the estimated restoration-cost shifter is 0.49 for AME and -0.61 for ENEL-SP. Since $\theta_{\psi,f}$ enters the log of the restoration-cost level, these estimates imply that, conditional on the other cost shifters, AME's restoration-cost level is approximately 3 times that of ENEL-SP. The estimated curvature parameter is $\psi_1 = 1.00$ (s.e. 0.15). With the restoration technology $\psi(d) = \psi_0 d^{-\psi_1}$, cutting average restoration time in half therefore multiplies restoration costs by approximately 2.

Most importantly, the estimates reject the benchmark in which distributors' local reliability decisions depend only on variable profit. The estimated consumer-surplus weights range from approximately 0.17 to 0.33 across distributors, with an average of 0.23. Because the coefficient on variable profit is normalized to one, these estimates imply that distributors behave as if one

Table 5: Supply-Side Parameter Estimates

	AME	COELBA	CEMIG	ENEL-RJ	ENEL-SP	LIGHT
Panel A. Transformer Replacement Costs						
Replacement cost, RC_f	26.5	25.3	24.4	25.9	23.8	22.9
	(2.0)	(1.8)	(1.7)	(2.1)	(1.6)	(1.9)
Maintenance cost, $\theta_{c,f}$	0.46	0.44	0.39	0.45	0.40	0.38
	(0.08)	(0.07)	(0.06)	(0.08)	(0.06)	(0.07)
Scale of replacement-cost shock, $\sigma_{\varepsilon,f}$	9.2	8.8	8.3	9.0	8.1	8.2
	(1.2)	(1.1)	(1.0)	(1.2)	(0.9)	(1.0)
Panel B. Outage-Restoration Costs						
Restoration-cost level, $\theta_{\psi,f}$	0.49	0.18	-0.04	0.03	-0.61	-0.45
	(0.18)	(0.15)	(0.14)	(0.17)	(0.12)	(0.13)
Restoration-cost curvature, ψ_1				1.00		
				(0.15)		
Panel C. Firm Objective						
Weight on consumer surplus, λ_f	0.19	0.17	0.33	0.23	0.20	0.25
	(0.08)	(0.07)	(0.07)	(0.09)	(0.07)	(0.08)
Discount factor, δ	0.95 (fixed)					
Demand parameters	Taken from Table 4					
Reliability technology	Estimated in preliminary Poisson model					

Notes: This table reports estimates of the parameters governing distributors' transformer-replacement and outage-restoration decisions. RC_f is the fixed replacement cost in thousands of reais, $\theta_{c,f}$ governs the age-dependent cost of continuing to operate existing equipment, and $\sigma_{\varepsilon,f}$ is the scale of the type-I extreme-value replacement-cost shock. Its standard deviation is $\pi\sigma_{\varepsilon,f}/\sqrt{6}$. The parameter $\theta_{\psi,f}$ governs the level of outage-restoration costs, while ψ_1 governs their curvature. The coefficient on variable profit in the firm's objective is normalized to one, so λ_f measures the value placed on one real of consumer surplus relative to one real of variable profit. Standard errors are reported in parentheses and account for estimation of the demand parameters and reliability technology. The discount factor is fixed at 0.95.

real of consumer surplus is worth roughly R$ 0.23 of firm profit on average. The estimates are substantially below the planner benchmark $\lambda_f = 1$, however, implying that firms internalize only part of the value that improvements in reliability generate for consumers.

This wedge is central for the counterfactual analysis. Payment evasion reduces the private return to reliability because electricity consumed by evading households generates little or no contemporaneous billed revenue, while those households continue to benefit from improvements in service quality. The combination of $\lambda_f < 1$ and endogenous formal participation therefore creates scope for tariff and enforcement policies to affect welfare not only through household demand, but also by changing distributors' incentives to replace equipment and restore service.

9 Counterfactuals

The estimated model delivers the two objects required to evaluate policy: household preferences over consumption, reliability, and payment regime, and the weight each distributor places on consumer surplus when choosing how much reliability to supply. I use these estimates for two exercises. The first quantifies the externality that payment evasion imposes on households that continue to pay, through the firm's investment and restoration decisions. The second evaluates optimal uniform and region-specific electricity tariffs when payment behavior and reliability are jointly determined. For both exercises, I report aggregate outcomes and examine heterogeneity between low- and high-income areas.

Both exercises require solving for an equilibrium in which household and firm behavior are mutually consistent. Given a tariff p and an enforcement technology, an equilibrium in region j consists of a reliability level r_j, a restoration time d_j, and a replacement policy $\{y(x)\}$ such that households choose their payment regime and consumption optimally given (p, r_j), with expected and realized reliability coinciding in the stationary equilibrium. These choices determine billed demand $D_j(p, r_j)$, stolen consumption $Q_j^S(p, r_j)$, consumer surplus $CS_j(p, r_j)$, and the share of households evading payment. The firm chooses d_j from the first-order condition (18) and $\{y(x)\}$ from the dynamic problem (22), given these demand objects and the estimated supply parameters $\hat{\theta}^S$. Reliability must then be consistent with the resulting frequency and duration of outages:

$$r_j \;=\; 1 - \frac{F_j d_j}{T}, \qquad F_j \;=\; \sum_{k \in j} \frac{N_k}{N_j} F^{\mathrm{tr}}(x_k) \;+\; F_j^{\mathrm{net}}\big(Q_j^S\big), \tag{34}$$

where the age distribution $\{x_k\}$ is the stationary distribution induced by the replacement policy and $F_j^{\mathrm{net}}(\cdot)$ is the response of non-transformer faults to stolen consumption. Reliability affects payment evasion, while evasion changes both the return to providing reliability and, through stolen load, the incidence of network faults. The counterfactual equilibrium is the fixed point of this mapping. I solve it by iterating on r_j until convergence and computing the stationary age distribution at each iteration from the replacement policy and transition rule (19).

Because this feedback could generate multiple fixed points, I start the iteration from several initial values—observed reliability, full availability, and low availability—and check whether they converge to the same point. I also compare the baseline solution at observed tariffs

with observed evasion, reliability, and replacement rates. The exercises compare stationary equilibria; the observed transformer stock need not already have reached its stationary age distribution.

9.1 Measuring the Externality on Formal Consumers

Payment evasion transfers resources away from the distributor and toward households that do not pay for the energy they consume. It can also impose a cost on households that continue to pay: nonpayment weakens the distributor's incentive to maintain the shared network, lowering the reliability received by formal and evading households alike. The first counterfactual isolates this reliability externality.

I remove the evasion regime from households' choice sets, so all households participate in the formal market, while holding electricity tariffs fixed at their observed levels. Firms then re-optimize transformer replacement and outage restoration, and reliability adjusts to the new equilibrium. Let $\mathcal{F}_j$ denote the set of households in the formal market in the baseline, and let r_j^{obs} and r_j^0 denote equilibrium reliability in the observed economy and under no evasion, respectively. The externality borne by baseline-formal households in region j is

$$E_j \;=\; \sum_{i \in \mathcal{F}_j} \mathbb{E}\Big[\, v_i^F\big(p, r_j^0\big) - v_i^F\big(p, r_j^{\text{obs}}\big) \;\Big|\; a_i = 0 \text{ in the baseline} \Big], \tag{35}$$

where both terms evaluate money-metric surplus in the formal regime at the observed tariff. The expectation integrates over the distribution of tastes conditional on formal participation in the baseline. I keep the set of baseline-formal households and their payment regime fixed when evaluating the welfare change. This distinction matters because eliminating evasion also brings previously evading households into the formal market: including their welfare change would conflate the externality on existing paying households with the direct effect of prohibiting evasion. The measure in (35) captures only the reliability gain of households that would have paid in either environment. Its money-metric value depends on the reference consumption parameter q_0 in the demand model.

Two mechanisms can contribute to this externality. First, there is a *behavioral* channel: evading households benefit from reliability but do not generate billed revenue, weakening the private return to replacement and restoration when the firm values consumer surplus less than profits ($\lambda^f < 1$). Removing evasion expands billed demand, although it also removes an incentive to improve reliability in order to attract consumers into the formal market; the net response is determined by the estimated model. Second, there is a *physical* channel: stolen consumption can increase the frequency of network faults through $F_j^{\text{net}}(Q_j^S)$. To distinguish these mechanisms, an intermediate counterfactual r_j^A recomputes the firm's demand objects with all households paying but holds non-transformer fault frequency at its observed level. The change from r_j^{obs} to r_j^A measures the behavioral component; the remaining change to r_j^0 captures the physical component. The latter requires an estimate of how faults respond to stolen energy, and is sensitive to whether tariff-induced variation in stolen energy has effects on faults other

than through illegal consumption.

A separate channel works through regulated tariffs. Distribution tariffs recognize an allowed level of nontechnical losses, some of which is recovered from paying customers. If p^0 is the tariff with no allowance for these losses, the associated surplus change is

$$E_j^\tau \;=\; \sum_{i \in \mathcal{F}_j} \mathbb{E}\Big[\, v_i^F\big(p^0, r_j^0\big) - v_i^F\big(p, r_j^0\big) \;\Big|\; a_i = 0 \text{ in the baseline} \Big]. \tag{36}$$

This tariff component combines a transfer from paying consumers with the consumption distortion generated by the higher price. It is conceptually distinct from the reliability externality E_j. The tables below report the reliability externality, holding observed tariffs fixed, rather than adding a hypothetical change in the recognized-loss allowance.

Table 6 reports the aggregate counterfactual. Removing payment evasion reduces its measured rate from 12.3% to zero. Firms respond by increasing transformer replacement from 4.5% to 4.9% and reducing average restoration time from 9.8 to 8.8 hours. Together, these changes reduce the outage-frequency index from 100 to 88.2, an 11.8% decline. The reliability gain is larger in places where evasion is more prevalent: outage frequency falls by 23% in high-evasion areas. For households that were paying in the baseline, the average reliability gain is worth R$ 5.0 per month, or 3.4% of the average electricity bill. This is a welfare cost of evasion beyond the direct loss of distributors' revenue.

