Game Theoretic Models of Electricity Theft Detection in Smart Utility Networks

Citation: Amin, Saurabh, Galina A. Schwartz, Alvaro A. Cardenas, and S. Shankar Sastry. "Game-Theoretic Models of Electricity Theft Detection in Smart Utility Networks: Providing New Capabilities with Advanced Metering Infrastructure." IEEE Control Syst. 35, no. 1 (February 2015): 66–81.

As Published: http://dx.doi.org/10.1109/MCS.2014.2364711

Publisher: Institute of Electrical and Electronics Engineers (IEEE)

Persistent URL: http://hdl.handle.net/1721.1/101377

Version: Original manuscript: author's manuscript prior to formal peer review

Game Theoretic Models of Electricity Theft Detection in Smart Utility Networks

Saurabh Amin, Galina A. Schwartz, Alvaro A. Cárdenas, S. Shankar Sastry

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 consumers. One of the key technologies being widely deployed on the consumption side of the grid is the Advanced Metering Infrastructure (AMI).

AMI refers to the modernization of the electricity metering system by replacing old mechanical meters by *smart meters*. Smart meters are new embedded devices that provide two-way communications between the utility and the consumer. These devices have advanced communication and computational capabilities, with a potential to enable new functionalities such as improved service choices, transparencies, etc. Distribution utilities (or distributors) using AMIs for monitoring and billing of electricity consumption can avoid sending their employees to read the meters on-site. Importantly, AMIs provide several new capabilities: monitoring of network-wide and individual electricity consumption, faster remote diagnosis of outages (with analog meters, utilities

learned of outages primarily by consumer call complaints), remote disconnect options, and automated power restoration. The AMIs also improve the consumers' access to their energy usage information (including the sources of electricity, renewables or otherwise) and promote the implementation of demand response schemes.

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 actual range of cost of devices widely varies: $\$100-\$400$ per device (excluding installation and maintenance costs). Hardening these devices by adding hardware co-processors 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 costs of millions of devices. Thus, creating a business case for improving 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 of the customers are fraudulent or misbehaving, 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. Indeed, they have clear potential to improve metering, billing and collection processes, and detection of fraud and unmetered connections.

Electricity theft in distribution networks

Historically, widespread energy theft is characteristic for developing countries. Indeed, according to a World Bank report [6], the theft of electricity reaches up to 50% in some jurisdictions of developing countries. Traditional ways range from comprise of physical security of meters to directly connecting loads to the electricity distribution lines. Default of payments was a major problem due to suboptimal levels of monitoring and enforcement. Lack of technology and insufficient distributor incentives were the major causes of this problem.

Non-technical losses and electricity theft

In general, distribution utilities can incur non-technical losses due to actions of (i) a utility personnel/operator (administrative losses due to errors in accounting and record keeping), (ii) customer theft (fraud or willful pilferage by bona-fide customers), (iii) customer non-payment (default), and lastly (iv) the theft by the outsiders (non-customers). The administrative errors can be strategic (i.e., intentional) when made with a purpose of assisting customer theft.

For a distribution utility, the non-technical losses (e.g., electricity theft, fraud, or uncollected/defaulted bills) contribute to costs. The consumers who acquire electricity via stealing fail to pay for electricity or defaulting on their bills, obtain the electricity at near zero prices. Effectively, the electricity consumption of these non-paying entities is subsidized, because their consumption is paid by the distribution utility and/or other consumers, or in some cases, by subsidies from local taxes. Overall, the consumption of these non-paying entities is paid by the society at large.

The non-technical losses can be recovered by (i) imposing higher electricity tariffs on other (paying) consumers, (ii) decreasing profit margins of the distributor; (iii) distributing the burden on the entire society, for example, by increasing taxes. The actual means depend on the security and recovery technologies that are available to the distributor, his choices to invest in them, and the regulatory environment. But when the distributor ends up being the net bearer of losses for a prolonged period of time and no regulatory resources exist to recover these losses, his incentives and capabilities to invest into the network and its maintenance are jeopardized. Such underinvestment negatively affects the long-run efficiency of distribution system. Thus, to improve efficiency of distribution systems, both technological and regulatory means to limit non-technical 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 cyber vulnerabilities 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, with resulting losses of millions for dollars for a single US utility, as reported by a cyber-intelligence bulletin issued by the FBI [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 FBI report also predicts and conjectures with medium confidence that as smart grid deployments continues, the cyber-means of electricity theft 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 consumers.

Still, in regulated environments, new investments required for 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 consumers, publicizing violations for sharper public scrutiny, increasing consumers awareness that electricity theft is a cognizable offense, and disconnecting customers for fraud/debts and reconnecting their service only after the blue remittance of the required payments. Examples of regulatory measures include fixing the 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 consumers, and are developing analytics capabilities for improved situational awareness [11]. Meter Data Management (MDM) vendors are providing analytics services to the utilities to turn their data into actionable information; see Fig. 5. An important MDM service is called revenue assurance. It provides data-analytics software to identify suspected electricity theft through 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 two reasons: First, it can be easily retrofitted with an existing distribution infrastructure. Second, unlike other security technologies it does not require major capital investments needed by other security technologies such as balance meters. Third, the extra security is a by product of the main reason of the MDM's popularity because it has its own value added due to its data storage and processing capabilities.

Related work

Early research on detection of electricity theft focused on the role of a set of trusted balance meters and looked at electricity consumption traces to check meters'

accuracy [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], [15], [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 a 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].

