UC Santa Cruz

UC Santa Cruz Previously Published Works

Title

Game-theoretic models of electricity theft detection in smart utility networks:
Providing new capabilities with advanced metering infrastructure

Permalink

<https://escholarship.org/uc/item/3658w184>

Journal

IEEE Control Systems, 35(1)

ISSN

1066-033X

Authors

Amin, S
Schwartz, GA
Cardenas, AA
et al.

Publication Date

2015-02-01

DOI

10.1109/MCS.2014.2364711

Peer reviewed

IMAGE LICENSED BY GRAPHIC STOCK

Game-Theoretic Models of Electricity Theft Detection in Smart Utility Networks

SAURABH AMIN,
GALINA A. SCHWARTZ,
ALVARO A. CÁRDENAS,
and S. SHANKAR SASTRY

PROVIDING NEW CAPABILITIES WITH ADVANCED METERING INFRASTRUCTURE

Digital Object Identifier 10.1109/MCS.2014.2364711
Date of publication: 19 January 2015

The smart grid refers to the modernization of the power grid infrastructure with new technologies, enabling a more intelligently networked automated system with the goal of improving efficiency, reliability, and security, while providing more transparency and choices to electricity customers. A key technology being widely deployed on the consumption side of the grid is advanced metering infrastructure (AMI).

AMIs refer to the modernization of the electricity metering system by replacing old mechanical meters with *smart meters*. Smart meters are new embedded devices that provide two-way communication between the utility and the customer. These devices have advanced communication and computational capabilities, with the potential to enable new functionalities, such as improved service choices and transparency. Distribution utilities (or distributors) using AMI for the monitoring and billing of electricity consumption can avoid sending their employees to read the meters on site. Importantly, AMIs provide several new capabilities, including monitoring network-wide and individual electricity consumption, faster remote diagnosis of outages (with analog meters, utilities learned of outages primarily by customer call complaints), remote disconnect options, and automated power restoration. AMIs also improve the customers' access to their energy usage information (including the sources of electricity, renewable or otherwise) and promote the implementation of demand-response schemes.

The widespread deployment of smart meters, by necessity, entails installing low-cost commodity devices in physically insecure locations [1], with an expected operational lifetime in the range of several decades. The cost of AMIs range from US\$100 to US\$400 per device, excluding installation and maintenance costs. Hardening these devices by adding hardware coprocessors and tamper-resilient memory might moderately increase the per-unit price of smart meters. However, this can significantly increase the distribution utility's cost of deploying and operating millions of devices. Thus, creating a business case for improving the security of smart-grid deployments is a difficult task for most electric distribution utilities. Consequently, these additions are not considered cost-effective in practice and are not even recommended as a priority [2]. To realize the promise of trusted computing in smart embedded devices, new technologies need to be developed and deployed [3].

Detecting electricity theft has been traditionally addressed by physical checks of tamper-evident seals by field personnel and by using balance meters [4]. Although these techniques reduce unmeasured and unbilled consumption of electricity, they are insufficient. Indeed, tamper-evident seals can be easily defeated [5], and although balance meters can detect that some customers are fraudulent, they cannot identify the culprits exactly. Despite the security vulnerabilities of smart meters, the higher-resolution data collected by them is seen as a promising technology that will complement traditional detection tools. They have clear potential to improve metering, billing and collection processes, and the detection of fraud and unmeasured connections.

ELECTRICITY THEFT IN DISTRIBUTION NETWORKS

Historically, widespread energy theft is characteristic for developing countries, with theft of electricity reaching up

to 50% in some jurisdictions [6]. Common methods of theft range from compromising the physical security of meters to directly connecting loads to electricity distribution lines. Default of payments has been a major problem, due to sub-optimal levels of monitoring and enforcement. The lack of technology and insufficient distributor incentives were the major contributors to this problem.

Nontechnical Losses and Electricity Theft

In general, distribution utilities can incur nontechnical losses due to

- » actions of utility personnel or an operator, that is, administrative losses due to errors in accounting and record keeping
- » customer theft
- » customer nonpayment
- » theft by noncustomers.

The administrative errors can be strategic (that is, intentional) when made for the purpose of assisting customer theft.

For a distribution utility, the nontechnical losses (such as electricity theft, fraud, or uncollected/defaulted bills) contribute to costs. The customers who acquire electricity via stealing or defaulting on their bills obtain the electricity at zero or near-zero prices. Effectively, the electricity consumption of these nonpayers is subsidized by the distribution utility and/or other customers, or in some cases, by subsidies from local taxes. Overall, the consumption of nonpaying entities is paid by the society at large.

The nontechnical losses can be recovered by some combination of

- » imposing higher electricity tariffs on other (paying) customers
- » decreasing the profit margins of the distributor(s)
- » distributing the burden to the entire society, for example, by increasing taxes.

The actual means depend on the security and recovery technologies that are available to the distributor, how much the distributor invests in them, and the regulatory environment. When the distributor bears losses for a prolonged period of time and no regulatory resources exist to recover these losses, the distributor's incentives and capabilities to invest in the network and its maintenance are jeopardized. Such underinvestment negatively affects the long-term efficiency of distribution system. Thus, to improve efficiency of distribution systems, both technological and regulatory means to limit nontechnical losses are desirable.

Technological and Regulatory Solutions

In recent years, basic protective measures such as tamper-evident seals and secure-link communications have been developed for AMIs. Still, they are not enough to prevent successful attacks during the meter lifespan. Security researchers have recently identified cybervulnerabilities in smart meters [7], [8] and were even able to perform rogue remote firmware updates [9]. Notably, hacked smart meters have been used to

steal electricity, resulting in losses of millions for dollars for a single U.S. utility [10]. Malicious insiders and outside hackers with only a moderate level of computer knowledge are likely able to compromise and reprogram meters with low-cost tools and software readily available on the Internet. The report [10] also predicts and conjectures with medium confidence that as smart grid deployments continue, the cyber theft of electricity will also rise. The most likely reasons for this rise are the lower costs of intrusion and high overall financial benefit for both hackers and customers.

Still, in regulated environments, new investments required for the effective deployment and enforcement of technological solutions is possible only when the necessary institutional and regulatory measures are enacted. Examples of required institutional measures include prosecution of fraudulent customers, publicizing violations for sharper public scrutiny, increasing customers' awareness that electricity theft is an identifiable offense, and disconnecting customers for fraud or debts and reconnecting their service only after the remittance of the required payments. Examples of regulatory measures include fixing skewed tariff structures, providing coordination and transparency in distribution operations, and creating mechanisms to improve investments in security upgrades.

AMI-ENABLED ANOMALY DETECTION

Distribution utilities are collecting fine-grained data from their networks, devices, and customers and are developing

the analytical capabilities for improved situational awareness [11]. Meter data management (MDM) vendors are providing analytical services to the utilities to turn their data into actionable information; see Figure 1. An important MDM service is called revenue assurance. It provides data-analytics software to identify suspected electricity theft through the detection and isolation of abnormal consumption trends [12]. Such anomaly detection schemes can become a cost-effective tool to complement the use of balance meters (which are still necessary to detect theft through unauthorized connections to the power distribution lines) and physical checks of tamper-evident seals by utility personnel.

Thus, the MDM system is emerging as a focus of many AMI deployment projects for three reasons. First, it can be easily retrofitted with an existing distribution infrastructure. Second, unlike other security technologies, it does not require the major capital investments needed by other security technologies such as balance meters. Third, the extra security provided by MDMs is a by-product. The main reason MDMs are popular is because they add value due to their data storage and processing capabilities.

Related Work

Early research on the detection of electricity theft focused on the role of a set of trusted balance meters and looked at electricity consumption traces to check the accuracy of meters [13]. Subsequently, the rise of smart meters and the possibility of high-frequency data collection by distribution utilities motivated the study of security of individual meters. Here, the focus was on the detection of abnormal electricity traces that are highly correlated with electricity theft. This work used a variety of machine-learning techniques, including support vector machines and extreme learning machines to identify suspicious energy traces [14]–[16]. More recent work has emphasized the need to consider consumption data anomalies as part of a diagnostic system, with the aim of enabling sensor fusion at the scale of an electricity distribution network, and reduce false positives [17]. Another new line of research focuses on explicitly modeling the objective of an adversary, whose goal is to steal electricity and yet evade the diagnostic system [18]. Here new metrics are proposed for evaluating a class of theft detection schemes in the presence of powerful attackers who can bypass these schemes. A broader picture of the electricity theft problem can be found in a recent survey article [19].

FIGURE 1 Centralized meter data management (MDM). When electric utilities deploy advanced metering infrastructures, they also need to deploy an MDM system in their back end to manage smart meter data storage and analytics (for example, forecasting and anomaly detection). One of the key services offered by popular MDM vendors is called revenue assurance, where data analytics software is used by the utility on the collected meter data to identify possible electricity theft situations and abnormal consumption trends. Leveraging MDM systems to collect indicators of electricity theft is a cost-effective way to complement the use of balance meters and personnel physically checking tamper-evident seals and reporting reprogramming or tampering attempts.