Table 6: Counterfactual I: Eliminating Payment Evasion

	Baseline	No Evasion	Change
Payment evasion (%)	12.3	0.0	−12.3 pp
Outage frequency (index)	100	88.2	−11.8%
Transformer replacement rate (%)	4.5	4.9	+8.9%
Average restoration time (hours)	9.8	8.8	−10.2%
Externality per formal household	—	—	R$ 5.0/month
Externality as % of average bill	—	—	3.4%

Notes: Payment evasion is eliminated while observed tariffs are held fixed. Firms re-optimize transformer replacement and outage restoration; reliability solves (34). Outage frequency is normalized to 100 in the baseline. Changes in replacement and restoration are percentage changes relative to baseline values. The externality is defined in (35) and is reported per household that was formal in the baseline, per month; the counterfactual does not include tariff reductions.

The average externality conceals substantial heterogeneity. To examine who bears the cost of evasion, I divide CONJs by local income and separately calculate outcomes for the bottom and top income quartiles. I continue to measure welfare only for households that were paying for electricity in the baseline. Table 7 shows that the feedback loop is considerably stronger in lower-income areas. Baseline evasion is 23% in the bottom quartile, versus 4.2% in the top quartile. Eliminating evasion reduces the outage-frequency index from 100 to 73 in the bottom quartile, compared with 100 to 97 in the top quartile. In the bottom quartile, the transformer replacement rate rises from 3.9% to 4.5%, and mean restoration time falls from 13.5 to 9.4

hours; changes in the highest-income quartile are much smaller.

These differences translate into a highly unequal externality. Households already paying in the bottom-income quartile gain R$ 21 per month from eliminating evasion, equivalent to approximately 20% of their average electricity bill. The corresponding gain in the top-income quartile is only R$ 1 per month (reported as approximately zero percent of the bill after rounding). Hence, households that continue to comply with the formal payment system in poorer areas bear a substantially larger reliability cost from their neighbors' evasion.

Table 7: Eliminating Payment Evasion: Heterogeneity by Income

	Bottom 25% Income			Top 25% Income		
	Baseline	No Evasion	Change	Baseline	No Evasion	Change
Payment evasion (%)	23.0	0.0	−23.0 pp	4.2	0.0	−4.2 pp
Outage frequency (index)	100	73	−27%	100	97	−3%
Transformer replacement rate (%)	3.9	4.5	+0.6 pp	5.3	5.3	0.0 pp
Average restoration time (hours)	13.5	9.4	−30.4%	6.6	6.6	0.0%
Externality (R$/month)	—	—	21.0	—	—	1.0
Externality (% of average bill)	—	—	20%	—	—	0%

Notes: The two groups are CONJs in the bottom and top quartiles of local income. The exercise holds observed tariffs fixed and removes the option to evade. Externalities are monthly gains per baseline-formal household in the corresponding income group; outage indices are separately normalized to 100 within each group. Changes are calculated from the displayed outcome levels: percentage-point changes are denoted "pp," and changes in outage frequency and restoration time are relative percentages. The top-quartile externality is R$ 1 per month; its share of the average electricity bill rounds to 0%.

9.2 Optimal Regional Tariffs

Tariffs are uniform within each distribution concession, although regions within the same concession differ in payment evasion, network reliability, and household income. I next ask what tariffs a planner would choose if it could set a separate price in each CONJ, while anticipating how households and the distribution company respond. Unlike the no-evasion counterfactual, this exercise leaves the estimated enforcement technology and the option to evade payment unchanged.

The planner maximizes total surplus, taking the firm's reliability choice $r_j(p_j)$ as an equilibrium response to tariffs. For a concession f with regions $\mathcal{J}_f$, the problem is

$$\begin{aligned} \max_{\{p_j\}_{j\in\mathcal{J}_f}} \quad & \sum_{j\in\mathcal{J}_f} \Big\{ \underbrace{p_j\big[D_j+\phi_4 Q_j^S\big] - w\big[D_j+Q_j^S\big] - C_j(r_j)}_{\pi_j(p_j,r_j)\text{: firm profit}} + CS_j \Big\} \\ \text{subject to} \quad & \sum_{j\in\mathcal{J}_f} \pi_j(p_j,r_j) \geq \bar{\pi}_f, \qquad r_j = r_j(p_j), \end{aligned} \tag{37}$$

where billed demand D_j, stolen consumption Q_j^S, and consumer surplus CS_j are evaluated at $(p_j, r_j(p_j))$. The firm supplies reliability using its estimated objective, which values consumer surplus at $\hat{\lambda}^f$; the planner instead places unit weight on consumer surplus. The steady-state

cost of reliability is

$$C_j(r_j) \;=\; \underbrace{RC \cdot \mathrm{E}[\text{replacements per period}]}_{\text{capital}} + \underbrace{\sum_{k \in j} \mathrm{E}[c(x_k;\theta_c)]}_{\text{maintenance}} + \underbrace{\psi(d_j;\theta_\psi,\Xi_j,X_j)}_{\text{repair effort}}, \tag{38}$$

evaluated at the stationary age distribution induced by the firm's optimal replacement policy. As in the firm's payoff, energy costs apply to both billed and stolen consumption. Expected repayments by evading households, $\phi_4 p_j Q_j^S$, are included in firm receipts and in households' expected cost of evasion, so they cancel as transfers in total surplus. The firm's payoff in (15) omits these relatively small receipts, and firm enforcement expenditures are not modeled. I exclude idiosyncratic replacement-cost shocks from C_j.

The break-even constraint requires the concession to recover at least $\bar{\pi}_f$ in profit, and links prices across its regions. For a given multiplier on this constraint, I search over candidate tariffs in each region, simulate the associated payment and consumption choices, solve the firm's replacement and restoration problems, and recompute equilibrium reliability from (34). The multiplier is then adjusted to satisfy the concession-level constraint. Regional tariffs can therefore redistribute revenue requirements across locations rather than simply lower the average price.

The optimal tariff trades off two forces. Reducing the tariff makes formal participation more attractive and increases billed demand, strengthening the firm's incentive to supply reliable electricity; raising the tariff collects more revenue from households that would remain in the formal market. Where payment participation is highly price-sensitive, the first force is more important. The induced improvement in reliability can further change the relative attractiveness of paying and evading. This general-equilibrium feedback means that a region's optimal tariff depends not just on demand elasticities but also on the firm's response to its changing payment base.

I compare the observed uniform tariff with two counterfactual pricing rules: an *optimal uniform* tariff, which is common across CONJs served by a distributor, and *optimal regional* tariffs, which can differ across its CONJs. Table 8 reports aggregate outcomes. The optimal uniform tariff lowers the average price from R$ 0.73 to R$ 0.71 per kWh, reducing evasion from 12.3% to 10.7% and raising the consumer-surplus index from 100 to 101.0. Allowing regional variation lowers evasion further, to 8.7%, while keeping the average tariff at R$ 0.73 per kWh; regional prices range from R$ 0.52 to R$ 0.94 per kWh. The outage-frequency index falls to 95, replacement rises from 4.5% to 4.6%, and mean restoration time falls from 9.8 to 9.0 hours. Aggregate consumer surplus increases by 2.4% relative to the observed tariff, of which 1.4 percentage points is gained beyond the optimal-uniform benchmark. Most of the consumer-surplus improvement thus comes from tailoring prices to local payment incentives rather than reducing their common level.

The aggregate gains mask sharp distributional differences. Table 9 compares the observed uniform and optimal regional tariffs in the lowest- and highest-income quartiles of CONJs.

Table 8: Counterfactual II: Optimal Tariffs

	Observed Uniform	Optimal Uniform	Optimal Regional
Average tariff (R$/kWh)	0.73	0.71	0.73
Minimum tariff (R$/kWh)	—	—	0.52
Maximum tariff (R$/kWh)	—	—	0.94
Payment evasion (%)	12.3	10.7	8.7
Outage frequency (index)	100	98	95
Transformer replacement rate (%)	4.5	4.5	4.6
Average restoration time (hours)	9.8	9.5	9.0
Consumer surplus (index)	100	101.0	102.4

Notes: Observed uniform is the baseline. Optimal uniform sets one tariff per concession; optimal regional sets a separate tariff for each CONJ. Both maximize (37) subject to concession-level break-even and allow the firm to re-optimize replacement and restoration. The household's option to evade and the estimated enforcement technology are held fixed. Outage frequency and consumer surplus are indexed to 100 in the observed-tariff equilibrium.

Under optimal regional pricing, the payment-evasion rate falls from 23% to 16% in the bottom quartile, versus only 4.2% to 4.1% in the top quartile. In poorer areas, the outage-frequency index declines from 100 to 83, the transformer replacement rate increases from 3.9% to 4.1%, and mean restoration time falls from 13.5 to 10.8 hours. Among richer areas, the corresponding changes in reliability and investment are modest.

The consumer-surplus consequences are particularly unequal. Regional pricing raises the surplus index from 100 to 109.4 in the bottom-income quartile, a gain of 9.4%, while the index falls from 100 to 99 in the top-income quartile, a decline of 1%. Thus, the aggregate gain of 2.4% does not represent an equal improvement everywhere: the policy benefits poorer areas much more strongly and entails a modest consumer-surplus loss in richer areas. Because enforcement is unchanged, these gains arise from the joint response of payment behavior and firms' reliability decisions to differentiated tariffs, rather than from a crackdown on illegal connections.

Table 9: Optimal Regional Tariffs: Heterogeneity by Income

	Bottom 25% Income		Top 25% Income	
	Observed Uniform	Optimal Regional	Observed Uniform	Optimal Regional
Payment evasion (%)	23.0	16.0	4.2	4.1
Outage frequency (index)	100	83	100	98
Transformer replacement rate (%)	3.9	4.1	5.3	5.3
Average restoration time (hours)	13.5	10.8	6.6	6.4
Consumer surplus (index)	100	109.4	100	99.0

Notes: Bottom and top quartiles are based on local CONJ income. Observed uniform denotes the current tariff regime; optimal regional denotes the counterfactual policy in Table 8. Outage frequency and consumer surplus are separately normalized to 100 within each income group in the baseline. Both consumer payment and firm reliability decisions adjust, but enforcement remains unchanged.

The regional-pricing exercise is a benchmark for the value of relaxing tariff uniformity, not necessarily a policy implementable under current regulations. It does not encompass instruments that target consumers within the same area, such as means-tested social tariffs, or interventions that directly change enforcement and loss reduction. Nevertheless, the results show that tariff design can affect more than transfers or electricity consumption. By changing who pays for service, tariffs also change firms' returns to maintaining infrastructure, equilibrium reliability, and the welfare of households that remain in the formal market.