Focus of the article

This article presents a game-theoretic framework for modeling the adversarial nature of the electricity theft problem. The model considers both pricing and investment decisions by the distribution utility (i.e., the distributor) who faces a population consisting of two consumer types: genuine and fraudulent. Both types of consumers derive identical utility from using electricity (preferences), but face different costs. The genuine consumers 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 consumers choose two amounts: first, the amount for which they will pay (as genuine ones do), and second, 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 received operating characteristic (ROC) curve (i.e., relationship between probability of detection and probability of false alarm). The probability of detection depends on two factors. First, it depends on the stolen amount (the probability increases with the stolen amount), and second, on the level of investment made by the distributor monitoring fraud. Higher level of investment by the distributor improves the probability of detection. The distributor chooses how much to invest in level of fraud monitoring and the price per unit quantity of billed electricity.

The article considers the two environments: 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 consumers. The article computes equilibria of both games, and compares the level of effort for unregulated monopoly and perfect competition. In both games, consumers 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 provisioning 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 consumers (when detected). The distributor's profit, i.e., revenue net costs, depends on both the level of investment and the per unit price he offers to consumers. The chosen level of investment and the consumers' 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 given distributor choice of price and level of investment, the consumers' response functions are derived. Finally, the optimal choices for the case when 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.

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 (re-programming) or cyber physical means (such as installing a rigged smart meter). Clearly, in response to such threats, the distributor can employ diagnostic schemes to find the fraudulent consumers. The game-theoretic framework proposed in this article can help analyze equilibrium consumer and distributor choices in scenarios where the assumptions on consumer utilities and distributor's profit function are applicable.

Modeling Consumer Preferences

Let $\mathcal{N} = \{1, \ldots, n\}$ denote the population of consumers that are served by the distributor. The security level of individual meters may vary across the population. For simplicity, assume that each consumer is either of type-f or of type-g. The AMIs of type-f consumers posses certain security vulnerabilities and/or installation defects that can be exploited for economic gains. When these consumers are successful in stealing electricity, the distributor does not fully recover electricity bills and incurs more non-technical losses. The type-g consumers do not have the technological means to exploit AMI security vulnerabilities or their AMIs are just harder to exploit because of the due care taken during the installation process. Each type-g consumer fully pays for the electricity he consumes. Thus, type-f and type-g consumers can be called "fraudulent" and "genuine", respectively. Let $\mathcal{N}_f \subset \mathcal{N}$ and $\mathcal{N}_g = \mathcal{N} \backslash \mathcal{N}_f$ denote the sets of these consumer types, and let λ be the fraction of type-f consumers, that is, $\lambda = \frac{\mathcal{N}_f}{\mathcal{N}}$. The distributor (a monopolist) knows the fraction λ, but cannot distinguish between type-f and type-g consumers without investing in monitoring and enforcement efforts.

Genuine consumers

Suppose that each type-g consumer has the following utility function:

$$U_g = u(q_g) - T(q_g) \quad \text{[Secure AMIs]}, \tag{1}$$

where the function $u(\cdot)$ (assumed to be same for all type-g consumers) satisfies $u(0) = 0$, $u'(q) > 0$, and $u''(q) < 0$, i.e., there is a decreasing marginal utility of electricity

consumption. If the distributor offers a tariff schedule $T(\cdot)$, a type-g consumer chooses *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 consumer surplus is given by:

$$v_g \equiv \max_{q_g \geq 0} \left[u(q_g) - T(q_g) \right], \tag{2}$$

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 which can be interpreted as a connection fee and p is constant per unit price (usage charge). For the purpose of analytical derivations, this article assumes that consumer preference is given by a quadratic function $u(q_g) = 2\sqrt{q_g}$. In this case, the chosen consumption and optimal surplus of a type-g consumer becomes:

$$q_g(p) = \frac{1}{p^2}, \quad v_g^*(p) = \left(\frac{1}{p} - A \right). \tag{3}$$

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

$$Q_g(p) = \frac{n(1-\lambda)}{p^2}. \tag{4}$$

Fraudulent consumers

Consider the following utility function for each type-f consumer:

$$U_f = u(q_f^B + q_f^S) - T(q_f^B) - \rho_D(\ell, q_f^S) F^r(q_f^S) \quad \text{[Insecure AMIs]}, \tag{5}$$

where $u(\cdot)$ and $T(\cdot)$ are same as in (1), q_f^B and q_f^S respectively denote the *expected* billed and stolen (or unpaid) quantities for a type-f consumer, $\rho_D(\ell, q_f^S)$ the probability that a fraudulent consumer is detected when distributor's level of investment in monitoring of fraud is $\ell \in \mathbb{R}_+$, and $F^r(\cdot)$ 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 consumers 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 consumer 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 consumer surplus is given by

$$v_f \equiv \max_{q_f^B \geq 0, q_f^S \geq 0} \left[u(q_f^B + q_f^S) - T(q_f^B) - \rho_D(\ell, q_f^S) F^r(q_f^S) \right], \tag{6}$$

and the first-order conditions (FOCs) are given by :

$$\partial_{q_f^B}[u(q_f^B + q_f^S)] = T'(q_f^B), \quad \partial_{q_f^S}[u(q_f^B + q_f^S)] = \partial_{q_f^S}[\rho_D(\ell, q_f^S) F^r(q_f^S)]$$