Focus of the Article

The goal of this article is to develop a model that allows investigating distributor monitoring choices when customers are strategic, and a known fraction of consumers engages in stealing. This model allows finding jointly optimal choices of distributors and customers, under the imposition of only mild (aggregate) requirements on distributor information about customer preferences.

More broadly, this article develops a game-theoretic framework that explicitly models the adversarial nature of the electricity theft problem. The model considers both pricing and investment decisions by the distribution utility (that is, the distributor) who faces a population with both genuine and fraudulent customers. All customers derive identical utility from using electricity (preferences) but face different costs. The genuine customers pay their entire bill. They choose how much to consume (equal to the amount billed), depending on their preferences and the price of electricity. The fraudulent customers choose two amounts: 1) the amount for which they will pay (as genuine ones do) and 2) the amount that they will steal. The second choice depends on the probability of detection and on the amount of fine that they pay if detected.

The probabilistic rate at which fraud is successfully detected depends on the diagnostic scheme implemented in the distributor's MDM system. In particular, the performance of a diagnostic scheme is governed by the received operating characteristic (ROC) curve (that is, the relationship between the probability of detection and the probability of false alarm). The probability of detection depends on two factors. First, it depends on the amount of electricity stolen (the probability increases with the amount stolen) and, second, on the level of investment made by the distributor to monitor fraud. More investment by the distributor increases the probability of detection. The distributor chooses how much to invest in fraud monitoring and the price per unit quantity of billed electricity.

This article considers the two environments: the unregulated monopoly and perfect competition. In both cases, the game is a leader-follower game, in which the distributor (leader) chooses first, given a known fraction of fraudulent customers. The article computes equilibria for both cases (games) and compares the level of effort for unregulated monopoly and perfect competition. In both games, customers make their choices after they learn the pricing (tariff) and distributor level of investment in monitoring fraud. The distributor's operational costs are affected by the level of investment in fraud monitoring, in addition to the traditional cost of providing the total quantity of electricity demanded by the population. Thus, the distributor's revenue function aggregates the revenue generated from billed electricity and the *expected* fines collected from fraudulent customers (when detected). The distributor's profit, that is, revenue net costs, depends on both the level of investment and the per unit price offered to customers.

The chosen level of investment and the customers' equilibrium consumption levels determine the diagnostic scheme's operating point on the ROC curve and hence the distributor's efficiency in recovering costs by monitoring and collecting fines. For a given distributor choice of price and level of investment, the customers' response functions are derived. Finally, the optimal choices for the case when the distributor is an unregulated monopolist are compared with the choices in the case of perfect competition. Although perfect competition is seldom achieved in electricity distribution systems, it offers a standard benchmark. The case of regulated monopolist is also briefly introduced.

The results indicate that for environments with a monopolist distributor, the fraudulent customers are likely to steal more electricity (in equilibrium) as the distributor's level of effort in monitoring and enforcement decreases. The stealing level also increases as the per-unit usage charge of electricity increases or the fine exercised by the distributor decreases. For a given marginal cost of monitoring and fixed fraction of fraudulent customers, the distributor's equilibrium profit increases with level of investment. For cases when optimal investment levels are lower, the monopolist distributor chooses a higher per-unit price of electricity. Such cases are relevant when the level of investment is constrained by limits on the false alarm rate.

The distributor's collection efficiency in equilibrium shows interesting behavior. Specifically, for lower (respectively, higher) values of fines, the collection efficiency of the distributor increases (respectively, decreases) with his level of investment. This indicates the necessity of certain regulatory impositions that can enable socially desirable levels of collection efficiency. The framework presented in this article can be used to compare the optimal level of investment for different regulatory regimes and to design mechanisms to improve the distributor incentives to implement socially optimal monitoring choices.

Although this article does not deal with attack models that have been tested on real AMIs, the proposed game-theoretic framework is motivated by practical attack models, such as rigging the electricity consumption signal via cyber (reprogramming) or cyberphysical means (such as installing a rigged smart meter). Clearly, in response to such threats, the distributor can employ diagnostic schemes to find the fraudulent customers. The proposed game-theoretic framework can help analyze equilibrium customer and distributor choices in scenarios where the assumptions on customer utilities and distributor's profit function are applicable.

MODELING CUSTOMER PREFERENCES

Let $\mathcal{N} = \{1, \dots, n\}$ denote the population of customers that are served by the distributor. All customers belong to the same socioeconomic class and thus have the same valuation or preference of electricity consumption. However, the security level of individual meters varies across the population. For simplicity, assume that each customer has a

meter that is either of *low* or *high* security. Customers whose meters have low (respectively, high) security levels are respectively referred as type-f (respectively, type-g) customers. On one hand, the meters of type-f customers possess certain security vulnerabilities and/or configuration defects that can be exploited for economic gain. Thus, type-f customers are more likely to steal electricity. When they are successful in stealing, the distributor does not fully recover the cost of the electricity used and incurs non-technical losses. On the other hand, the meters of type-g customers are harder to exploit, either because of the due care taken during their manufacturing and installation process, or due to the customers' lack of technological knowledge required to exploit the high-security-level meters. Each type-g customer fully pays for the electricity consumed. Thus, type-f and type-g customers can be called "fraudulent" and "genuine," respectively.

Let $N_f \subset N$ and $N_g = N \setminus N_f$ denote the sets of these customer types, and let λ be the fraction of type-f customers, that is, $\lambda = |N_f| / |N|$. The distributor (a monopolist) knows the fraction λ but cannot distinguish between type-f and type-g customers. An estimate of λ is

$$\lambda = \frac{Q_u}{Q_T}, \quad (1)$$

where Q_u denotes the total unrecovered quantity due to stealing by type-f customers when the level of investment in monitoring fraud is negligibly small, and Q_T is the total quantity of electricity provided by the distributor. To reduce the stealing losses, the distributor deploys a diagnostic scheme and has to incur the costs of monitoring (for distinguishing between fraudulent and genuine customers) and enforcement (for recovery of money from customers who are successfully detected as fraudulent).

Genuine Customers

Suppose that each type-g customer has the utility function

$$U_g = u(q_g) - T(q_g) \quad [\text{Secure AMIs}], \quad (2)$$

where the function $u(\cdot)$ (assumed to be same for customers) satisfies $u(0) = 0$, $u'(q) > 0$, and $u''(q) < 0$, that is, there is a decreasing marginal utility of electricity consumption. If the distributor offers a tariff schedule $T(\cdot)$, a type-g customer chooses the *expected* quantity q_g and pays $T(q_g)$ to the distributor. Assume $T(\cdot)$ is increasing in q_g . In general, the distributor can offer a nonlinear tariff schedule. The customer surplus is

$$v_g \equiv \max_{q_g \geq 0} [u(q_g) - T(q_g)], \quad (3)$$

and the first-order-condition is $u'(q_g) - T'(q_g) = 0$. Consider a two-part tariff schedule given by $T(q_g) = A + pq_g$. Here A is a fixed charge that can be interpreted as a connection fee, and p is constant per unit price (usage charge). For the pur-

pose of analytical derivations, this article assumes that customer preference is given by a square-root function $u(q_g) = 2\sqrt{q_g}$. Under this assumption, the chosen consumption and optimal surplus of a type-g customer becomes

$$q_g(p) = \frac{1}{p^2}, \quad v_g^*(p) = \left(\frac{1}{p} - A\right). \quad (4)$$

The customer surplus decreases as the distributor charges more per unit price p . Of course, the fixed charge A is constrained by $A < (p)^{-1}$. Since $|N_g| = n(1 - \lambda)$, the total quantity of electricity consumed by genuine customers is

$$Q_g(p) = \frac{n(1 - \lambda)}{p^2}. \quad (5)$$

Fraudulent Customers

Consider the following utility function for each type-f customer

$$U_f = u(q_f^B + q_f^S) - T(q_f^B) - \rho_D(\ell, q_f^S)F^r(q_f^S) \quad [\text{Insecure AMIs}], \quad (6)$$

where $u(\cdot)$ and $T(\cdot)$ are same as in (2), q_f^B and q_f^S , respectively, denote the *expected* billed and stolen (or unpaid) quantities for a type-f customer, $\rho_D(\ell, q_f^S)$ is the probability that a fraudulent customer is detected when distributor's level of investment in monitoring of fraud is $\ell \in \mathbb{R}_+$, and $F^r(\cdot)$ is the fine schedule exercised by the distributor upon successful fraud detection. Consistent with common practice of regulating distributors, the $F^r(\cdot)$ schedule is increasing in q_f^S . It is fixed by a regulating entity and is known to all customers and the distributor. The probability of detection increases with ℓ and q_f^S . If the stolen electricity q_f^S were perfectly detectable, the customer would pay $F^r(q_f^S)$ to the distributor. However, under imperfect detection, the distributor only recovers for $\rho_D(\ell, q_f^S)q_f^S < q_f^S$ via fine (in expectation), and the remaining quantity is stolen. The customer surplus is

$$v_f \equiv \max_{q_f^B \geq 0, q_f^S \geq 0} [u(q_f^B + q_f^S) - T(q_f^B) - \rho_D(\ell, q_f^S)F^r(q_f^S)], \quad (7)$$

and the first-order conditions (FOCs) are

$$\begin{aligned} \partial_{q_f^B} [u(q_f^B + q_f^S)] &= T'(q_f^B), \\ \partial_{q_f^S} [u(q_f^B + q_f^S)] &= \partial_{q_f^S} [\rho_D(\ell, q_f^S)F^r(q_f^S)]. \end{aligned}$$

That is, a small increase in total quantity ($q_f = q_f^B + q_f^S$) consumed by a type-f customer generates a marginal surplus $u'(q_f)$ equal to marginal payment $T'(q_f^B)$ (respectively, *expected* marginal fine $\partial_{q_f^S} [\rho_D(\ell, q_f^S)F^r(q_f^S)]$) for a small increase in the billed (respectively, stolen) quantity.