10 Conclusion

This paper studies the feedback between electricity service quality and payment evasion in Brazilian distribution. Two reduced-form results establish the ingredients of the loop. In lower-income areas, accumulated service interruptions raise the share of unmetered electricity consumption. Across areas, higher non-technical losses are followed by fewer transformer replacements, conditional on the age of the installed stock. The first result is a behavioral response of households to service quality. The second concerns the firm's investment rather than realized reliability, and the link between investment and reliability runs through the age profile of transformer-caused interruptions.

The structural model built on these results serves three purposes: to price the value of reliability to households, formal and evading; to recover the weight distributors place on that value when they replace equipment and restore service; and to evaluate tariff rules when informality is endogenous. Eliminating payment evasion would lead firms to replace transformers more often and restore service faster, reducing outage frequency by 23 percent in high-evasion areas; households that pay would gain R$ 5.0 per month, or 3.4 percent of the average bill. Regional tariffs, lower where payment evasion is price-sensitive and higher where it is not, reduce evasion from 12.3 to 8.7 percent and raise consumer surplus by 2.4 percent while leaving the average tariff unchanged. Most of this gain comes from letting tariffs respond to local payment incentives rather than from changing their level.

The distributional results show that the costs of payment evasion are particularly concentrated in poorer areas. Eliminating evasion reduces outage frequency by 27 percent in the bottom income quartile, compared with only 3 percent in the top quartile. For households already paying for electricity, these improvements translate into welfare gains of R$ 21 per month in the bottom quartile, versus R$ 1 in the top. Similarly, optimal regional tariffs increase consumer surplus by 9.4 percent in the poorest quartile, while slightly reducing it in the richest. These differences highlight an important consequence of the feedback loop: payment evasion imposes costs not only on distributors but also on households that remain in the formal market, particularly those living in areas where nonpayment is widespread.

These findings suggest that electricity reliability regulation and policies addressing payment evasion should not be considered independently. Improving service quality can strengthen incentives to pay, while policies that increase formal participation can improve the private returns to infrastructure investment. Tariff regulation therefore affects more than electricity consump-

tion and the distribution of payments: it also shapes the reliability that firms choose to provide. The regional-pricing counterfactual illustrates the potential gains from accounting for these interactions, although it is a benchmark rather than a policy directly implementable under current regulations. The analysis holds enforcement technology fixed and does not evaluate the optimal combination of enforcement and reliability investment. Nevertheless, the central implication extends beyond electricity distribution: when consumers' willingness to pay depends on service quality, and firms' incentives to provide quality depend on payment, policies addressing either side of the market can have consequences for the other.

References

Abito, Jose Miguel, "Measuring the welfare gains from optimal incentive regulation," *The Review of Economic Studies*, 2020, *87* (5), 2019–2048.

Alm, James, Gary H McClelland, and William D Schulze, "Why do people pay taxes?," *Journal of public Economics*, 1992, *48* (1), 21–38.

ANEEL, "Relatório de Perdas de Energia," `https://aneel.gov.br` 2026. Accessed: 2026-10-04.

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

Balboni, Clare, Oriana Bandiera, Robin Burgess, Maitreesh Ghatak, and Anton Heil, "Why do people stay poor?," *The Quarterly Journal of Economics*, 2022, *137* (2), 785–844.

Banerjee, Abhijit V and Andrew F Newman, "Occupational choice and the process of development," *Journal of political economy*, 1993, *101* (2), 274–298.

Berkouwer, Susanna B, Pierre E Biscaye, Maya Mikdash, Steven L Puller, and Catherine D Wolfram, "Voltage quality and economic activity," *Work. Pap., Wharton Sch., Univ. Penn., Philadelphia*, 2024.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, "Robust Inference with Multiway Clustering," *Journal of Business & Economic Statistics*, 2011, *29* (2), 238–249.

Cisse, Abdoulaye, "The Value of Electricity Reliability: Evidence from a Natural Experiment in Senegal," Technical Report, Technical report, Working Paper 2025.

ConJur, "'"See how the Superior Court of Justice (STJ) has judged actions on energy supply"'," *Consultor Jurídico* October 2019. Accessed August 5, 2026.

Corbae, Dean and Andrew Glover, "Employer credit checks: Poverty traps versus matching efficiency," *Review of Economic Studies*, 2025, *92* (3), 1661–1698.

dos Santos, Carolina Tojal Ramos and Bruna Morais Guidetti, "Expansion of Piped Water and Sewer Networks: The Effects of Regulation," IDB Working Paper IDB-WP-1734, Inter-American Development Bank July 2025.

Dubin, Jeffrey A and Daniel L McFadden, "An econometric analysis of residential electric appliance holdings and consumption," *Econometrica: Journal of the Econometric Society*, 1984, pp. 345–362.

Dwenger, Nadja, Henrik Kleven, Imran Rasul, and Johannes Rincke, "Extrinsic and intrinsic motivations for tax compliance: Evidence from a field experiment in Germany," *American Economic Journal: Economic Policy*, 2016, *8* (3), 203–232.

Dzansi, James, Steven L Puller, Brittany Street, and Belinda Yebuah-Dwamena, "The vicious circle of blackouts and revenue collection in developing economies: Evidence from Ghana," *Technical Report*, 2018.

Fowlie, Meredith, "Emissions trading, electricity restructuring, and investment in pollution abatement," *American Economic Review*, 2010, *100* (3), 837–869.

G1, "Consumidor ficou em média 10,4 horas sem energia em 2023: veja ranking de distribuidoras com melhores e piores desempenhos," `https://g1.globo.com/economia/noticia/2024/03/16/consumidor-ficou-em-media-104-horas-sem-energia-em-2023-veja-ranking-de-distribuidoras.ghtml` 2024. Accessed: October 6, 2026.

Gagnepain, Philippe and Marc Ivaldi, "Incentive regulatory policies: the case of public transit systems in France," *RAND Journal of Economics*, 2002, pp. 605–629.

Gowrisankaran, Gautam, Ashley Langer, and Mar Reguant, "Energy transitions in regulated markets," Technical Report, National Bureau of Economic Research 2024.

Hotz, V Joseph and Robert A Miller, "Conditional choice probabilities and the estimation of dynamic models," *The Review of economic studies*, 1993, *60* (3), 497–529.

Ito, Koichiro, "Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing," *American Economic Review*, 2014, *104* (2), 537–563.

Jack, Kelsey and Grant Smith, "Charging ahead: Prepaid metering, electricity use, and utility revenue," *American Economic Journal: Applied Economics*, 2020, *12* (2), 134–168.

Jessoe, Katrina and David Rapson, "Knowledge is (less) power: Experimental evidence from residential energy use," *American Economic Review*, 2014, *104* (4), 1417–1438.

Lee, Kenneth, Edward Miguel, and Catherine Wolfram, "Does household electrification supercharge economic development?," Technical Report, National Bureau of Economic Research 2019.

_ , _ , and _ , "Experimental evidence on the economics of rural electrification," *Journal of Political Economy*, 2020, *128* (4), 1523–1565.

Lim, Claire SH and Ali Yurukoglu, "Dynamic natural monopoly regulation: Time inconsistency, moral hazard, and political environments," *Journal of Political Economy*, 2018, *126* (1), 263–312.

Lipscomb, Molly, A Mushfiq Mobarak, and Tania Barham, "Development effects of electrification: Evidence from the topographic placement of hydropower plants in Brazil," *American Economic Journal: Applied Economics*, 2013, *5* (2), 200–231.

Luttmer, Erzo FP and Monica Singhal, "Tax morale," *Journal of economic perspectives*, 2014, *28* (4), 149–168.

McRae, Shaun, "Infrastructure quality and the subsidy trap," *American Economic Review*, 2015, *105* (1), 35–66.

Reiss, Peter C and Matthew W White, "Household electricity demand, revisited," *The Review of Economic Studies*, 2005, *72* (3), 853–883.

Resende, Davi, Gabriel Richter, Marcelo Castello Branco Sant'Anna, and André Trindade, "Pricing and Informality: Evidence from Energy Theft in Brazil," *Available at SSRN*, 2025.

Rust, John, "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher," *Econometrica: Journal of the Econometric Society*, 1987, pp. 999–1033.

Seim, Katja and Joel Waldfogel, "Public monopoly and economic efficiency: evidence from the Pennsylvania liquor control board's entry decisions," *American Economic Review*, 2013, *103* (2), 831–862.

Szabó, Andrea and Gergely Ujhelyi, "Reducing nonpayment for public utilities: Experimental evidence from South Africa," *Journal of Development Economics*, 2015, *117*, 20–31.

Timmins, Christopher, "Measuring the dynamic efficiency costs of regulators' preferences: Municipal water utilities in the arid west," *Econometrica*, 2002, *70* (2), 603–629.

Weigel, Jonathan L, "The participation dividend of taxation: How citizens in Congo engage more with the state when it tries to tax them," *The Quarterly Journal of Economics*, 2020, *135* (4), 1849–1903.

Wilson, Kyle, "Does public competition crowd out private investment? evidence from municipal provision of internet access," *American Economic Journal: Microeconomics*, 2025, *17* (4), 399–431.