That is, a small increase in total quantity ($q_f = q_f^B + q_f^S$) consumed by a type-f consumer generates a marginal surplus $u'(q_f)$ equal to marginal payment $T'(q_f^B)$ (resp. *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 (resp. 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 quadratic consumer preferences $u(q_f) = 2\sqrt{q_f}$, the FOCs imply that quantities q_f^B and q_f^S satisfy:

$$\rho_D(\ell, q_f^S)p_f + \partial_{q_f^S}[\rho_D(\ell, q_f^S)] \left[F + p_f q_f^S \right] = p, \quad q_f^B = \frac{1}{p^2} - q_f^S. \tag{7}$$

Note that $q_f = q_f^B + q_f^S = q_g = \frac{1}{p^2}$. This results from the assumption that each consumer's valuation of the total quantity of electricity does not depend on consumer type, i.e., $u(\cdot)$ is same for both type-g and type-f consumers. For the case when, upon detection, the fraudulent consumer pays a fixed fine F that is much larger than $p_f q_f^S$, i.e., $F^r(\cdot) \approx F$, the FOCs (7) simplify to:

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

The probability of detection $\rho_D(\ell, q_f^S)$ is a property of the diagnostic scheme employed by the distributor, and because of the variability of meter measurements received from genuine and fraudulent consumers, a high value of $\rho_D(\ell, q_f^S)$ also entails a high value of the probability ρ_F of false positive (or false alarm). The statistical decision theory models this trade-off between ρ_D and ρ_F values of a diagnostic scheme as a received operating characteristics (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 probability of false alarm ρ_F increases as distributor's level of effort ℓ in monitoring of fraud increases, i.e., $\rho_F(\ell) \in (0, 1)$ is increasing in $\ell \in \mathbb{R}_+$. Furthermore, let $\rho_F(\cdot)$ be a continuously differentiable and invertible function.

For the purpose of analytical tractability, consider the following ROC curve:

$$\rho_D(\ell, q_f^S) = 1 - [1 - \rho_F(\ell)]^{\left(\frac{q_g}{q_f^B}\right)} = 1 - [1 - \rho_F(\ell)]^{\left(\frac{1}{1 - p^2 q_f^S}\right)}, \tag{9}$$

where the second equality follows from $q_g = (q_f^B + q_f^S) = p^{-2}$. As the stolen quantity $q_f^S \to 0$, $\rho_D \to \rho_F$, i.e., the diagnostic scheme uses random guessing, and as $q_f^S \to q_g$, $\rho_D \to 1$, i.e., the diagnostic scheme detects fraud with high probability. In fact, the ROC curve (9)

corresponds to the case when the meter measurements received by the distributor for type-g and type-f consumers follow exponential distributions with parameters $1/q_g$ and $1/q_f^B$, $q_g \geq q_f^B$, respectively. For this assumption, the probability density functions of meter measurements, q, collected from type-g and type-f can be written as:

$$f_g(q) = \frac{1}{q_g} \exp\left(-\frac{q}{q_g}\right), \text{ and } f_f(q) = \frac{1}{q_f^B} \exp\left(-\frac{q}{q_f^B}\right). \tag{10}$$

Consider that the diagnostic scheme employed by the distributor uses meter measurements and a threshold value τ to detect fraudulent consumers. A consumer is classified as fraudulent if $q < \tau$ for that consumer, and genuine otherwise. It follows that

$$\rho_D = \int_0^\tau \frac{1}{q_f^B} \exp\left(-\frac{q}{q_f^B}\right) dq = 1 - \exp\left(-\frac{\tau}{q_f^B}\right),$$
$$\rho_F = \int_0^\tau \frac{1}{q_g} \exp\left(-\frac{q}{q_g}\right) dq = 1 - \exp\left(-\frac{\tau}{q_g}\right).$$

ρ_D can be expressed as a function of ρ_F by eliminating τ; also see Fig. 1. Thus, (7) represents the ROC curve of distributor's diagnostic scheme. By the Neyman-Pearson lemma, for a given distributor level of investment ℓ in fraud monitoring, the threshold value τ can be determined as follows:

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

It is important to note that, under the stated assumptions, the type g (resp. type-f) consumers influence the distributions of their meter readings only by choosing the mean parameter $1/q_g$ (resp. $1/q_f^S$) of the exponential distribution which characterizes their consumption patterns. In other words, consumers do not alter the probabilistic form of their distribution, but only the mean parameter.

The following definitions are introduced for notational convenience:

$$\alpha \equiv \frac{q_g}{q_f^B} = \frac{1}{1 - p^2 q_f^S}, \quad \bar{\rho}_F(\ell) \equiv (1 - \rho_F(\ell)).$$

Using the ROC curve (9) in FOCs (8) provides that α satisfies:

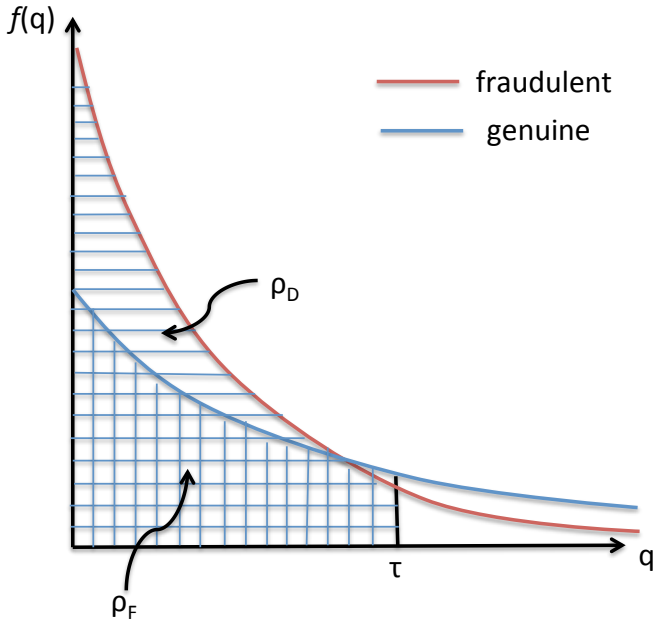

Figure 1. Detection probability ρ_D for given level of false alarm probability ρ_F.