Again consider a two-part tariff schedule $T(q_f^B) = A + pq_f^B$ and a similar fine schedule $F^r(q_f^S) = F + p_f q_f^S$. Assuming square-root customer preferences $u(q_f) = 2\sqrt{q_f}$, the FOCs imply that quantities q_f^B and q_f^S satisfy

$$\rho_D(\ell, q_f^S) p_f + \partial_{q_f^S} [\rho_D(\ell, q_f^S)] [F + p_f q_f^S] = p, \quad q_f^B = \frac{1}{p^2} - q_f^S. \quad (8)$$

Thus, all customers have the same average total consumption, that is,

$$q_f = q_f^B + q_f^S = q_g = \frac{1}{p^2}. \quad (9)$$

This results from the assumption that each customer's valuation of the total quantity of electricity does not depend on customer type, that is, $u(\cdot)$ is the same for both type-g and type-f customers. For the case when, upon detection, the fraudulent customer pays a fixed fine F that is much larger than $p_f q_f^S$, that is, $F'(\cdot) \approx F$, the FOCs (8) simplify to

$$\partial_{q_f^S} [\rho_D(\ell, q_f^S)] = \frac{p}{F}, \quad q_f^B = \frac{1}{p^2} - q_f^S. \quad (10)$$

DIAGNOSTIC SCHEME

The distributor employs a diagnostic scheme that requires a level of investment $\ell \in \mathbb{R}_+$. The cost of stolen electricity is partially recovered by imposing the predetermined fine schedule on the customers who are correctly diagnosed to be fraudulent. Specifically, the expected fine collected from a type-f customer is $\rho_D(\ell, q_f^S) F'(\cdot)$, where the probability of detection $\rho_D(\ell, q_f^S)$ is a property of the diagnostic scheme employed by the distributor.

Due to the inherent variability of meter measurements received from genuine and fraudulent customers, a high value of $\rho_D(\ell, q_f^S)$ also entails a high probability of false positive (or false alarm), denoted as ρ_F . The statistical decision theory models this tradeoff between ρ_D and ρ_F values as the ROC curve. That is, a diagnostic scheme with higher ρ_D will result in a higher ρ_F , and vice versa. Let ρ_D be concave increasing in ρ_F . It is reasonable to expect that the probability of false alarm ρ_F increases as the distributor increases investment ℓ in monitoring of fraud, that is, $\rho_F(\ell) \in [0, 1]$ is increasing in $\ell \in \mathbb{R}_+$. Furthermore, assume $\rho_F(\cdot)$ to be a continuously differentiable and invertible function, and $\rho_F(0) = 0$.

Note that ℓ models the distributor's level of effort in monitoring and enforcement for a *fixed* diagnostic scheme. Thus, ℓ can be viewed as the distributor's willingness to act on the alerts of the diagnostic scheme by investigating customer meters and/or their communication with the centralized MDM system for security compromises. It is important to note that, in the current game-theoretic framework, a higher level of ℓ *does not* imply a better diagnostic performance.

Another effect of the distributor's level of investment can be a better diagnostic scheme. In this case, a higher level of ℓ will improve the tradeoff between the probability of missed detection ($1 - \rho_D$) and probability of false alarm ρ_F , that is, as ℓ increases, the ROC curve will shift toward left. The models presented in this article do not consider this effect. See "Practical Evaluation of Electricity Theft

FIGURE 2 The computation of detection probability ρ_D and false alarm probability ρ_F .

Detection Schemes" for a discussion on the evaluation of different electricity theft detection schemes.

Detecting Electricity Theft

A practical implementation of the diagnostic scheme is outlined below. The distributor collects a time series of electricity consumption measurements y_t^i reported by the i th customer's meter, where $i \in \mathcal{N}$ and $t \in \{1, \dots, K\}$. Here K denotes the number of time intervals in one billing cycle. For simplicity, the time interval between subsequent measurements, $(y_{t+1}^i - y_t^i)$, is assumed to be fixed and known a priori. Assume that the meter measurements y_t^i are independent and drawn from identically distributed random variables, $Y_1^i, \dots, Y_K^i$, where y_t^i is the realization of the random variable Y_t^i . Under this assumption, the reported meter measurements are independently and identically distributed (iid). Let the probability density function (pdf) of meter measurements, y_t^i , of type-g (respectively, type-f) customers be denoted by f_g (respectively, f_f). For the given customer preferences, the expected value of Y_t^i for a type-g (respectively, type-f) customer is q_g (respectively, q_f^B); see Figure 2. Consider the two *simple* hypotheses

$$\begin{aligned} H_g: Y_1^i, \dots, Y_K^i &\stackrel{\text{iid}}{\sim} f_g, & \mathbb{E}[Y_t^i] &= q_g, \\ H_f: Y_1^i, \dots, Y_K^i &\stackrel{\text{iid}}{\sim} f_f, & \mathbb{E}[Y_t^i] &= q_f^B < q_g, \end{aligned}$$

where q_g and q_f^B satisfy (4) and (8), respectively. The likelihood-ratio test takes the form

$$\frac{\prod_{t=1}^K f_g(y_t^i)}{\prod_{t=1}^K f_f(y_t^i)} \stackrel{H_g}{\underset{H_f}{\gtrless}} \gamma, \quad (11)$$

where f_g (respectively, f_f) denotes the pdfs of meter measurements y_t^i of type-g (respectively, type-f) customers, and γ reflects the distributor's tradeoff between the probability of missed detection ($1 - \rho_D$) and the probability of false alarm ρ_F . The diagnostic scheme employed by the distributor uses the previous K meter measurements and a threshold value τ to detect fraudulent customers. That is, a customer i is classified as fraudulent if $\sum_{t=1}^K y_t^i < \tau$, and

Practical Evaluation of Electricity Theft Detection Schemes

A significant practical challenge for designing accurate electricity theft detectors is dealing with an adversarial environment where the attacker can design fake electricity consumption traces that will not be detected by the detector.

In the game-theoretic formulation proposed in this article, the fraudulent customers (attackers) choose q_f^S such that the marginal payment for a small increase in the billed quantity of electricity is equal to the expected marginal fine for the small increase in q_f^S ; see FOCs (8). This holds for the case when customer valuations for total electricity consumed are the same for fraudulent and genuine customers. Essentially, the choice of q_f^S determines how different the distribution of meter measurements sent by type-f customers will be from the distribution of measurements sent by type-g customers.

An alternative problem formulation is to find a distribution of compromised meter measurements that maximizes the quantity of stolen electricity subject to the constraint that individual meter measurements will be undetected with high probability. This is the basis of recent work [18], where electricity traces of 108 residential customers were obtained and analyzed from a real AMI deployment.

The design of an optimal attack signal was based on the adversarial model that maximized stolen electricity without being detected. In other words, the probability that a compromised meter measurement would be detected by any of these algorithms is negligibly small. However, to remain undetected, the attacker must place a cap on the maximum amount of electricity that can be stolen.

A new performance metric was proposed in [18]. Here, in contrast to traditional ROC curves, the detection scheme's performance was evaluated based on how the distributor's total loss due to undetected attacks, including the loss of revenue from net stolen electricity, varies with the probability of false alarm. Figure S1 provides a comparison of five detection schemes according to this performance metric. Experimental results indicate superior performance of ARMA-GLR, as it is the test that minimizes the amount of stolen electricity among all possible undetected attacks. Additionally, [18] addresses issues related to concept drift (the fact that the statistical distribution of electricity consumption changes with time) and with training data set poisoning attacks (where the attacker can feed a profiling algorithm malicious data).

FIGURE S1 The performance evaluation of several detection schemes: auto-regressive moving average with a generalized likelihood ratio (ARMA-GLR) test, a simple average consumption test, the nonparametric cumulative sum (CUSUM) algorithm, an exponential weighted moving average (EWMA) detector, and an outlier detection algorithm called local outlier factor (LOF).

The anomaly detection schemes proposed in [18] should be used as part of a more comprehensive electricity theft detection system. A limitation of this approach is the assumption of the adversary model, where the fraudulent customers continue to use electricity as genuine customers but will try to send lower meter readings to the utility to minimize their payments. This model, however, does not cover an attacker that increases electricity consumption but sends signals corresponding to their previous consumption. This type of attack can be detected by adding new balance meters and having frequent site inspections.