Appendix A: Tables and Figures

Figure A1: Illegal Grid Taps

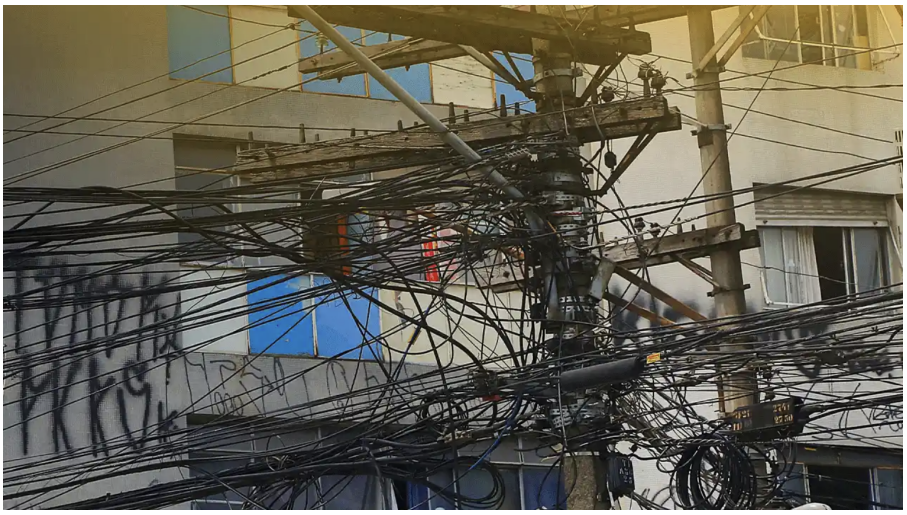
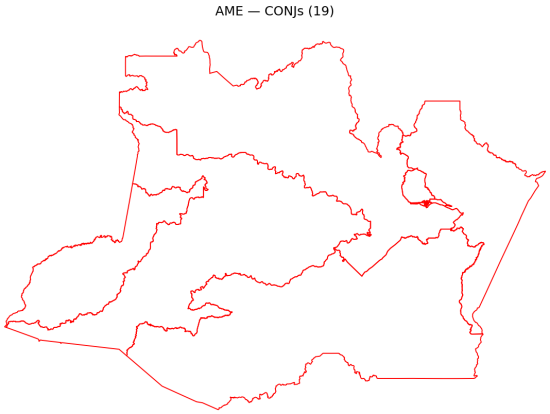
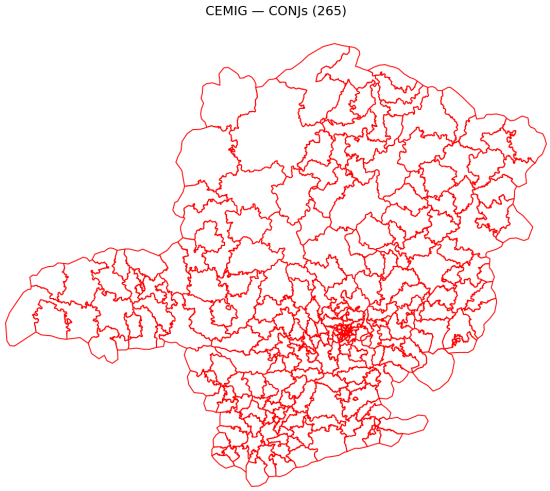


Figure Notes: Example of illegal grid connections.

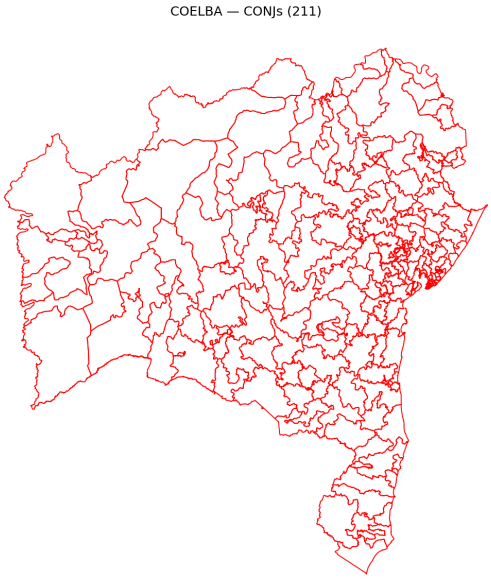
Figure A2: Distribution Areas



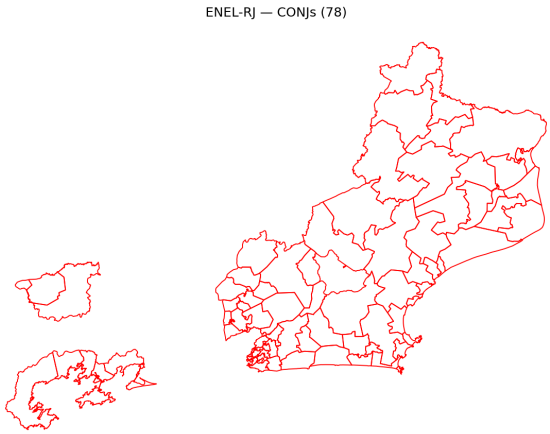
(a) AME



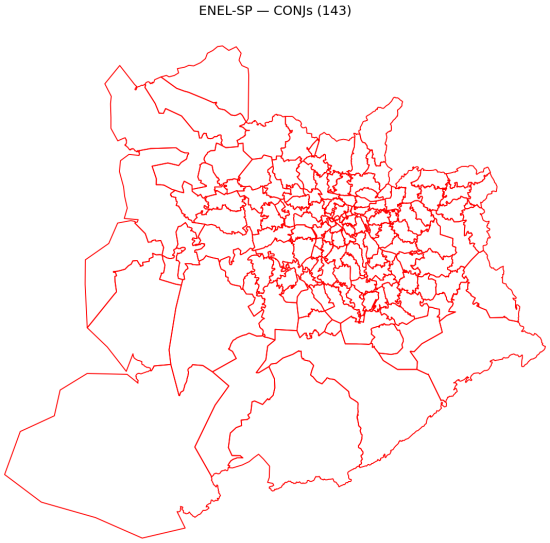
(b) CEMIG



(c) COELBA



(d) ENEL-RJ



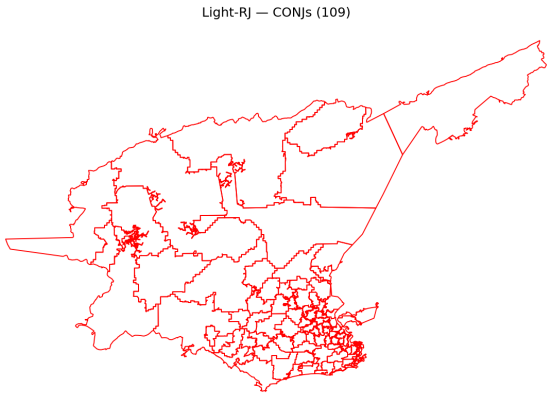
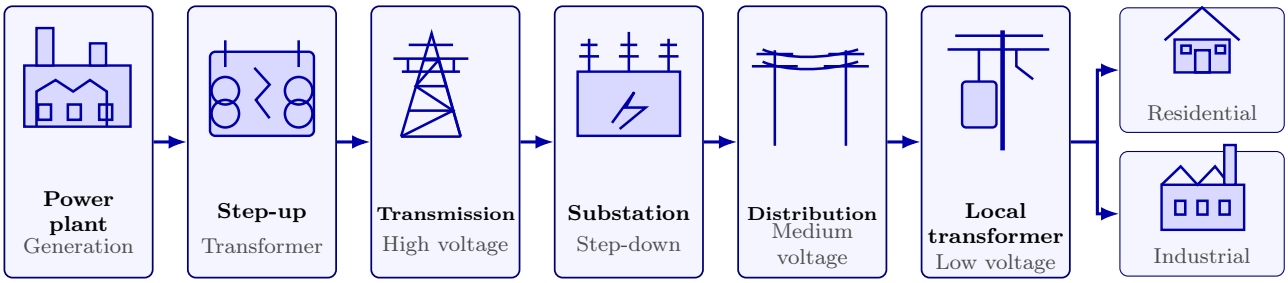
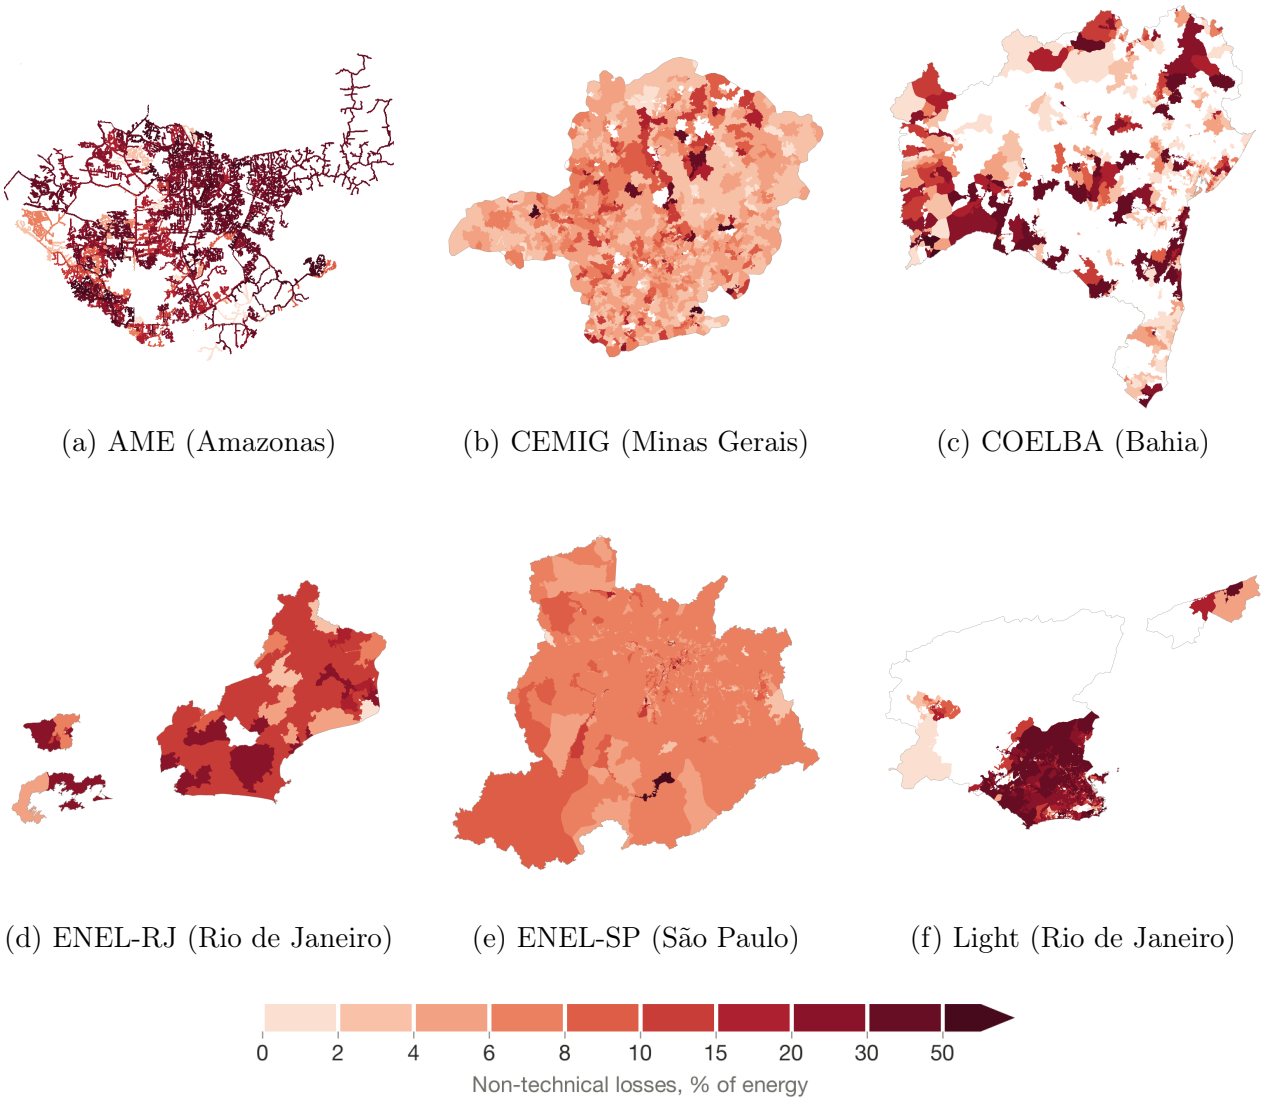


Figure A3: Electricity path from generation to final consumption



Notes: Schematic illustration; voltage levels and network configurations vary across systems. Source: Author's illustration.