$$\alpha^2 (1 - \rho_F(\ell))^\alpha \ln(1 - \rho_F(\ell)) = -\frac{1}{Fp}. \tag{11}$$

The solution for $\alpha(p, \ell)$ is given by

$$\alpha(p, \ell) = \frac{2W\left(\frac{1}{2}\sqrt{-\frac{\ln(1-\rho_F(\ell))}{Fp}}\right)}{\ln(1 - \rho_F(\ell))}, \tag{12}$$

where W is the *product logarithm* function defined as inverse function of $f(W) = We^W$.

Thus, for a choice $\rho_F(\ell)$ (or, equivalently, ℓ) and per unit price p of the distributor, and for quadratic consumer valuation $u(q_f) = 2\sqrt{q_f}$ and fixed fine schedule $F^r(\cdot) \approx F$,

the type-f consumer's chosen (optimal) consumptions q_f^S and q_f^B are:

$$q_f^B(p,\ell) = \frac{1}{p^2\alpha(p,\ell)}, \qquad q_f^S(p,\ell) = \frac{1}{p^2}\left(1 - \frac{1}{\alpha(p,\ell)}\right), \tag{13}$$

where $\alpha(p,\ell)$ is given by (12).

TABLE I

$\alpha(p,\ell)$ FOR DIFFERENT $\rho_F(\ell)$

$\rho_F(\ell)$	Equation to solve for $x(\ell)$
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	$Fp\alpha^2 \exp(-2.302\alpha) = 0.434$

The optimal surplus of a type-f consumer becomes:

$$v_f^*(p,\ell) = \frac{1}{p}\left(2 - \frac{1}{\alpha(p,\ell)}\right) - (A + F) + F(1 - \rho_F(\ell))^{\alpha(p,\ell)}. \tag{14}$$

A necessary condition for any type-f consumer to remain fraudulent, his optimal surplus v_f^* should be at least v_g^* (the type-g consumer's optimal surplus), that is:

$$v_f^* \geq v_g^*. \tag{15}$$

Consumers are indifferent between types when $v_f^* = v_g^*$. Equivalently, from (3) and (14), the necessary condition (15) becomes:

$$(1 - Fp)\,\alpha(p,\ell) + Fp\alpha(p,\ell)\,(1 - \rho_F(\ell))^{\alpha(p,\ell)} \geq 1. \tag{16}$$

For given p and ℓ choices of the distributor, the FOC (11) and the necessary condition (16) determine $\alpha(p, \ell)$, and hence the optimal consumer response $q_f^S(p, \ell), q_f^B(p, \ell)$. Figure 2 indicates how the fraction $q_f^S(p, \ell)/q_g(p)$ varies with $z \equiv Fp$ and $y \equiv \rho_F(\ell)$. For a given ℓ (resp. p), type-f consumers steal less as p (resp. ℓ) increases. Figure 3

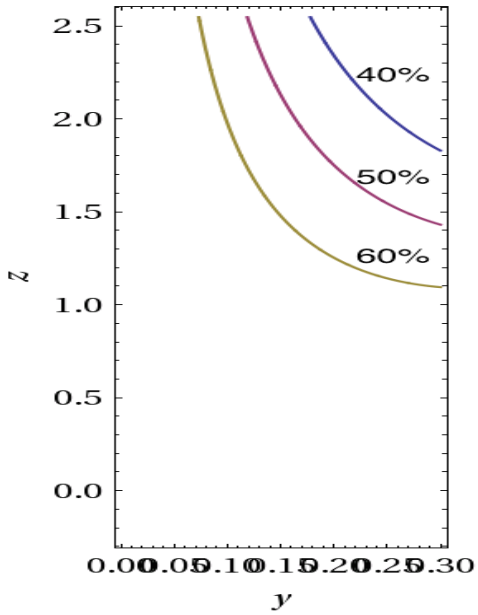

Figure 2. $z = Fp$ vs $y = \rho_F(\ell)$ for fractions $\frac{q_f^S}{q_g} = 0.4, 0.5, 0.6$.

plots $z \equiv Fp$ versus $y \equiv \rho_F(\ell)$ for $\alpha = 2.5$ (i.e., q_f^S is 60% of q_g) and $\alpha = 3.0$ (i.e., q_f^S is 66.6% of q_g). For $\alpha = 2.5$ (resp. $\alpha = 3.0$), the point ($z = 1.35$, $y = 0.16$) (resp. ($z = 1.15$, $y = 0.14$)) corresponds to maximum price p and minimum investment ℓ. From (9), the probability of detection of fraud $\rho_D = 0.38$ for $\alpha = 2.75$ and $\rho_D = 0.368$ for $\alpha = 0.30$. When $\alpha \in (2.275, 3.05)$, there exists at least one distributor choice of (p, ℓ) such that the FOC (11) and the necessary condition (16) are satisfied; this corresponds to q_f^S in the range 56% – 67% of q_g. Thus, under the stated assumptions, certain distributor

choices of price and investment levels can permit high levels of stolen electricity, for e.g., with $\rho_D = 0.38$ and $q_f^S/q_g = 0.6$, the expected unbilled quantity is $\approx 39\%$ of q_g.