Some recent research has also focused on improving the privacy of electricity customers [20], [21]. The idea of these schemes is to shape the electricity usage signal to prevent inferences that can be made with nonintrusive load monitors [22]. It is still not clear if any of these systems will ever see a significant deployment; however, because these schemes do not change the total consumption of electricity, the game-theoretic framework presented in this article can still be applied to analyze equilibrium customer choices.

genuine otherwise. Figure 2 illustrates the resulting probabilities ρ_D and ρ_F

$$\rho_D = \int_0^\tau f_f(y) dy, \quad \rho_F = \int_0^\tau f_g(y) dy. \quad (12)$$

For the purpose of analytical tractability, assume that the meter measurements received by the distributor for type-g and type-f customers follow exponential distributions with parameters $1/q_g$ and $1/q_f^B$, $q_g > q_f^B$, respectively,

$$f_g(y^i) = \frac{1}{q_g} \exp\left(-\frac{y^i}{q_g}\right), \text{ and } f_f(y^i) = \frac{1}{q_f^B} \exp\left(-\frac{y^i}{q_f^B}\right). \quad (13)$$

Admittedly, the assumption that the meter measurements for type-g (respectively, type-f) customers are iid exponentially distributed with parameters $1/q_g$ (respectively, $1/q_f^B$) might not be consistent with real-world consumption patterns. However, this assumption greatly eases the development of analytical expressions for the equilibrium choices of the distributor and customers and helps to highlight the

FIGURE 3 Received operating characteristic curves (ρ_D versus ρ_F) for different levels (that is, $q_f^S/q_g \times 100\%$) of stealing by type-f customers.

interplay between the diagnostic scheme's performance and the customers' optimal choices.

Substituting (13) into the likelihood-ratio test (11), and taking the logarithm on both sides, yields the simplification

$$\sum_{i=1}^K y_i^t \geq_{H_f} \tau, \quad (14)$$

where $\tau = ((q_g - q_f^S)/(q_g q_f^S)) \ln(\gamma(q_g/q_f^S)^K)$. Under the assumption (13), ρ_D and ρ_F become

$$\rho_D = 1 - \exp\left(-\frac{\tau}{q_f^B}\right), \quad \rho_F = 1 - \exp\left(-\frac{\tau}{q_g}\right).$$

For a given ℓ , the threshold τ can be expressed as a function of $\rho_F(\ell)$

$$\tau(\ell) = -q_g \ln(1 - \rho_F(\ell)). \quad (15)$$

Using this expression, the probability ρ_D can be expressed as

$$\rho_D(\ell, q_f^S) = 1 - [1 - \rho_F(\ell)]^{\left(\frac{q_g}{q_f^B}\right)} = 1 - [1 - \rho_F(\ell)]^{\left(\frac{1}{1 - p^2 q_f^S}\right)}, \quad (16)$$

where the second equality follows from (9). Equation (16) represents the ROC curve of distributor's diagnostic scheme.

Figure 3 plots the ROC curve for different fractions q_f^S/q_g . Note that when the stolen quantity is negligible ($q_f^S \rightarrow 0$), the diagnostic scheme uses random guessing ($\rho_D \rightarrow \rho_F$), and as the quantity of stolen electricity reaches close to the average consumption ($q_f^S \rightarrow q_g$), fraud is detected with almost certainty ($\rho_D \rightarrow 1$).

Optimal Choices of Type-f Customers

For given distributor choices of price p and level of investment ℓ , the optimal response of type-f customers can be

Table 1 $\alpha(p, \ell)$ for different $\rho_F(\ell)$.

$\rho_F(\ell)$	Equation to solve for α
0.1	$Fp\alpha^2 \exp(-0.105\alpha) = 9.490$
0.25	$Fp\alpha^2 \exp(-0.287\alpha) = 3.476$
0.50	$Fp\alpha^2 \exp(-0.693\alpha) = 1.442$
0.75	$Fp\alpha^2 \exp(-1.386\alpha) = 0.721$
0.90	$Fp\alpha^2 \exp(-2.302\alpha) = 0.434$

derived from (16) and (10). The following definition is introduced for notational convenience

$$\alpha \equiv \frac{q_g}{q_f^B} = \frac{1}{1 - p^2 q_f^S}.$$

Combining the ROC curve (16) with FOCs (10) provides that α satisfies

$$\alpha^2 (1 - \rho_F(\ell))^\alpha \ln(1 - \rho_F(\ell)) = -\frac{1}{Fp}, \quad (17)$$

from which the closed-form solution

$$\alpha(p, \ell) = \frac{2W\left(\frac{1}{2}\sqrt{-\frac{\ln(1 - \rho_F(\ell))}{Fp}}\right)}{\ln(1 - \rho_F(\ell))} \quad (18)$$

can be derived. In (18), W is the *product logarithm* function. It is defined as the inverse function of $f(W) = We^W$. As an alternative to (18), Table 1 lists nonlinear equations that can be solved to obtain $\alpha(p, \ell)$ for given F, p , and $\rho_F(\ell)$.

Assume a square-root customer preference function $u(q_f) = 2\sqrt{q_f}$ and a fixed fine schedule $F'(\cdot) \approx F$. Then, for any chosen $\rho_F(\ell)$ (or, equivalently, ℓ) and per unit price p of the distributor, the optimal consumptions q_f^S and q_f^B of the type-f customers are

$$q_f^B(p, \ell) = \frac{1}{p^2 \alpha(p, \ell)}, \quad q_f^S(p, \ell) = \frac{1}{p^2} \left(1 - \frac{1}{\alpha(p, \ell)}\right), \quad (19)$$

where $\alpha(p, \ell)$ is given by (18). The optimal surplus of a type-f customer becomes

$$v_f^*(p, \ell) = \frac{1}{p} \left(2 - \frac{1}{\alpha(p, \ell)}\right) - (A + F) + F(1 - \rho_F(\ell))^{\alpha(p, \ell)}. \quad (20)$$

Note that when the level of investment is negligible ($\ell \rightarrow 0$), the ability of a diagnostic scheme to detect fraud is reduced ($\rho_F \rightarrow 0$ and $\rho_D \rightarrow 0$) and type-f customers' stolen quantity increases ($q_f^S \rightarrow q_g$ and $q_f^B \rightarrow 0$). For this case, (20) simplifies to

$$v_f^*(p, 0) = \left(\frac{2}{p} - A\right),$$

which is greater than the type-g customer surplus $v_g^*(p)$; see (4).

Figure 4 provides contour plots for a range of fractions, $q_f^s(p, \ell)/q_g(p)$, using the FOC (17). The x-y coordinates respectively reflect the distributor's investment levels ℓ (represented by the corresponding false alarm rates $\rho_F(\ell)$), and the usage charge p (scaled by regulator chosen F). For a given $\rho_F(\ell)$ (respectively, F_p), the type-f customers steal less as p (respectively, ℓ) increases. That is, fraudulent customers are likely to steal more as the distributor's level of investment in monitoring decreases, or as the usage charge or fine increases. Note that the contour plots entail no information on the distributor's optimal choices of p and ℓ . As explained later, the distributor will choose p and ℓ to maximize profit, given the customers' best responses.

Finally, since $|N_f| = n\lambda$, the total consumption by fraudulent customers becomes

$$Q_f(p) = \frac{n\lambda}{p^2}. \quad (21)$$

From (5) and (21), the total quantity provided by the distributor is

$$Q_T(p) \equiv Q_g(p) + Q_f(p) = \frac{n}{p^2}.$$

Under the stated assumptions, $Q_T(p)$ does not depend on λ or ℓ , and decreases with p . The total stolen (or unrecovered) quantity is

$$Q_f^s(p, \ell) = n\lambda(1 - \rho_D(\ell, q_f^s(p, \ell)))q_f^s(p, \ell), \quad (22)$$

where $(1 - \rho_D(\ell, q_f^s(p, \ell)))q_f^s(p, \ell)$ is the unrecovered quantity from a type-f customer. The distributor's collection efficiency can be expressed as

$$\eta(p, \ell) \equiv 1 - \frac{Q_f^s(p, \ell)}{Q_T(p)}, \quad (23)$$

FIGURE 4 Isolines of different levels of stealing by type-f customers, that is, contour plots of $q_f^s(p, \ell)/q_g(p)$ for 30, 40, 50, 60, and 70%. Environments with lower ℓ and p are likely to result in higher levels of stealing.

where Q_f^s is given by (22). Note that as $\ell \rightarrow 0$, the entire consumption of fraudulent customers is unrecovered by the distributor ($Q_f^s \rightarrow Q_f$). For this case, the quantity Q_U in (1) can be computed as

$$Q_U := Q_f^s(p, 0) = \frac{n\lambda}{p^2}, \quad (24)$$

and the collection efficiency achieves its minimum value $\eta(p, 0) = 1 - \lambda$.