Figure A4: Non-Technical Losses by Feeder

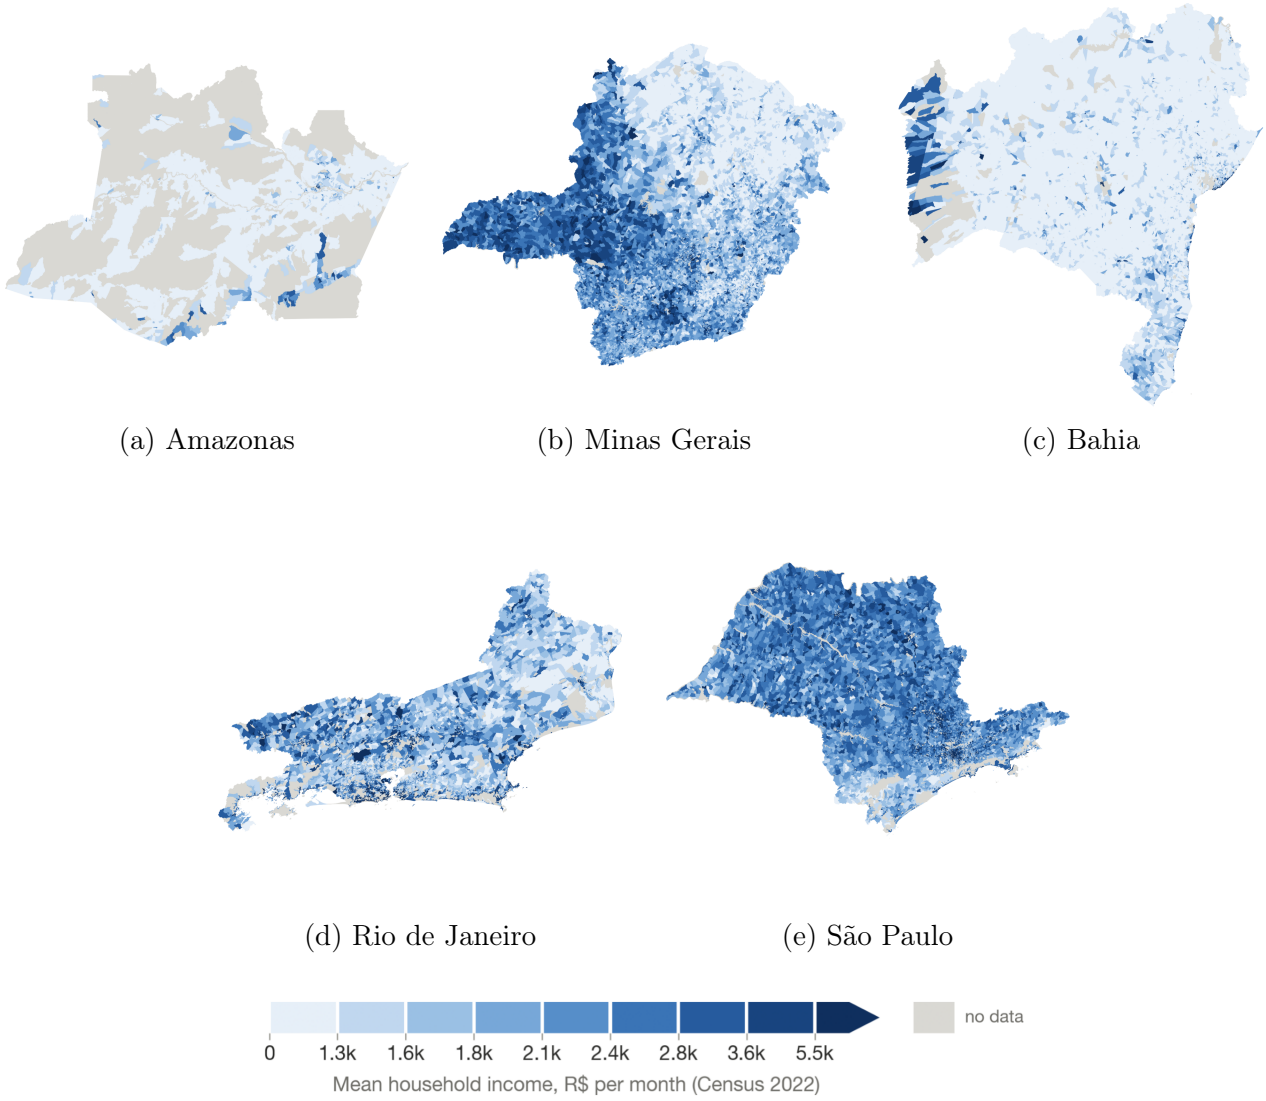


(a) AME (Amazonas) (b) CEMIG (Minas Gerais) (c) COELBA (Bahia)

(d) ENEL-RJ (Rio de Janeiro) (e) ENEL-SP (São Paulo) (f) Light (Rio de Janeiro)

Figure Notes: Each panel shows the non-technical loss share of medium-voltage feeders (CTMT), $(\text{PNT}^{BT} + \text{PNT}^{MT})/(\text{PNT}^{BT} + \text{PNT}^{MT} + E)$, where E is annual paid energy, from ANEEL's BDGD for 2024. Panels (b)–(f) assign each point of the concession area to the nearest medium-voltage line (SSDMT) and color it by that line's feeder; panel (a) shows the medium-voltage network of Manaus, colored by feeder. White marks feeders without a valid share: zero injected energy, negative reported losses, or placeholder losses (1 kWh per month).

Figure A5: Mean Household Income by Census Tract

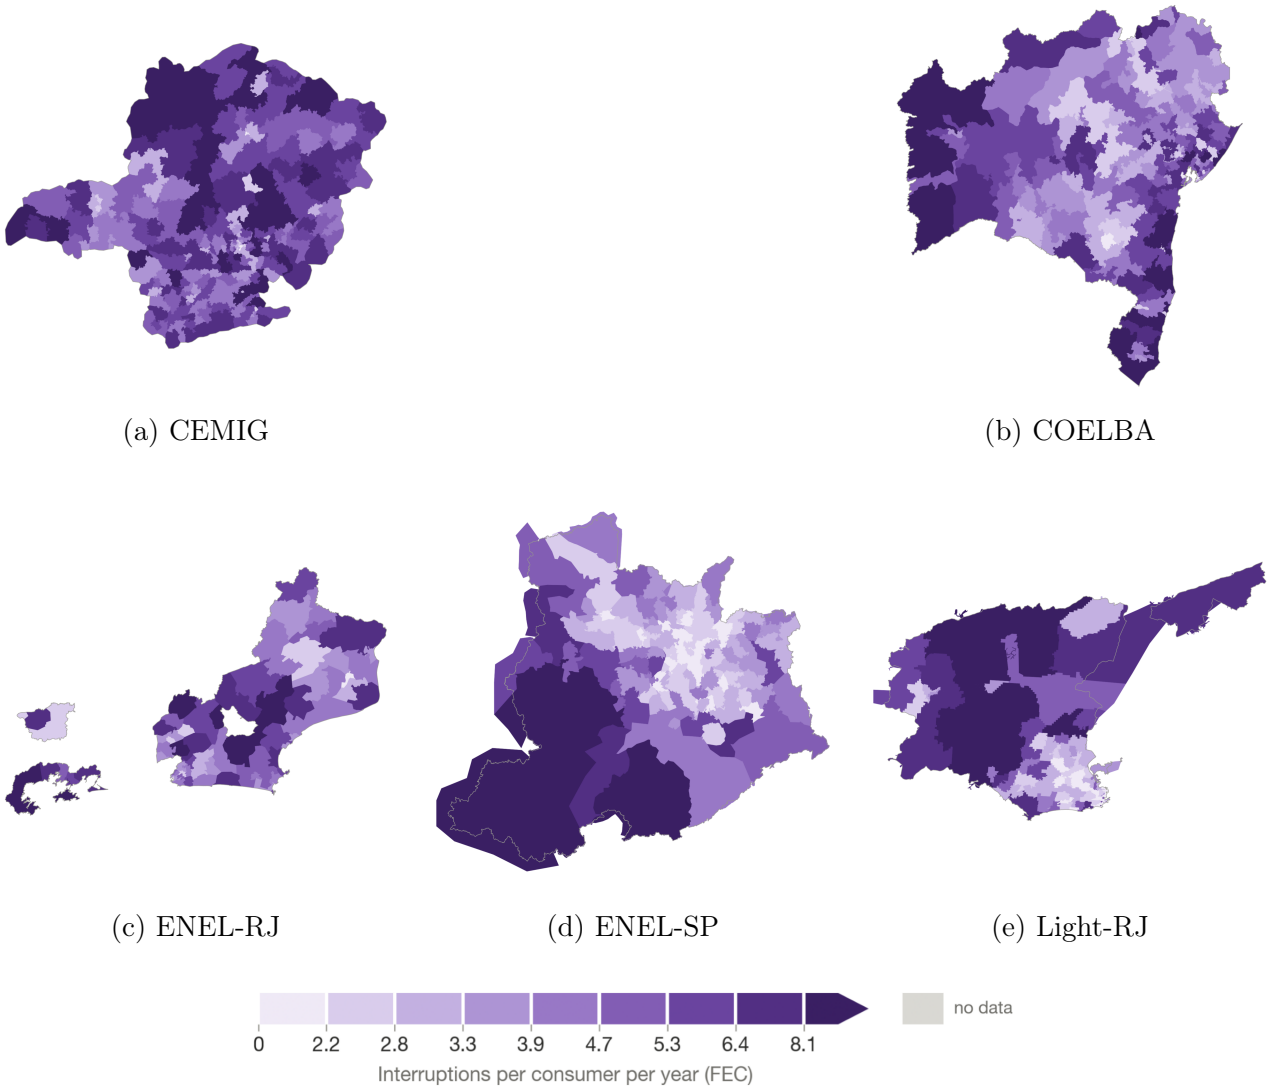


(a) Amazonas (b) Minas Gerais (c) Bahia

(d) Rio de Janeiro (e) São Paulo

Figure Notes: Mean monthly nominal income of the persons responsible for private households (V06004), from the 2022 Census tract aggregates. Classes are deciles of the pooled tracts of the five states. Grey marks tracts with suppressed or missing income.