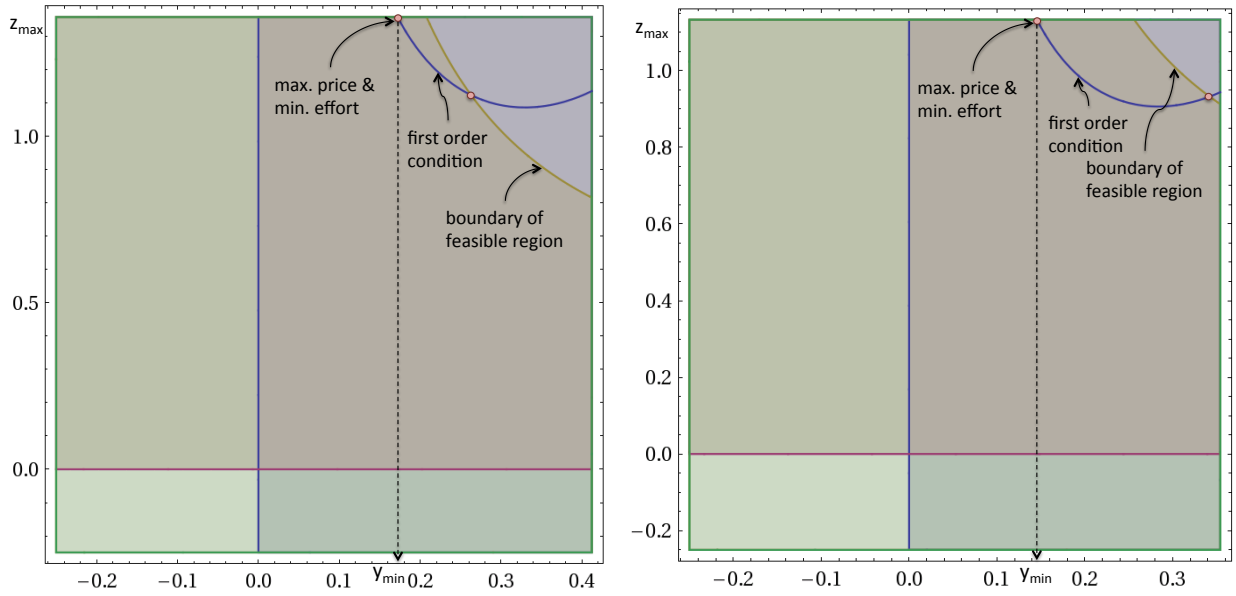

Figure 3. $z = Fp$ vs $y = \rho_F(\ell)$ for $\alpha = 2.5$ (left) and $\alpha = 3.0$ (right), where $\alpha = \frac{q_g}{q_f^B}$.

Finally, since $|\mathcal{N}_f| = n\lambda$, the total and unrecovered quantities of electricity consumed by fraudulent consumers are given by:

$$Q_f(p) = \frac{n\lambda}{p^2}, \quad Q_f^S(p,\ell) = n\lambda\left(1 - \rho_D(\ell, q_f^S(p,\ell))\right) q_f^S(p,\ell), \tag{17}$$

where $\left(1 - \rho_D(\ell, q_f^S(p,\ell))\right) q_f^S(p,\ell)$ is the unrecovered quantity from a type-f consumer.

Modeling Costs of Distribution Utility

Monopolist distributor

From (4) and (17), the total quantity provisioned 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 distributor's collection efficiency can be expressed as:

$$\eta(p, \ell) \equiv 1 - \frac{Q_f^S(p, \ell)}{Q_T(p)}, \tag{18}$$

where Q_f^S is given by (17).

For the quantity Q_T provisioned by the distributor, let revenue $R_\lambda(p, \ell)$ be his total revenue, when he offers a tariff schedule $T(q_f^B) = A + pq_f^B$ and implements a fine schedule $F^r(\cdot)$ to recover the quantity $n\lambda q_f^S \rho_D(\ell, q_f^S)$ from fraudulent consumers. Here the notation $R_\lambda(p, \ell)$ emphasizes the dependence of revenue on the distributor choice variables: the per unit price p and the level of fraud monitoring ℓ when the he faces λn fraudulent consumers. The following analysis considers $F^r(q_f^S) = F$ (i.e., $p_f \approx 0$). The total revenue is the sum of revenues generated from genuine and fraudulent consumers, i.e.,

$$\begin{aligned} R_\lambda(p, \ell) &= n(1-\lambda)\left[A + \frac{1}{p}\right] + n\lambda\left[A + p\left(\frac{1}{p^2} - q_f^S\right) + F\rho_D(\ell, q_f^S)\right] \\ &= n\left[A + pq_g(p) + \lambda\left(-pq_f^S + F\rho_D(\ell, q_f^S)\right)\right] \end{aligned} \tag{19}$$