Distribution of Meter Measurements

The analysis presented in this article is also extensible to other parametric assumptions on the pdfs of meter measurements. For example, consider a more realistic assumption that the probability density function of meter measurements of type-g (respectively, type-f) customers follow a normal (Gaussian) distribution with mean parameter q_g (respectively, q_f^B), and *known* variance, $\sigma^2 > 0$. Again, the assumption $q_g > q_f^B > 0$ distinguishes the two distributions. The likelihood ratio test for this case is similar to (14), but with the threshold value $\tau_{\text{ndf}} = (\sigma^2/q_f^s) \ln(\gamma_{\text{ndf}}) + K(q_g - q_f^s/2)$, where γ_{ndf} is the parameter that governs the tradeoff between ρ_D and ρ_F for the case of a normal distribution. After solving for the value of τ_{ndf} in terms of $\rho_F(\ell)$ the ROC curve can be expressed as

$$\rho_D(\ell, q_f^s) = 1 - Q\left(\sqrt{K\left(\frac{q_f^s}{\sigma}\right)^2} - Q^{-1}(1 - \rho_F(\ell))\right), \quad (25)$$

where $Q(z) := \int_{u \geq z} (1/2\pi) \exp(-u^2/2) du$ is the tail probability of the standard normal distribution (mean zero and unit variance). The Q function is invertible. Equation (16) can be replaced by (25) for normally distributed meter measurements; however, this case is not pursued further in this article for the sake of brevity.

It is important to note that, for both exponential and normal distribution assumptions, the type-g (respectively, type-f) customers influence the distributions of their meter readings only by choosing the mean parameters of the pdf, which characterizes their consumption patterns. In other words, under the stated assumptions, the customers' choices only influence the mean parameter. They do not alter the form or high-order moments of the probability distribution of meter measurements. Extensions such as relaxing these assumptions on customer choices and including a broader class of meter distributions in the current game-theoretic framework are a part of future work.

Optimal Surplus of Type-g Versus Type-f Customers

Consider combinations of distributor choices p and ℓ for which the optimal surplus of a type-f customer v_f^* exceeds the type-g customer's optimal surplus v_g^* , given that both types of customers respond optimally. It is expected that such combinations of p and ℓ favor higher levels of electricity

The proposed game-theoretic framework is motivated by practical attack models, such as rigging the electricity consumption signal via cyber (reprogramming) or cyberphysical means.

theft. In other words, a favorable condition for a type-f customer to remain fraudulent is

$$v_i^* \geq v_g^*. \quad (26)$$

Customers are indifferent between types when $v_i^* = v_g^*$. Using (4) and (20), condition (26) becomes

$$(1 - F_p)\alpha(p, \ell) + F_p\alpha(p, \ell)(1 - \rho_F(\ell))^{\alpha(p, \ell)} \geq 1. \quad (27)$$

Figure 5 marks the type-f customers' favorable region of fines and investment levels [that is, condition (26) holds] for $\alpha = 2.5$ (that is, q_f^S is 60% of q_g), and for $\alpha = 3.0$ (that is, q_f^S is 66.6% of q_g), given that both types of customers respond optimally to the distributor's choices. The points of intersection of the contour plots for 60% and 66.7% stealing levels (similar to Figure 4) with the corresponding favorable regions are also depicted. It can be checked that when $\alpha \in (2.275, 3.05)$, there exists at least one distributor choice of (p, ℓ) such that the FOC (17) and condition (27) are satisfied; this corresponds to q_f^S in the range 56–67% of q_g . Note that although the consumer types are assumed to be fixed in the current model, in reality, they could switch between the types, and if for a given fine level, the distributor's investment level is in type-f customers' favorable region, more stealing is expected.

Other Extensions for Modeling Customer Preferences

While the proposed game-theoretic model considers a population of only two customer types (type-f and type-g) and identical customer preferences $u(\cdot)$, the analysis presented here is extensible to a more general case of multiple security levels and heterogeneous customer preferences. The following two generalizations can be pursued within the current framework.

- 1) The distributor faces a detection problem where there are multiple customer types with the same preference and, thus, the same expected total consumption. The set of customer types is $\mathcal{L} = \{1, \dots, l\}$, where $l > 2$. This set covers the range of AMI security levels that are deployed in the population. For each $j \in \mathcal{L}$, the size of $\mathcal{N}_j$ is known to the distributor. Let λ_j denote the fraction of j -type customers in the total population: $\lambda_j = N_j / N$. Then, $\sum_{j \in \mathcal{L}} \lambda_j = 1$. Analogous to (2) and (6), type- j customer chooses q_j and q_j^S to maximize the objective

$$U_j = u(q_j) - T(q_j - q_j^S) - \rho_D(\ell, q_j^S) F_j'(q_j^S) \quad [\text{Security level } j], \quad (28)$$

where $F_j'(\cdot)$ denotes the fine schedule for customers with meter security level $j \in \mathcal{L}$.

- 2) In addition to the AMI security levels, the customers also differ in their preferences of consuming electricity. Here the population $\mathcal{N}$ can be subdivided into subpopulations or categories $\mathcal{N}^1, \dots, \mathcal{N}^m$, where each category $\mathcal{N}^k$ consists of customers with identical preferences $u_k(\cdot)$ (for example, the same socioeconomic class). However, the preferences, and thus the expected total consumption, differ across categories. The distributor can determine preference type of each customer based on known parameters (such as demographics and electricity profiling of the household). Similar to (28), a customer with AMI security level $j \in \mathcal{L}$ and preferences $u_k(\cdot)$ chooses total consumption q_{jk} and stealing quantity q_{jk}^S to maximize the objective

FIGURE 5 Isolines of 60% and 67% stealing levels and favorable regions in which type-f customers obtain higher surplus, relative to type-g customers. For $F_p > 1.1$ and $\rho_F(\ell) < 0.265$, the (solid-boxed) isoline for 60% stealing level is inside the region bounded by the dashed-boxed line, that is, the region where type-f customers obtain higher surplus, relative to type-g customers. The corresponding ranges for 67% stealing level, that is, the intersection of the (solid-oval) isoline with the (dashed-oval) favorable region boundary, are $F_p > 0.9$ and $\rho_F(\ell) < 0.34$.

In recent years, basic protective measures such as tamper-evident seals and secure-link communications have been developed for AMIs.

$$U_{jk} = u_k(q_{jk}) - T_k(q_{jk} - q_{jk}^S) - \rho_D(\ell, q_{jk}^S) F_j^r(q_{jk}^S) \quad (29)$$

[Security level j , preference type k],

where $T_k(\cdot)$ denotes the tariff schedule for customers with preferences $u_k(\cdot)$.

MODELING THE COSTS OF THE DISTRIBUTION UTILITY

Monopolist Distributor

For the quantity Q_T provided by the distributor, let revenue $R_\lambda(p, \ell)$ be the total revenue, when tariff schedule $T(q_f^B) = A + pq_f^B$ is offered and fine schedule $F^r(\cdot)$ is implemented to recover the quantity $n\lambda q_f^S \rho_D(\ell, q_f^S)$ from fraudulent customers. Here the notation $R_\lambda(p, \ell)$ emphasizes the dependence of revenue on the distributor's choice for the variables p (the per unit price) and ℓ (the level of fraud monitoring) when facing λ_n fraudulent customers. The following analysis considers $F^r(q_f^S) = F$ (that is, $p_f \approx 0$). The total revenue is the sum of revenues generated from genuine and fraudulent customers

$$R_\lambda(p, \ell) = n(1 - \lambda) \left[A + \frac{1}{p} \right] + n\lambda \left[A + p \left(\frac{1}{p^2} - q_f^S \right) + F \rho_D(\ell, q_f^S) \right] = n[A + pq_g(p) + \lambda(-pq_f^S + F \rho_D(\ell, q_f^S))]. \quad (30)$$

There are several different costs that the distributor pays. These costs include detection, monitoring, and recovery costs, that is, the costs of collecting fines. Consider the following two main operational costs to the distributor:

- i) For providing the electricity to meeting the total demand Q_T in each billing period, the distributor faces the cost $C(Q_T)$, where C is an increasing function of Q_T .
- ii) For a level ℓ of investment in fraud monitoring, the distributor faces a cost $\Psi(\ell)$, where $\Psi'(\ell) > 0$.

The following costs are not included in the current model:

- iii) the cost of deploying secure AMIs to ensure that a fraction $(1 - \lambda)$ of the population is type-g customers and the cost of security upgrades of insecure AMIs of type-f customers
- iv) the cost of false alarms which penalizes the distributor for higher false positive rates.

Costs iii)–iv) can be readily introduced in the current model. Their inclusion will likely change the equilibrium results. Still, the results presented here can be justified for environments when iii) can be considered as a sunk cost or

a subsidy by the regulator and iv) is relatively small, that is, when the distributor does not have to pay a high penalty for false positives.

For the sake of simplicity, consider linear cost of providing $C(Q_T) = cQ_T$, $c > 0$, and a linear cost of monitoring fraud $\Psi(\ell) = n\psi\ell$, where $\psi > 0$. The average (per customer) profit for an unregulated monopolist is

$$\pi_\lambda^m(p, \ell) \equiv \frac{\Pi_\lambda^m(p, \ell)}{n} = R_\lambda(p, \ell) - C(Q_T) - \Psi(\ell) = A + (p - c)q_g(p) + \lambda(-pq_f^S + F \rho_D(\ell, q_f^S)) - \psi\ell, \quad (31)$$

where the superscript m on π_λ emphasizes the monopolist profit. The problem of choosing optimal (p, ℓ) that maximizes the distributor's profit becomes

$$\pi_\lambda^{m*} = \max_{p \geq 0, \ell \geq 0} [(p - c)q_g(p) + \lambda(-pq_f^S(p, \ell) + F \rho_D(\ell, q_f^S(p, \ell))) - \psi\ell], \quad (32)$$

subject to (16),

the ROC curve of the diagnostic scheme,

(4) and (17)–(19), the optimal customer response

$$q_g(p), q_f^B(p, \ell), q_f^S(p, \ell),$$

$v_f \geq 0$, $v_g \geq 0$, nonnegative customer valuations.