Figure A6: Interruption Frequency by CONJ

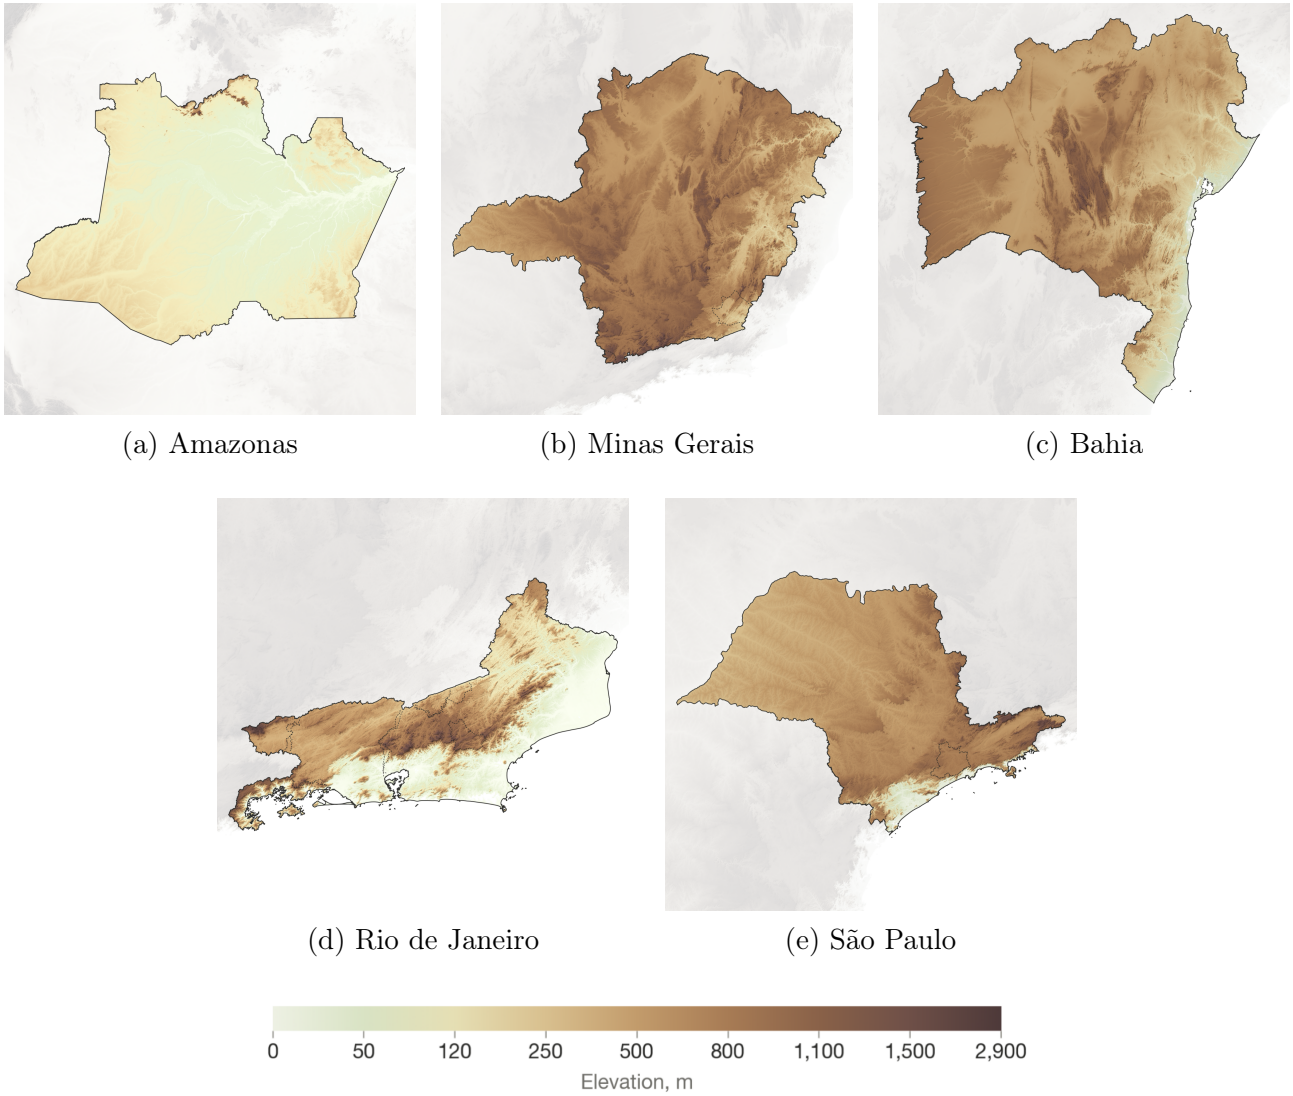


(a) CEMIG (b) COELBA

(c) ENEL-RJ (d) ENEL-SP (e) Light-RJ

Figure Notes: Interruptions lasting at least three minutes per consumer unit per year (FEC), by consumer set (CONJ), averaged over the months observed between 2021 and 2024 and annualized. Classes are deciles of the pooled CONJs of the six distributors. Grey marks CONJs without records.

Figure A7: Topography of the Concession States

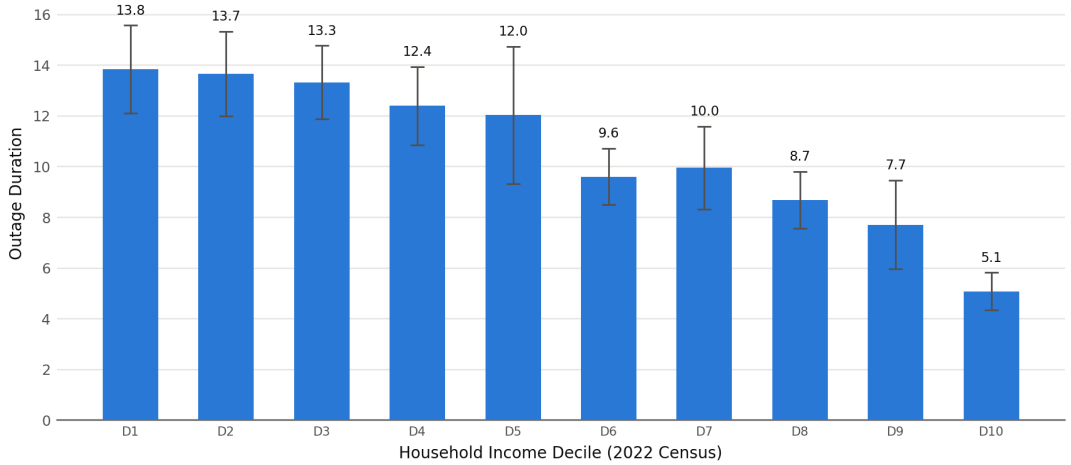


(a) Amazonas (b) Minas Gerais (c) Bahia

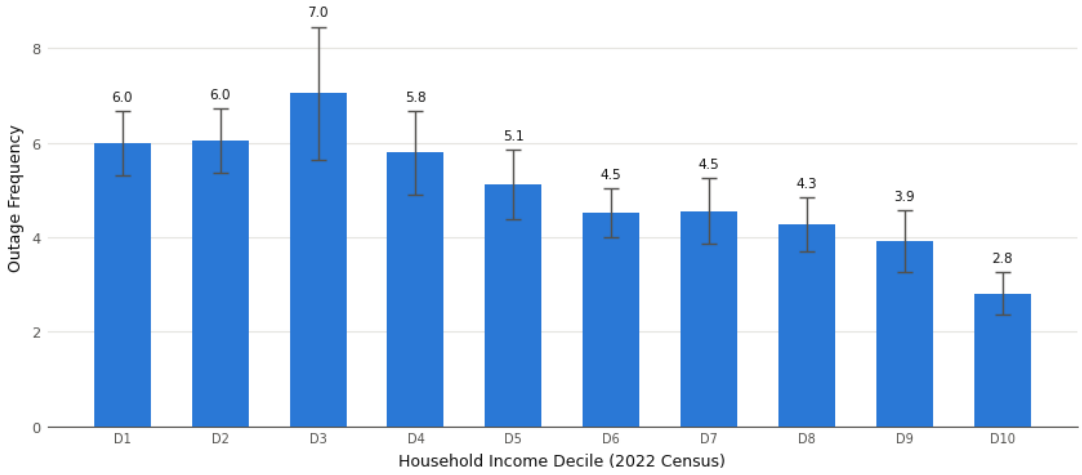
(d) Rio de Janeiro (e) São Paulo

Figure Notes: Elevation from the Copernicus GLO-90 digital elevation model (90-meter resolution), shaded from the northwest. The color scale is common to the five states. Dashed lines mark concession boundaries inside each state. Copernicus DEM © DLR e.V. 2010–2014 and © Airbus Defence and Space GmbH 2014–2018, provided under COPERNICUS by the European Union and ESA.

Figure A8: Electricity Reliability by Income Decile



(a) Average Outage Duration



(b) Average Monthly Outage Frequency

Figure Notes: Panel (a) shows average outage duration by income decile. Panel (b) shows average monthly outage frequency by income decile. Both panels pool observations across distribution companies.