There are two main operational costs to the distributor:

i) For provisioning the electricity to meeting the total demand Q_T in each billing period, the distributor faces the cost $C(Q_T)$, where C is 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 cost of deploying secure AMIs to ensure that a fraction $(1 - \lambda)$ of population is type-g consumers is a sunk cost, and is not considered here. For sake of simplicity, consider linear cost of provisioning $C(Q_T) = cQ_T$, $c > 0$, and a linear cost of monitoring fraud $\Psi(\ell) = n\psi\ell$, where $\psi > 0$. The average (per-consumer) profit for an unregulated monopolist is:

$$\begin{aligned}\pi_\lambda^m(\mathrm{p},\ell) \equiv \frac{\Pi_\lambda^m(\mathrm{p},\ell)}{n} &= \mathrm{R}_\lambda(\mathrm{p},\ell) - \mathrm{C}(\mathrm{Q_T}) - \Psi(\ell) \\ &= \mathrm{A} + (\mathrm{p}-c)\mathrm{q_g}(\mathrm{p}) + \lambda\left(-\mathrm{p}\mathrm{q}_\mathrm{f}^\mathrm{S} + F\rho_\mathrm{D}(\ell, \mathrm{q}_\mathrm{f}^\mathrm{S})\right) - \psi\ell, \end{aligned} \tag{20}$$

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:

$$\pi_\lambda^{m*} = \max_{\mathrm{p}\geq 0, \ell \geq 0}\left[(\mathrm{p}-c)\mathrm{q_g}(\mathrm{p}) + \lambda\left(-\mathrm{p}\mathrm{q}_\mathrm{f}^\mathrm{S}(\mathrm{p},\ell) + F\rho_\mathrm{D}(\ell, \mathrm{q}_\mathrm{f}^\mathrm{S}(\mathrm{p},\ell))\right) - \psi\ell\right], \text{ subject to} \tag{21}$$

(9) : ROC curve of diagnostic scheme,

(3) and (11) – (13) : optimal consumer response $\mathrm{q_g}(\mathrm{p}), \mathrm{q}_\mathrm{f}^\mathrm{B}(\mathrm{p},\ell), \mathrm{q}_\mathrm{f}^\mathrm{S}(\mathrm{p},\ell)$,

$v_\mathrm{f} \geq 0, \quad v_\mathrm{g} \geq 0$: nonnegative consumer valuations.

The above optimization problem can be solved by ignoring the constraints initially, but

verifying them *ex post*. The distributor FOCs with respect to p and ℓ are:

$$\partial_p\left((p-c)q_g(p)\right)+\lambda\partial_p\left(-pq_f^S(p,\ell)+F\rho_D(\ell,q_f^S(p,\ell))\right)=0 \tag{22}$$

$$\lambda\partial_\ell\left(-pq_f^S(p,\ell)+F\rho_D(\ell,q_f^S(p,\ell))\right)-\psi=0. \tag{23}$$

Taking into account the optimal consumer responses, these FOCs can be simplified. In particular, re-writing FOC (22) and using the FOC for type-f consumers (7):

$$q_g(p)+(p-c)q_g'(p)+\lambda\underbrace{\partial_{q_f^S}\left(-pq_f^S+F\rho_D(\ell,q_f^S)\right)}_{=0,\text{ from (7).}}\partial_p q_f^S=0$$

$$\Rightarrow\frac{1}{p^2}-\frac{2(p-c)}{p^3}=0$$

$$\Rightarrow p^*=2c. \tag{24}$$

Thus, in equilibrium, the price p^* is determined solely by consumer preferences. With monopolist distributor, the price reflects a monopolistic markup $(p-c)=c$. From (3), the optimal total consumption for a type g or type f consumer is:

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

Next, substituting p^* into (11):

$$\frac{(1-\rho_F(\ell))^{\frac{1}{1-4c^2q_f^S}}}{(1-4c^2q_f^S)^2}\ln(1-\rho_F(\ell))=-\frac{1}{2Fc}. \tag{25}$$

Similarly, the distributor FOC (22) (with respect to ℓ) can be simplified using (7):

$$\underbrace{\partial_{q_f^S}\left(-pq_f^S+F\rho_D(\ell,q_f^S)\right)}_{=0,\text{ from (7).}}\partial_p q_f^S+F\partial_\ell^1\rho_D(\ell,q_f^S)=\frac{\psi}{\lambda},$$

$$\Rightarrow\partial_\ell^1\rho_D(\ell,q_f^S)=\frac{\psi}{F\lambda},$$

where the notation $\partial^1(\cdot,\cdot)$ indicates partial derivative with respect to the first argument. Now, from (9):

$$\frac{(1-\rho_F(\ell))^{\frac{1}{1-4c^2q_f^S}}}{(1-4c^2q_f^S)}\frac{\rho_F{}'(\ell)}{(1-\rho_F(\ell))}=\frac{\psi}{F\lambda}. \tag{26}$$

Solving (25)-(26) gives equilibrium $q_f^{S^*}$ and ℓ^*. In fact, $q_f^{S^*}$ can be expressed as:

$$q_f^{S^*}=\frac{1}{4c^2}\left[1-\frac{2\psi c(1-\rho_F(\ell^*))}{\lambda\rho_F{}'(\ell)_{|\ell^*}}\ln\left(\frac{1}{1-\rho_F(\ell^*)}\right)\right]. \tag{27}$$

A reasonable model of differentiable and increasing function $\rho_F(\cdot)$ is shown in Fig. 4. For this model of false alarm probability, the following expressions hold:

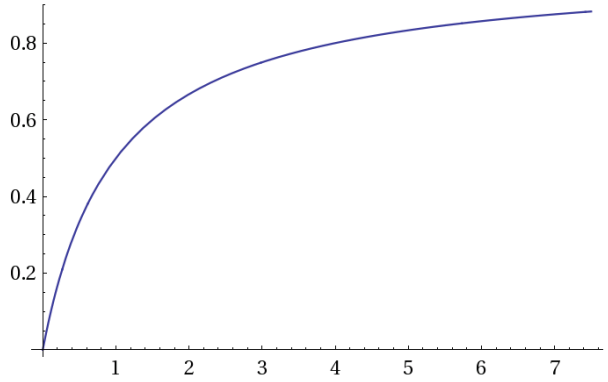

Figure 4. Plot of false alarm probability ρ_F vs. level of investment ℓ according to the model $\rho_F(\ell)=\frac{\ell}{1+\ell}$.