The optimization problem (32) can be solved by ignoring the constraints initially, but verifying them ex post. The distributor FOCs with respect to p and ℓ are

$$\partial_p((p - c)q_g(p)) + \lambda \partial_p(-pq_f^S(p, \ell) + F \rho_D(\ell, q_f^S(p, \ell))) = 0 \quad (33)$$

$$\lambda \partial_\ell(-pq_f^S(p, \ell) + F \rho_D(\ell, q_f^S(p, \ell))) - \psi = 0. \quad (34)$$

Taking into account the optimal customer responses, these FOCs can be simplified. In particular, rewriting FOC (33) and using the FOC for type-f customers (8)

$$\begin{aligned} q_g(p) + (p - c)q_g'(p) + \lambda \underbrace{\partial_{q_f^S}(-pq_f^S + F \rho_D(\ell, q_f^S))}_{=0, \text{ from (8)}} \partial_p q_f^S &= 0, \\ \Rightarrow \frac{1}{p^2} - \frac{2(p - c)}{p^3} &= 0, \\ \Rightarrow p^* &= 2c. \end{aligned} \quad (35)$$

Thus, in equilibrium, the price p^* is determined solely by customer preferences. With a monopolist distributor, the price reflects a monopolistic markup $(p - c) = c$. From (4), the optimal total consumption for a type-g or type-f customer is

$$q_g^* = q_f^* = \frac{1}{p^*} = \frac{1}{4c^2}.$$

FIGURE 6 A plot of false alarm probability ρ_F versus level of investment ℓ for the model (39).

Next, substituting p^* into (17),

$$\frac{(1 - \rho_F(\ell)) \frac{1}{1 - 4c^2 q_f^S}}{(1 - 4c^2 q_f^S)^2} \ln(1 - \rho_F(\ell)) = -\frac{1}{2Fc}. \quad (36)$$

Similarly, the distributor FOC (33) (with respect to ℓ) can be simplified using (8)

$$\begin{aligned} \underbrace{\partial_{q_f^S}(-p q_f^S + F \rho_D(\ell, q_f^S))}_{=0, \text{ from (8)}} \partial_p q_f^S + F \partial_\ell^1 \rho_D(\ell, q_f^S) &= \frac{\psi}{\lambda}, \\ \Rightarrow \partial_\ell^1 \rho_D(\ell, q_f^S) &= \frac{\psi}{F\lambda}, \end{aligned}$$

where the notation $\partial^1(\cdot, \cdot)$ indicates a partial derivative with respect to the first argument. From (16),

$$\frac{(1 - \rho_F(\ell)) \frac{1}{1 - 4c^2 q_f^S}}{(1 - 4c^2 q_f^S)^2} \frac{\rho_F(\ell)}{(1 - \rho_F(\ell))} = \frac{\psi}{F\lambda}. \quad (37)$$

Solving (36) and (37) gives equilibrium q_f^{S*} and ℓ^* . In fact, q_f^{S*} can be expressed as

$$q_f^{S*} = \frac{1}{p^{*2}} \left[1 - \frac{2\psi c(1 - \rho_F(\ell^*))}{\lambda \rho_F(\ell^*)} \ln\left(\frac{1}{1 - \rho_F(\ell^*)}\right) \right]_+. \quad (38)$$

A reasonable model of a differentiable and increasing function $\rho_F(\cdot)$ is

$$\rho_F(\ell) = \frac{\ell}{1 + \ell}. \quad (39)$$

For this model of false alarm probability (see Figure 6),

$$\rho_F(\ell) = (1 + \ell)^{-2}, \quad (1 - \rho_F(\ell)) = (1 + \ell)^{-1}.$$

The optimal q_f^{S*} and q_f^{B*} are

$$q_f^{S*} = \frac{1}{p^{*2}} \left[1 - \frac{2\psi c(1 + \ell^*)}{\lambda} \ln(1 + \ell^*) \right]_+, \quad q_f^{B*} = \frac{1}{p^{*2}} - q_f^{S*}, \quad (40)$$

where the monopolist distributor's optimal level of investment ℓ^* satisfies

$$p^* \beta \ln\left(\frac{p^*}{F} [\beta(1 + \ell^*)]^2 \ln(1 + \ell^*)\right) + \frac{1}{1 + \ell^*} = 0, \quad (41)$$

FIGURE 7 Optimal (equilibrium) choices of a monopolist distributor: per unit price p^* and ℓ (or, equivalently, $\rho_D(\ell^*)$) for $\beta = 1.0, 1.5, 2.0, 2.5$, and 3.0 and $F = 2.0$.

FIGURE 8 A monopolist distributor's optimal profit $\pi_{0.5}^{m*}(p^*, \ell^*)$ versus $\rho_D(\ell^*)$ for $\beta = 1.0, 1.5, 2.0, 2.5$, and 3.0 and $F = 2.0$.

where $\beta \equiv \psi/\lambda$. In summary, (4) and (18)–(19) characterize the customers' optimal choices q_s , and q_f^B, q_f^S for a given distributor choice (p, ℓ) , and (35)–(41) characterize a monopolist distributor's optimal choices (p^*, ℓ^*) .

Figure 7 shows equilibrium choices (p^*, ℓ^*) of a monopolist distributor for $F = 2.0$ and different values of β . For a given β (that is, the ratio of ψ and λ), the distributor chooses a higher per unit price p^* for lower levels of investment ℓ^* in fraud monitoring. This conclusion also holds for other values of F . The optimal profit $\pi_{0.5}^{m*}(p^*, \ell^*)$ (net A) of the distributor also increases with $\rho_D(\ell^*)$, as shown in Figure 8. That is, for a given β (ratio of ψ and λ), the distributor's profit increases with level of investment ℓ^* in fraud monitoring. Figures 7 and 8 also show that for fixed F and a chosen $\rho_D(\ell^*)$, the monopolist distributor prefers higher p^* as β increases (that is, ψ increases or λ decreases). These observations are consistent with the price-setting behavior of a profit-maximizing monopolist.

This article develops a game-theoretic framework that explicitly models the adversarial nature of the electricity theft problem.

It is instructive to study how the distributor's collection efficiency η [see (23)] varies with different equilibrium choices (p^*, ℓ^*) . Figure 9 (respectively, Figure 10) shows $\eta(p^*, \ell^*)$ versus $\rho_F(\ell^*)$ for different values of β when the fixed fine $F = 0.5$ (respectively, $F = 2.0$) and the fraction of type-f customers is $\lambda = 0.5$. Interestingly, the collection efficiency η versus $\rho_F(\ell^*)$ shows different behavior for different fine levels. In particu-

lar, for lower values of F , the collection efficiency η increases with ℓ^* . In contrast, for higher values of F , η decreases with ℓ^* . This behavior can be partly explained from the fact that the ratio q_f^*/q_g^* (where $q_g^* = (p^*)^{-2}$) evolves differently with respect to $\rho_F(\ell^*)$ for low and higher values of F . In addition, for fixed F and chosen $\rho_D(\ell^*)$, the efficiency η decreases as β (or equivalently, ψ) increases.

Regulation of Electricity Distribution Utilities

Electricity is delivered to the customers by distribution utilities (or distributors), which are firms operating as regulated monopolists. Each distributor is an exclusive franchise. It is subjected to tariff and performance regulations by the public utility commission (or regulator). The principles for tariff regulation are broadly similar across different utilities [23]. Ideally, the regulator would like to achieve operational efficiency to ensure reliable delivery at the lowest cost, dynamic efficiency to meet future demand, and consumption efficiency to ensure the lowest prices subject to cost recovery of maintenance and provision of short-term cost and long-term investment by the distributor [24], [25].

The design of regulatory requirements would be an easy task if the regulator were perfectly informed about the distributor's costs and the customer demand [26]. In reality, the distributor has an informational advantage over the regulator about both aggregate customer demand and its own operational costs; see Figure S2. In such cases, the regulatory design can become extremely subtle and fragile to changes in the regulator's assumptions about the distributor's efficiency and costs. There exists a well-developed body of work dedicated to designing optimal regulatory policies for a monopolistic distributor who has privileged information about its own technological capabilities and customers' demand and when the regulator has well-defined intertemporal commitment powers [27]. Here the regulator is not subject to a time-inconsistent optimal policy. However, such a normative analysis assumes that the imperfectly informed regulator perfectly *knows the structure* of the regulated environment and has a formal model of information asymmetry between the regulator and the distributor.

In practice, the precise nature of information asymmetry and the full set constraints that affect the regulator and the distributor are not known a priori. Hence, "well-designed" regulatory policy must be robust, that is, it must perform "reasonably well" under broad conditions, although such a policy may be suboptimal in each particular case [28]. Two main regulatory regimes have been adapted for distribution utilities: i) rate of return (dominant regime in the United States) and ii) price cap (dominant regime in European Union and developing countries).