Table A1: Service Reliability and Payment Evasion: Topography IV

Dep. Variable: Payment Evasion	(1) OLS	(2) IV
Panel A. Payment Evasion		
Average Outage Duration	0.053**	0.033**
	(0.022)	(0.0138)
Panel B. First Stage: Average Outage Duration		
Std Dev of Elevation	–	0.004
	–	(0.00)
Std Dev of Elevation	–	3(10^−4)
	–	(4(10^−5))
First-Stage F-Statistic	–	29.5
Observations	952	952
Demographic Controls	✓	✓

Notes: Panel A reports OLS and IV estimates of the relationship between average outage duration and payment evasion. Column (2) instruments average outage duration using the mean and standard deviation of elevation within each location. Panel B reports the first-stage coefficients and the F-statistic for the excluded instruments. Both specifications include demographic controls. Standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Appendix B: Data Cleaning and Variable Construction

The CONJ-month panel covers the 825 CONJs of the six distributors from January 2021 to December 2024, or 39,600 CONJ-months. Energy balances, continuity indicators, and emergency-service indicators are monthly. The BDGD is released as one snapshot per distributor and year, so network assets are observed annually, while consumer-unit records carry monthly consumption and interruption fields within each snapshot. This appendix states the cleaning rules and the construction of each variable; Table A2 reports the sample at each step.

B.1 Non-Technical Losses

The non-technical-loss share of CONJ j in month t is

$$ntl_{jt} = \frac{PNT_{jt}}{PNT_{jt} + ENE_{jt}},$$

where PNT_{jt} is the non-technical losses in the low-voltage networks and ENE_{jt} is the energy in the feeder energy balance, all in kWh and aggregated from feeders to the CONJ. Losses are injected energy minus billed energy minus estimated technical losses, and are negative in months in which billed energy or estimated technical losses exceed injected energy. I apply five rules.

1. *Negative entries.* 1,099 CONJ-months (2.8%) have a negative $PNTBT$, $PNTMT$, or ENE, which yields shares below zero (955 CONJ-months) or above one (55). I discard every CONJ-month with a negative entry. This removes all shares outside $[0, 1]$.

2. *Persistent negatives.* Negative entries are isolated in most CONJs: among the 241 CONJs with at least one, the median CONJ has two. In 29 CONJs (23 in Light, five in COELBA, one in CEMIG) they occur in more than 20% of months, and they cluster: a negative month is followed by a month 0.46 standard deviations below the CONJ's own mean rather than by a compensating spike, and the within-CONJ autocorrelation of losses is 0.39, against 0.65 elsewhere. I treat the balances of these CONJs as mismeasured and drop the CONJs. Their mean share is close to that of the retained CONJs of the same distributor (0.28 against 0.30 in Light), so the exclusion does not select on evasion. Dropping every CONJ with any negative month would instead discard 241 CONJs, disproportionately poor ones.

3. *Undefined shares.* I drop 58 CONJ-months with a missing balance or with zero losses and zero energy.

4. *Outliers.* I winsorize ntl_{jt} at its 99th percentile, 0.38, which affects 377 CONJ-months.

The annual share in Section 3.2 is the mean of the monthly shares of the CONJ-year.

Table A2: Sample Construction: Non-Technical-Loss Share

	AME	COELBA	CEMIG	ENEL-RJ	ENEL-SP	LIGHT	Total
Panel A. CONJs							
In the panel	19	211	265	78	143	109	825
Dropped: persistent negative entries	0	5	1	0	0	23	29
Without Census link	0	56	88	7	2	7	160
Panel B. CONJ-months							
In the panel	912	10,128	12,720	3,744	6,864	5,232	39,600
Dropped: persistent negative entries	0	240	48	0	0	1,104	1,392
Dropped: other negative entries	0	440	16	0	0	91	547
Dropped: missing or zero balance	0	57	0	1	0	0	58
Retained	912	9,391	12,656	3,743	6,864	4,037	37,603
Excluded: reporting breaks	228	0	3,168	0	3,432	1,023	7,851
Retained, excluding breaks	684	9,391	9,488	3,743	3,432	3,014	29,752
Mean non-technical-loss share	0.219	0.063	0.065	0.152	0.069	0.240	0.097

Notes: The panel runs from January 2021 to December 2024. Persistent negative entries: CONJs with a negative $PNTBT$, $PNTMT$, or ENE in more than 20% of months. Reporting breaks: AME 2021, CEMIG 2023, ENEL-SP 2023–2024, and Light 2022. Without Census link: CONJs to which Census tracts cannot be apportioned; 159 of them also lack continuity, emergency-service, and asset records before 2024 (before 2023 in Light). The mean share is the unweighted mean of the winsorized share over retained CONJ-months outside reporting breaks.

B.2 Interruptions and Restoration

Frequency and duration. FEC_{jt} is the number of interruptions lasting at least three minutes per consumer unit in CONJ j and month t, and DEC_{jt} the hours of interruption per consumer unit. The accumulated regressor in (1) sums the twelve preceding months and is defined only when all twelve are observed, so the sample of Section 3.1 starts in January 2022. The event records list each interruption with its cause, duration, and number of consumers affected; I aggregate them to the CONJ-month and classify an event as weather-related when its recorded cause is a weather event.

Restoration time. I measure restoration effort by the mean execution time of emergency services in the CONJ-month (TME), from the regulator's emergency-service indicators, which also report the number of emergency occurrences and the mean preparation and travel times of crews. I do not use the ratio DEC_{jt}/FEC_{jt}, whose denominator carries the measurement error in FEC_{jt}.

Missing indicators. For 159 CONJs (88 in CEMIG, 55 in COELBA, seven in ENEL-RJ, seven in Light, and two in ENEL-SP), continuity indicators, event records, emergency-service indicators, and asset records start in 2024 (in 2023 for Light), and these CONJs, together with one more in COELBA, cannot be linked to Census tracts. They enter only specifications that use none of these variables.

B.3 Transformers and Replacement

The asset registry has two layers. A transformer unit (UNTRMT) is a fixed position in the network, linked to its feeder and CONJ; the equipment installed at the unit (EQTRMT) carries

its own code, nominal capacity, and capitalization date. Age is the snapshot year minus the year of the capitalization date of the installed equipment, set to zero when the date is later than the snapshot.

A unit is replaced between snapshots t and $t+1$ when the equipment recorded at the unit in $t+1$ has a different code and its age does not advance by one year. A code change with continuing age is a re-coding of the same equipment: in Light, the equipment code changes in 7.9% of unit-years, but in three quarters of these cases the age advances by exactly one year, and the replacement rate is 2.0%. Units absent from the $t+1$ snapshot are removals and units that first appear in $t+1$ are additions; neither counts as a replacement.

The replacement rate R_{jt} in (2) divides the replacements in CONJ j and year t by the number of installed units, and the age moments describe the stock in the snapshot of year $t-1$. Replacement requires two consecutive snapshots, so it is not defined in 2021; it is also missing for AME in 2022 and for CEMIG in 2023. For the reliability technology, each unit is assigned to the CONJ and feeder of the majority of the registered consumer units it serves, N_k is the number of those units, and the unit's interruptions in a month are the mean individual interruption count of the units it serves.

B.4 Consumer Units and Household Counts

Consumer units. From the low-voltage consumer-unit records (UCBT), I keep active units (SIT_ATIV = AT) in the residential subclasses (CLAS_SUB RE1–RE6, which include the low-income subclasses of the Social Electricity Tariff). I discard unit-months with zero, negative, or missing consumption or with consumption of 5,000 kWh or more, and units with missing or nonpositive installed capacity (CAR_INST, in kW). Mean log consumption in the CONJ-month uses all retained units.

Reference consumption. I set q_0 in (4) to the first percentile of monthly consumption across the household-months of the demand sample, 1 kWh. Because q_0 cancels from demand and regime choices, it affects only the level of surplus and of $\partial CS_j/\partial r$; I report welfare and supply estimates also at the 5th and 10th percentiles, 8.2 and 27.3 kWh.

Availability. Individual interruption hours (DIC) and counts (FIC) are bounded at 730 hours and 200 interruptions per month. F_{jt} and $F_{jt}d_{jt}$ are the means of FIC and DIC across the registered residential units of the CONJ, and availability is $r_{jt} = 1 - F_{jt}d_{jt}/T$, with $T = 730$ hours. Average availability $\bar{r}_{jt}$ in (3) is defined only when all twelve preceding months are observed.

Household counts. The household count N_j^{CNEFE} is the number of private dwellings in the 2022 CNEFE (COD_ESPECIE = 1), in which each apartment is a separate record. I assign each dwelling to the CONJ of the nearest registered consumer unit, with distances computed in a projected coordinate system (SIRGAS 2000 / Brazil Polyconic), and drop dwellings more than 500 meters from any registered unit.

Formal share. The formal share is $N_{jt}^{UCBT}/N_j^{CNEFE}$, where N_{jt}^{UCBT} counts active residential units.

Measurement. The dwelling count is a 2022 snapshot, while registered units are counted in each year. The mapping treats one registered unit as one household, and leaves in the losses the unmetered consumption of non-residential low-voltage users.

B.5 Census, Tariffs, Weather, and Topography

Census. Census variables are 2022 tract-level aggregates. Let N_{sj} be the number of registered consumer units of CONJ j located in tract s, and $N_s = \sum_j N_{sj}$. Tract counts are apportioned as $X_j = \sum_s (N_{sj}/N_s)\, X_s$, and tract means are averaged with weights proportional to N_{sj}. Mean household income is the mean monthly nominal income of the persons responsible for private households (variable V06004). The income groups of Section 3.1 split CONJs at the median of this variable, with one observation per CONJ.

Tariffs. The price is the conventional residential (B1) tariff approved by ANEEL for each distributor, the sum of its energy (TE) and network-use (TUSD) components plus taxes in R$ per MWh. It changes with each annual readjustment and periodic review. The flag surcharge is added to it: zero under the green flag, positive under the yellow and the two red flags, and R$ 142 per MWh under the water-scarcity flag of September 2021 to April 2022. In a month in which the flag changes, the surcharge is weighted by the days under each flag.

Weather and topography. Weather variables come from the ERA5 reanalysis on a 0.25-degree grid, averaged over the grid cells that cover each CONJ: 2-meter temperature, converted to degrees Celsius, precipitation, in millimeters per day, and 10-meter wind speed. Extreme-wind days are days on which sustained wind speed in the CONJ reaches 64 km/h for at least one hour. Elevation and slope are means, standard deviations, and maxima over the CONJ polygon from Topodata, which is based on the Shuttle Radar Topography Mission.

Note on Consumer Surplus

Utility is quasi-linear, so surplus is in reais per household-month. Let $P(x; A, r) = Ar^{\gamma}x^{-1/\eta}$ denote the household's marginal willingness to pay for its x-th kWh, with $\alpha \equiv 1 - 1/\eta$.