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

The optimal $q_f^{S^*}$ and $q_f^{B^*}$ can be obtained as:

$$q_f^{S^*}=\frac{1}{4c^2}\left[1-\frac{2\psi c(1+\ell^*)}{\lambda}\ln(1+\ell^*)\right],\quad q_f^{B^*}=\frac{1}{4c^2}\frac{2\psi c(1+\ell^*)}{\lambda}\ln(1+\ell^*) \tag{28}$$

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

$$\mathrm{p}^*\beta \ln\left(\frac{\mathrm{p}^*}{F}\left[\beta(1+\ell^*)\right]^2 \ln(1+\ell^*)\right) + \frac{1}{1+\ell^*} = 0, \tag{29}$$

where $\beta \equiv \frac{\psi}{\lambda}$. In summary, (3) and (12)-(13) characterize the consumers' optimal choices $\mathrm{q_g}$, and $\mathrm{q_f^B}, \mathrm{q_f^S}$ for a given distributor choice (p, ℓ), and (24)-(29) characterize a monopolist distributor's optimal choices (p^*, ℓ^*).

Table II lists optimal distributor choices (p^*, ℓ^*) and optimal consumer responses $\mathrm{q_g}^*, \mathrm{q_f^{S*}}$ for some combinations of c, β, and F. *Ceteris paribus* the optimum level of investment ℓ^* increases as F increases, and decreases as ψ increases. For a fixed ψ and F, there exists a cost for which the ℓ^* attains a maximum value and thereafter, ℓ^* decreases with c. Distributor's performance metrics such detection rate ρ_{D}, collection efficiency η, and profit π_λ^m can be computed using these values.

Perfect competition

Consider a benchmark case of perfect competition, where the average per consumer profit is zero, i.e.,

$$\pi_\lambda^c = \mathrm{A} + (\mathrm{p} - c)\mathrm{q_g}(\mathrm{p}) + \lambda\left(-\mathrm{p}\mathrm{q_f^S} + F\rho_{\mathrm{D}}(\ell, \mathrm{q_f^S})\right) - \psi\ell = 0, \tag{30}$$

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

TABLE II

(p^*, ℓ^*) FOR DIFFERENT c, $\alpha = \psi/\lambda$, F.

c	F	$\frac{\psi}{\lambda}$	p^*	ℓ^*	${q_g}^*$	$q_f^{S^*}$
0.1	1	2	0.2	0.113	25	23.8
0.2	1	2	0.4	0.169	6.25	5.33
0.3	1	2	0.6	0.162	2.78	2.19
0.4	1	2	0.8	0.148	1.56	1.16
0.5	1	2	1.0	0.133	1.00	0.717
0.5	*1.1*	2	1.0	0.145	1.00	0.689
0.5	*1.5*	2	1.0	0.190	1.00	0.585
0.5	*2.0*	2	1.0	0.241	1.00	0.464
0.5	*2.5*	2	1.0	0.290	1.00	0.343
0.5	*3.0*	2	1.0	0.335	1.00	0.228
0.5	1	*1*	1.0	0.311	1.00	0.644
0.5	1	*1.5*	1.0	0.195	1.00	0.680
0.5	1	*2.0*	1.0	0.133	1.00	0.717
0.5	1	*2.5*	1.0	0.097	1.00	0.746
0.5	1	*3.0*	1.0	0.073	1.00	0.773

respectively. Following set of conditions lead to zero profit:

$$(p^\dagger - c)q_g(p^\dagger) = 0 \quad \Rightarrow \quad p^\dagger = c, \quad (\textit{using } (3))$$

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

Note that these conditions are not the only conditions that ensure zero distributor profit. With (9) and assuming $\rho_F(\ell) = \frac{\ell}{1+\ell}$, the second condition above can be re-written as:

$$p^{\dagger} q_f^S + F\left(\frac{1}{1+\ell^{\dagger}}\right)^{\frac{1}{1-(p^{\dagger})^2 q_f^S}} = A + F - \frac{\psi \ell^{\dagger}}{\lambda}. \tag{31}$$

Next, substituting $\rho_F(\ell) = \frac{\ell}{1+\ell}$ into fraudulent consumer's FOC (11):

$$\frac{1}{(1-(p^{\dagger})^2 q_f^S)^2}\left(\frac{1}{1+\ell^{\dagger}}\right)^{\frac{1}{1-(p^{\dagger})^2 q_f^S}} = \frac{1}{F p^{\dagger} \ln(1+\ell^{\dagger})} \tag{32}$$

Optimal $(p^{\dagger}, \ell^{\dagger})$ can be determined by feasible solutions of equations (31)–(32) 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:

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

where $y(p^{\dagger}, \ell^{\dagger}) = 0.5\sqrt{\ln(1+\ell^{\dagger})}\left(\sqrt{\ln(1+\ell^{\dagger})} + \sqrt{\ln(1+\ell) + 4\left[(A+F)p^{\dagger} - \beta \ell^{\dagger} p^{\dagger} - 1\right]}\right)$. Plugging $q_f^{S\dagger}$ into (32) gives $\ell^{\dagger}$.