FIGURE S2 The players in regulated electricity distribution. A central issue in the regulation of distribution utilities is the presence of asymmetric information between the three entities affected the electricity distribution system: the regulator, distributors, and end customers.

Below each regime is outlined briefly. This article considers an average revenue constraint imposed by the regulator, which is an example of price-cap regulation.

RATE OF RETURN VERSUS PRICE CAP REGULATION

Under *rate-of-return regulation*, the distribution utility is given a prespecified a rate of return, and the tariff structures for the electricity are adjusted as the distributor's cost changes to ensure that the distributor will be able to earn the authorized rate of return. Here, the regulator bears the onus of setting the prices and ensures that the rate of return does not deviate significantly from the target rate. Since the prices are directly linked to the distributor's costs, the distributor lacks incentives to engage in cost-reducing activities. A classical example is the Averch-Johnson effect, which demonstrates that under the rate-of-return regulation, the distributor deviates from cost minimization. However, since the distributor faces a limited risk of expropriation of sunk investments by the regulator, upgrades to the distribution network can

Perfect Competition

The optimal choices (p^*, ℓ^*) for the case when distributor is a monopolist are different than the optimal choices in case of perfect competition. Now consider a benchmark case of perfect competition, where the average per customer profit is zero, that is,

$$\pi_\lambda^c = A + (p - c)q_g(p) + \lambda(-pq_f^S + F\rho_D(\ell, q_f^S)) - \psi\ell = 0, \quad (42)$$

where the superscript c on π_λ emphasizes that the distributor's operating environment is that of perfect competition. Let $(p^\dagger, \ell^\dagger)$ denote the distributor's choice of per unit price of electricity and investment level in fraud monitoring under

perfect competition, respectively. The following set of conditions lead to zero profit

$$(p^\dagger - c)q_g(p^\dagger) = 0 \Rightarrow p^\dagger = c, \text{ using (4)} \quad (43)$$

$$A + F\rho_D(\ell^\dagger, q_f^S(p^\dagger, \ell^\dagger)) = \frac{\psi\ell^\dagger}{\lambda} + p^\dagger q_f^S(p^\dagger, \ell^\dagger). \quad (44)$$

Note that these conditions are not the only ones that ensure zero distributor profit. With (16), and assuming $\rho_F(\ell) = \ell/1 + \ell$, (44) can be rewritten as

$$p^\dagger q_f^S + F\left(\frac{1}{1 + \ell^\dagger}\right)^{\frac{1}{1 - (p^\dagger)^2 q_f^S}} = A + F - \frac{\psi\ell^\dagger}{\lambda}. \quad (45)$$

be sustained in this form of regulation. The investment incentives of the regulated distributor are especially important since the infrastructure upgrades (such as capacity expansion) and modernization (such as AMI installations) require substantial costs.

Under *price cap regulation*, the tariffs rate of the distributor to customers could increase, on average, at a specified rate during a prespecified time. The specified rate is typically linked to the inflation rate and may fail to reflect the distributor's short-term costs and/or profit. Typically, under a price cap regulatory regime, only average prices are controlled by the regulator, and the utility is given the flexibility to control the pattern of relative prices subject to predefined constraints. Since the tariff rates are specified for relatively long periods of time, the distributor has incentives to minimize its operating costs, and thus to operate efficiently.

Notice that a price cap does not directly provide incentives for long-term investments, such as distribution network upgrades and reduction of nontechnical losses. Similarly, a price cap does not incentivize the distributor to choose optimal allocation of service quality. To remedy this, additional requirements on service quality are frequently imposed. Still, the price cap regulation may fail to incentivize the distributor to invest in monitoring and enforcement efforts to reduce unbilled electricity (such as customer theft) at socially optimal levels.

When the pricing flexibility of price cap regulation is combined with the rewards (respectively, punishments) for performance improvement (respectively, deterioration) relative to the regulator's benchmark, the resulting regime is termed *performance-based* (or incentive) regulation. Indeed, in the face of a rapidly changing technological environment and evolving customer preferences, the regulated electricity distribution industry is moving toward incentive regulation. The goal of incentive regulation is to improve distributors' incentives by decoupling the regulated price structure from the need to know the exact operating and maintenance costs.

REGULATED DISTRIBUTOR

This article presents distributors' optimal choices (p, ℓ) for the case of an unregulated monopolist and the case of perfect competition. This analysis can be extended to a regulated distributor who is

subject to price cap or rate-of-return regulation. For example, the distributor could face an average revenue constraint imposed by a regulator, that is, with the tariff schedule $T(\cdot)$, fine schedule $F(\cdot)$, and the investment level ℓ in fraud monitoring. Then, the average revenue (per unit quantity) collected should be no more than a regulator-specified price cap $\bar{p}$. The computation of average revenue depends on audits and regulatory processes followed by the regulator.

The average revenue can be computed based on the total quantity Q_T provided by the distributor or the quantity $(Q_T - Q_f^S)$, which excludes the stolen quantity of electricity; see (22)–(23). In the later case, the regulator only accounts for the billed and recovered (via fines) quantity in setting the price cap for the distributor's average revenue. Thus, two possible designs of an average revenue constraint are

$$\frac{R_\lambda(p, \ell)}{Q_T(p)} \leq \bar{p}, \quad (S1)$$

$$\frac{R_\lambda(p, \ell)}{Q_T(p) - Q_f^S(p, \ell)} \leq \bar{p}, \quad (S2)$$

where the total revenue $R_\lambda(p, \ell)$ is given by (30). Clearly, (S2) imposes a stricter regulatory imposition on the distributor. In the case of (S2), the regulator does not account for the fraudulent customers' surplus resulting from the successfully stolen electricity Q_f^S . From the viewpoint of the distributor (respectively, regulator), the constraint (S1) [respectively, (S2)] is more desirable because it eases (respectively, tightens) the regulatory constraint. Using (22) and (30), and for special case $F'(\cdot) \approx F$, constraints (S1) and (S2) can be rewritten as

$$-p^2 q_f^S + F\rho_D(\ell, q_f^S) \leq \frac{1}{\lambda} \left(\frac{\bar{p}}{p} - Ap - 1 \right),$$

$$pq_f^S(-p + \bar{p}\rho_D(\ell, q_f^S)) + F\rho_D(\ell, q_f^S) \leq \frac{1}{\lambda} \left(\frac{\bar{p}}{p} - Ap - 1 \right).$$

When an average revenue constraint is imposed on the distributor, the regulated distributor's optimal price p^r and level of investment in fraud monitoring ℓ^r can be obtained by solving the constrained optimization problem (32), subject to ROC curve (16), customer responses (4) and (17)–(19), and the average revenue constraint (S1) or (S2).

FIGURE 9 The distributor's collection efficiency $\eta(p^*, \ell^*)$ versus $\rho_D(\ell^*)$ for $\beta = 1.0, 1.5, 2.0, 2.5$, and 3.0 when $F = 0.5$ (low fine) and $\lambda = 0.5$.

FIGURE 10 The distributor's collection efficiency $\eta(p^*, \ell^*)$ versus $\rho_D(\ell^*)$, for $\beta = 1.0, 1.5, 2.0, 2.5$, and 3.0 when $F = 2.0$ (high fine) and $\lambda = 0.5$.

Next, substituting $\rho_F(\ell) = \ell/1 + \ell$, into the fraudulent customer's FOC (17) gives

$$\frac{1}{(1 - (p^\dagger)^2 q_f^S)^2} \left(\frac{1}{1 + \ell^\dagger} \right)^{\frac{1}{1 - (p^\dagger)^2 q_f^S}} = \frac{1}{F p^\dagger \ln(1 + \ell^\dagger)}. \quad (46)$$

Optimal $(p^\dagger, \ell^\dagger)$ can be determined by feasible solutions of equations (45)–(46) for $p^\dagger = c$, and given parameters A, F, ψ , and λ . In particular, optimal $q_f^{S\dagger}$ (and hence, $q_f^{B\dagger}$) can be obtained as

$$q_f^{S\dagger} = \frac{1}{(p^\dagger)^2} (1 - y(p^\dagger, \ell^\dagger)),$$

where

$$y(p^\dagger, \ell^\dagger) = 0.5 \sqrt{\ln(1 + \ell^\dagger)} (\sqrt{\ln(1 + \ell^\dagger)} + \sqrt{\ln(1 + \ell) + 4[(A + F)p^\dagger - \beta \ell^\dagger p^\dagger - 1]}).$$

Substituting $q_f^{S\dagger}$ into (46) gives $\ell^\dagger$.

This completes the analysis of distributor choices when he is a monopolist or faces perfect competition. See “Regulation of Electricity Distribution Utilities” for a discussion of environments when the distributor is a regulated monopolist.

AN ALTERNATIVE FORMULATION

A second game-theoretic model, which is not fully covered here, assumes that all customers have the same initial preferences (utility function) and that they make a decision to become fraudulent, or stay genuine, depending on the probability of detection and the fine they would face if caught.

Once the customers make their decisions about which type they will be (genuine or fraudulent), they could be viewed as if they are playing the game described in this article. Therefore, this second model could also be viewed as a leader-follower game, where relative to the first model, the customers have to make an additional decision, that is, to choose whether they will be honest or fraudulent.