Formal household. A formal household buys $q^F = (Ar^{\gamma}/p)^{\eta}$ at tariff p. Its surplus is

$$v^F(p, r; A) \;=\; \int_{q_0}^{q^F} P(x; A, r)\, dx \;-\; p\, q^F \;=\; Ar^{\gamma}\, \frac{(q^F)^{\alpha} - q_0^{\alpha}}{\alpha} \;-\; p\, q^F, \tag{39}$$

the area under its inverse demand above the reference level q_0, minus the bill. Because demand is inelastic ($\eta < 1$, $\alpha < 0$), the area under P diverges as $x \to 0$; without q_0, surplus would fall with reliability although consumption rises with it. With q_0, $\partial v^F/\partial p = -q^F$ and

$$\frac{\partial v^F}{\partial r} \;=\; \frac{\gamma}{r}\, Ar^{\gamma}\, \frac{(q^F)^{\alpha} - q_0^{\alpha}}{\alpha} \;>\; 0 \quad\Longleftrightarrow\quad q^F > q_0. \tag{40}$$

Evading household. An evading household consumes up to the capacity of its connection, $\bar{q}_i = \exp(\kappa_0 + \kappa_1 \log c_i)$, pays no tariff, and bears the expected enforcement and repayment cost in $g(\cdot)$:

$$v^S(p, r; A) \;=\; Ar^{\gamma}\, \frac{\bar{q}^{\,\alpha} - q_0^{\alpha}}{\alpha} \;+\; g(r, \bar{q}, p), \qquad g = \phi_1 r + \phi_2 r^2 + \phi_3 r\bar{q} - \phi_4 p\bar{q}. \tag{41}$$

The q_0 terms are the same in both regimes, so q_0 cancels from $v^F - v^S$ and is not identified by demand or regime choices.

Regime choice and market surplus. Household i is formal, $d_i = 1$, when $v_i^F - v_i^S + \Delta\xi_{jt} + e_i \geq 0$, with $\Delta\xi_{jt}$ the market's unobserved taste for formal service and $e_i \sim \mathcal{N}(0, \sigma_e^2)$ correlated ρ with the consumption shock ν_i. Payoffs and shocks are all in reais, so formal service carries the premium $\Delta\xi_{jt} + e_i$, $U_i = \max\{v_i^F + \Delta\xi_{jt} + e_i,\, v_i^S\}$, and

$$CS_{jt} \;=\; N_j \int U_i\, dF(\nu_i, e_i \mid X_j), \tag{42}$$

with N_j the CNEFE household count.

What the supply side uses. The distributor's marginal payoff of reliability weights $\partial CS_j/\partial r$ by λ^f. Because $\Delta\xi$ and e do not depend on r and households on the margin between regimes are indifferent, the envelope theorem gives

$$\frac{\partial CS_{jt}}{\partial r} \;=\; N_j\, \mathbb{E}\!\left[\, d_i\, \frac{\partial v_i^F}{\partial r} \;+\; (1 - d_i)\, \frac{\partial v_i^S}{\partial r} \right], \tag{43}$$

so regime switching contributes nothing at the margin. This holds in the steady state, where the trailing availability $\bar{r}_{jt}$ that drives regime choice equals r_{jt}; with $\rho \neq 0$, each term averages

A_i over the households selected into that regime.

Appendix C: A Simple Model of the Feedback

Households decide whether to pay given the reliability they receive, and the distributor decides how much reliability to supply given the share of households who pay. In a minimal model, the interaction amplifies shocks and can sustain a low-reliability, high-evasion state.

Households. A unit mass of households shares one network with reliability $r \in [0,1]$, the fraction of time it supplies power. The tariff is linear: a formal household pays p per unit of power it receives. Household i values each unit at A_i, and each unit of downtime does damage D (spoiled food, damaged appliances). A formal household receives $(A_i - p)\, r - D(1-r)$. An evading household draws the same power from the same network and pays nothing, but with probability ϕ it is detected and disconnected and receives nothing; its expected payoff is $(1-\phi)\,[A_i r - D(1-r)]$. Household i pays if and only if the connection it risks losing is worth more than the bill it saves, $\phi\,[A_i r - D(1-r)] \geq p\, r$, that is,

$$A_i \;\geq\; A^*(r) \equiv \frac{p}{\phi} + D\,\frac{1-r}{r}. \tag{44}$$

The bill scales with the power received, but the value of the connection does not: on a network that fails often, the connection is worth little, and the threat of disconnection does not deter. Households with $A_i < p/\phi$ evade even when service is perfect. A share $1-\bar{s}$ cannot evade (apartment buildings, connections in a landlord's name); among the rest $A_i \sim H$, density h, and

$$s = S(r) \equiv \bar{s}\, H\big(A^*(r)\big), \qquad S'(r) = -\bar{s}\, h\big(A^*(r)\big)\,\frac{D}{r^2} < 0. \tag{45}$$

Distributor. The distributor buys energy at w per unit for every household it supplies, formal or undetected, bills only formal ones, and chooses reliability at maintenance cost $C(r)$, $C', C'' > 0$:

$$\Pi(r;s) = r\big[(1-s)\,p - (1-\phi s)\,w\big] - C(r) \quad\Longrightarrow\quad r = R(s), \qquad R'(s) = -\frac{p-\phi w}{C''(r)} < 0. \tag{46}$$

Undetected evasion is energy bought but not billed, so evasion lowers the margin on each unit of uptime and with it the return to maintenance. Equations (45) and (46) are the two arrows of Sections 3.1 and 3.2.

Steady states and amplification. Households choose their regime on past reliability and the distributor takes current evasion as given, so $s_t = S(r_{t-1})$, $r_t = R(s_t)$, and $r_t = \Phi(r_{t-1})$ with $\Phi \equiv R \circ S$ increasing. A steady state $r^* = \Phi(r^*)$ is stable when the loop gain

$$L \equiv R'(s^*)\, S'(r^*) = \frac{(p-\phi w)\,\bar{s}\, h(A^*(r^*))\, D}{C''(r^*)\,(r^*)^2} \tag{47}$$

is below one: the fall in reliability caused by a unit rise in evasion, times the rise in evasion caused by a unit fall in reliability. Let a shock δ raise marginal maintenance cost to $C'(r)+\delta$, as storm damage does. Holding evasion fixed, reliability falls by $\partial R/\partial\delta = -1/C''$. In equilibrium,

$$\frac{dr^*}{d\delta} = \frac{1}{1-L}\frac{\partial R}{\partial \delta}, \tag{48}$$

so at a stable steady state the feedback multiplies the direct effect by $1/(1-L) > 1$: households evade in response to the outage, and their evasion erodes the return to maintenance further. A fall in detection ϕ works on both arrows at once: it raises A^*, so more households evade at any r, and it raises the cost of each evader to the distributor.

A threshold. Where many households are close to indifferent, h is large and L can exceed one. Φ then crosses the 45-degree line three times (Figure A9a): a high steady state with reliable service and little evasion, a low one with poor service and widespread evasion, and an unstable threshold between them. A small outage shock to a network in the high state is absorbed; a shock that pushes reliability below the threshold is not, and the network settles in the low state after the shock has passed (Figure A9b). By the same logic, a temporary investment large enough to cross the threshold moves a network out of the low state.

Scope. The model states a mechanism, not a magnitude, and does not imply that any network in the data is in a low state. A distributor that anticipated the evasion response, or that weighted consumer surplus by λ^f, would supply more reliability than $R(s)$. Sections 4 and 5 estimates the two arrows and how far the distributor internalizes them.

Figure A9: The Feedback Between Reliability and Evasion

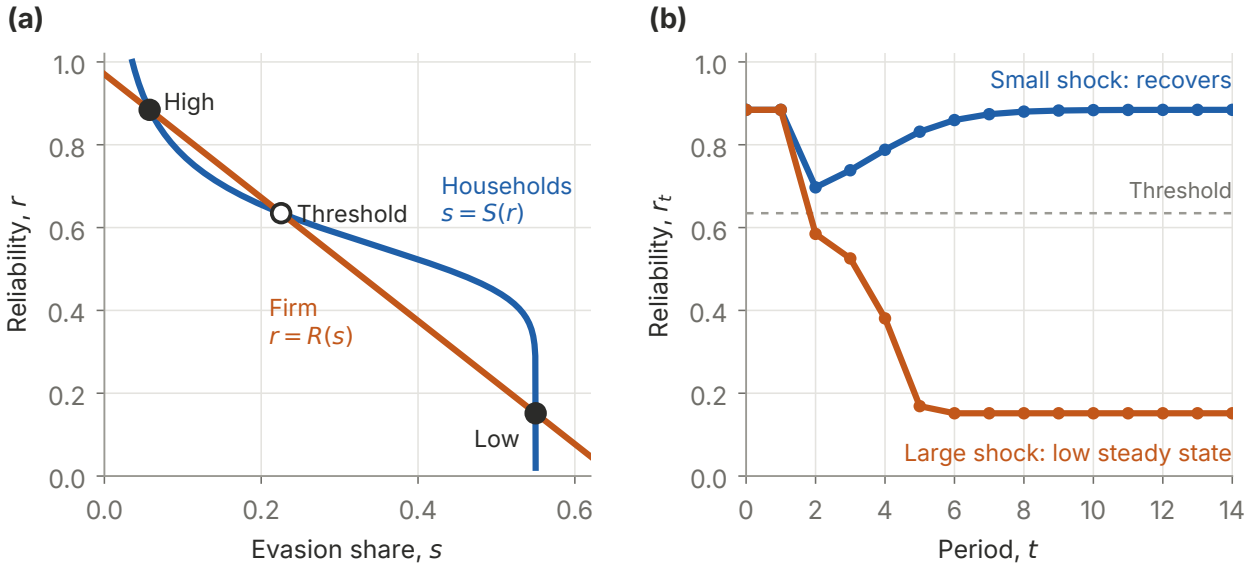


Figure Notes: Numerical example with $p = 0.5$, $\phi = 0.3$, $D = 0.2$, $w = 0.216$, $C(r) = 0.146\, r^2$, and $\bar{s} = 0.55$; among households who can evade, A_i is logistic with mean 1.8 and scale 0.05. Panel (a): filled circles are stable steady states ($r^* = 0.88$ and 0.15, $L = 0.39$ and 0.00); the hollow circle is the threshold ($r^* = 0.63$, $L = 1.97$). In the low state every household who can evade does. Panel (b): reliability after a one-time outage shock in period 2 to a network in the high steady state.