For any fixed fine and detection probability, it is possible to determine what fraction of customers will be fraudulent in equilibrium. Thus, it is possible to jointly solve the problem of the distributor's choice of security investment and find the corresponding fraction of customers that would choose to be fraudulent with a given security investment. Then, the problem becomes identical to the original formulation. This allows the distributor to compute expected profit as a function of security investment. Next, if the distributor is a monopolist, it maximizes its profit and chooses the equilibrium level of investment in monitoring fraud that achieves the highest profit.

ACKNOWLEDGMENT

The work was supported in part by CPS: Frontiers: Collaborative Research: Foundations of Resilient CybEr-Physical Systems, which receives support from the National Science Foundation under the Award CNS-1239054. The authors are grateful for useful advice and suggestions from four anonymous reviewers.

AUTHOR INFORMATION

Saurabh Amin (amins@mit.edu) is an assistant professor in the Department of Civil and Environmental Engineering, Massachusetts Institute of Technology (MIT). His research focuses on the design and implementation of high-confidence network control algorithms for infrastructure systems. He received the Ph.D. in systems engineering from the University of California, Berkeley, the M.S. from the University of Texas at Austin, and the B.Tech. from IIT Roorkee. He can be contacted at Massachusetts Institute of Technology, 77 Massachusetts Avenue 1-241, Cambridge, MA 02139 USA.

Galina A. Schwartz is a researcher in the Department of Electrical Engineering and Computer Sciences at the

Historically, widespread energy theft is characteristic for developing countries, with theft of electricity reaching up to 50% in some jurisdictions.

University of California, Berkeley. Her primary expertise is game theory and microeconomics. She studies the resilience of large-scale networked systems and their interfaces with humans using game-theoretic tools. She received the M.S. in mathematical physics from Moscow Institute of Engineering Physics and the Ph.D. in economics from Princeton University.

Alvaro A. Cárdenas is an assistant professor at the University of Texas (UT), Dallas. He holds M.S. and Ph.D. degrees from the University of Maryland, College Park, and a B.S. from the Universidad de los Andes. Prior to joining UT, Dallas, he was a researcher with Fujitsu Laboratories of America and a postdoctoral scholar at the University of California, Berkeley. His research interests include information security, cyberphysical systems, the smart grid, and intrusion detection.

S. Shankar Sastry is dean of engineering at the University of California (UC), Berkeley, and director of the Blum Center for Developing Economies. He received the Ph.D. in 1981 from UC Berkeley. His areas of research are embedded and autonomous software for unmanned systems, computer vision, nonlinear and adaptive control, robotic tele-surgery, control of hybrid and embedded systems, cyber-physical security, and critical infrastructure protection. He is a member of the National Academy of Engineering and the American Academy of Arts and Sciences.

REFERENCES

- [1] S. McLaughlin, D. Podkuiko, and P. McDaniel, "Energy theft in the advanced metering infrastructure," in *Proc. 4th Int. Conf. Critical Information Infrastructures Security*, 2009, pp. 176–187.
- [2] Guidelines for smart grid cyber security: Privacy and the smart grid, U. S. Dept. Commerce, Nat. Inst. Standards Technol., Interagency Rep. 7628, Aug. 2010, vol. 2.
- [3] M. LeMay and C. A. Gunter, "Cumulative attestation kernels for embedded systems," *IEEE Trans. Smart Grid*, vol. 3, no. 2, pp. 744–760, 2012.
- [4] E. de Buda, "System for accurately detecting electricity theft," U.S. Patent Application 12/351 978, Jan. 2010.
- [5] A. W. Appel, "Security seals on voting machines: A case study," *ACM Trans. Inform. Syst. Secur.*, vol. 14, no. 2, pp. 1–29, 2011.
- [6] P. Antmann, "Reducing technical and non-technical losses in the power sector," World Bank, Washington, D.C., Tech. Rep., July 2009.
- [7] S. McLaughlin, D. Podkuiko, S. Miadzvezhanka, A. Delozier, and P. McDaniel, "Multi-vendor penetration testing in the advanced metering infrastructure," in *Proc. Annu. Computer Security Applications Conf.*, Dec. 2010, pp. 107–116.
- [8] M. Davis. (2009, July). Smartgrid device security. Adventures in a new medium. [Online]. Available: <http://www.blackhat.com/presentations/bh-usa-09/MDAVIS/BHUSA09-Davis-AMI-SLIDES.pdf>
- [9] D. Peterson. (2012, Apr.). AppSecDC in review: Real-world backdoors on industrial devices. [Online]. Available: <http://www.digitalbond.com/2012/04/11/appsecdc-in-review/>
- [10] B. Krebs. (2012, Apr.). FBI: Smart meter hacks likely to spread. [Online]. Available: <http://krebsonsecurity.com/2012/04/fbi-smart-meter-hacks-likely-to-spread/>
- [11] A. Lesser. (2012, Mar.). When big IT goes after big data on the smart grid. [Online]. Available: <http://gigaom.com/cleantech/when-big-it-goes-after-big-data-on-the-smart-grid-2/>
- [12] C. Geschickter, "The emergence of meter data management: A smart grid information strategy report," GTM Res., Tech. Rep. 214831, 2010.
- [13] C. J. Bandim, J. E. R. Alves Jr., A. V. Pinto Jr., F. C. Souza, M. R. B. Loureiro, C. A. Magalhaes, and F. Galvez-Durand, "Identification of energy theft and tampered meters using a central observer meter: A mathematical approach," in *Proc. IEEE PES Transmission Distribution Conf. Expo.*, 2003, vol. 1, pp. 163–168.
- [14] A. Nizar and Z. Dong, "Identification and detection of electricity customer behaviour irregularities," in *Proc. Power Systems Conf. Expo.*, Mar. 2009, pp. 1–10.
- [15] J. Nagi, K. S. Yap, S. K. Tiong, S. K. Ahmed, and M. Mohamad, "Non-technical loss detection for metered customers in power utility using support vector machines," *IEEE Trans. Power Delivery Syst.*, vol. 25, no. 2, pp. 1162–1171, Apr. 2010.
- [16] S. S. S. R. Depuru, L. Wang, and V. Devabhaktuni, "Support vector machine based data classification for detection of electricity theft," in *Proc. IEEE Power Systems Conf. Expo.*, Mar. 2011, pp. 1–8.
- [17] S. McLaughlin, B. Holbert, A. Fawaz, R. Berthier, and S. Zonouz, "A multi-sensor energy theft detection framework for advanced metering infrastructures," *IEEE J. Sel. Areas Commun.*, vol. 31, no. 7, pp. 1319–1330, 2013.
- [18] D. Mashima and A. A. Cárdenas, "Evaluating electricity theft detectors in smart grid networks," in *Research in Attacks, Intrusions, and Defenses*. Berlin Heidelberg, Germany: Springer-Verlag, 2012, pp. 210–229.
- [19] R. Jiang, R. Lu, Y. Wang, J. Luo, C. Shen, and X. S. Shen, "Energy-theft detection issues for advanced metering infrastructure in smart grid," *Tsinghua Sci. Technol.*, vol. 19, no. 2, pp. 105–120, 2014.
- [20] S. E. McLaughlin, P. McDaniel, and W. Aiello, "Protecting consumer privacy from electric load monitoring," in *Proc. ACM Conf. Computer Communications Security*, 2011, pp. 87–98.
- [21] G. Kalogridis, C. Efthymiou, S. Z. Denic, T. A. Lewis, and R. Cepeda, "Privacy for smart meters: Towards undetectable appliance load signatures," in *Proc. 1st IEEE Int. Conf. Smart Grid Communications*, Oct. 2010, pp. 232–237.
- [22] D. C. Bergman, D. Jin, J. P. Juen, N. Tanaka, C. A. Gunter, and A. K. Wright, "Distributed non-intrusive load monitoring," in *Proc. IEEE PES, Innovative Smart Grid Technologies*, Jan. 2011, pp. 1–8.
- [23] J. G. Kassakian and R. Schmalensee. (2011). The future of electric grid: An interdisciplinary MIT study. Massachusetts Inst. Technol., Cambridge, MA, Tech. Rep., 2011. [Online]. Available: <http://mitei.mit.edu/publications/reports-studies/future-electric-grid>
- [24] P. L. Joskow, "Incentive regulation in theory and practice: Electricity distribution and transmission networks," in *Economic Regulation and its Reform: What Have We Learned?*, Nat. Bureau Econ. Res., Inc., Tech. Rep. 0514, 2011.
- [25] I. Vogelsang, "Electricity transmission pricing and performance-based regulation," CESifo Working Paper Series 1474, CESifo Group Munich, 2005.
- [26] J. J. Laffont and D. Martimort, *The Theory of Incentives: The Principal-Agent Model*. Princeton, NJ: Princeton Univ. Press, 2002.
- [27] P. Bolton and M. Dewatripont. *Contract Theory*, vol. 1. Cambridge, MA: MIT Press, 2005.
- [28] M. Armstrong and D. E. M. Sappington, "Recent developments in the theory of regulation," in *Handbook of Industrial Organization*, vol. 3. Amsterdam, The Netherlands: Elsevier, 2007, ch. 27, pp. 1557–1700.