An alternative formulation

A second game-theoretic model, which is only briefly outlined but not fully covered here, assumes that all consumers have the same initial preferences (utility function), and that they make a decision to become fraudulent or stay genuine depending on probability of detection and the amount of fine that they are facing in the case of being caught.

Once the consumers 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 consumers have to make an additional decision, i.e., to choose whether they will be honest or fraudulent.

In turns out that for any fixed fine and detection probability, one will be able to determine what fraction of consumers will be fraudulent in equilibrium. Thus, one can jointly solve the problem of distributor choice of security investment and find the corresponding fraction of consumers that would choose to be fraudulent with a given security investment. Then, the problem becomes identical to the first scenario. This allows the distributor to compute his expected profit as a function of his security investment. Next, if the distributor is a monopolist, he maximizes his profit and chooses equilibrium level of investment in monitoring fraud, i.e., investment for which he achieves the highest profit.

References

[1] S. McLaughlin, D. Podkuiko, and Patrick McDaniel. Energy theft in the advanced metering infrastructure. In *Proceedings of CRITIS 09, 4th International Conference on Critical Information Infrastructures Security*, 2009.

[2] Guidelines for smart grid cyber security: Vol 2., privacy and the smart grid. NIST IR 7628, Aug. 2010.

[3] Michael LeMay and Carl A Gunter. Cumulative attestation kernels for embedded systems. *Smart Grid, IEEE Transactions on*, 3(2):744–760, 2012.

[4] E. De Buda. System for accurately detecting electricity theft. US Patent Application 12/351978, Jan. 2010.

[5] Andrew Appel. Security seals on voting machines: A case study. *ACM Transactions on Information and Systems Security*, 14:1–29, 2011.

[6] Pedro Antmann. Reducing technical and non-technical losses in the power sector. Technical report, World Bank, July 2009.

[7] S. McLaughlin, D. Podkuiko, S. Miadzvezhanka, A. Delozier, and Patrick McDaniel. Multi-vendor penetration testing in the advanced metering infrastructure. In *Proceedings of the Annual Computer Security Applications Conference (ACSAC)*, Dec. 2010.

[8] Mike Davis. Smartgrid device security. adventures in a new medium. `http://www.blackhat.com/presentations/bh-usa-09/MDAVIS/BHUSA09-Davis-AMI-SLIDES.pdf`, July 2009.

[9] Dale Peterson. AppSecDC in review: Real-world backdoors on industrial devices.

http://www.digitalbond.com/2012/04/11/appsecdc-in-review/, April 2012.

[10] Brian Krebs. FBI: smart meter hacks likely to spread. http://krebsonsecurity.com/2012/04/fbi-smart-meter-hacks-likely-to-spread/, April 2012.

[11] Adam Lesser. When big IT goes after big data on the smart grid. http://gigaom.com/cleantech/when-big-it-goes-after-big-data-on-the-smart-grid-2/, March 2012.

[12] Chet Geschickter. *The Emergence of Meter Data Management (MDM): A Smart Grid Information Strategy Report*. GTM Research, 2010.

[13] CJ Bandim, JER Alves Jr, AV Pinto Jr, FC Souza, MRB Loureiro, CA Magalhaes, and F. Galvez-Durand. Identification of energy theft and tampered meters using a central observer meter: a mathematical approach. In *Transmission and Distribution Conference and Exposition, 2003 IEEE PES*, volume 1, pages 163–168. IEEE, 2003.

[14] A. Nizar and Z. Dong. Identification and detection of electricity customer behaviour irregularities. *Power Systems Conference and Exposition (PSCE)*, pages 1–10, March 2009.

[15] J. Nagi, K. S. Yap, S. K. Tiong, S. K. Ahmed, and M. Mohamad. Nontechnical loss detection for metered customers in power utility using support vector machines. *IEEE Transactions on Power Delivery Systems*, 25(2):1162–1171, April 2010.

[16] S.S.S.R. Depuru, Lingfeng Wang, and V. Devabhaktuni. Support vector machine based data classification for detection of electricity theft. In *Power Systems Conference and Exposition (PSCE), 2011 IEEE/PES*, pages 1 –8, march 2011.

[17] Stephen McLaughlin, Brett Holbert, Ahmed Fawaz, Robin Berthier, and Saman Zonouz. A multi-sensor energy theft detection framework for advanced metering

infrastructures. *Selected Areas in Communications, IEEE Journal on*, 31(7):1319–1330, 2013.

[18] Daisuke Mashima and Alvaro A Cárdenas. Evaluating electricity theft detectors in smart grid networks. In *Research in Attacks, Intrusions, and Defenses*, pages 210–229. Springer, 2012.

[19] Rong Jiang, Rongxing Lu, Ye Wang, Jun Luo, Changxiang Shen, and Xuemin Sherman Shen. Energy-theft detection issues for advanced metering infrastructure in smart grid. *Tsinghua Science and Technology*, 19(2):105–120, 2014.

[20] Stephen E. McLaughlin, Patrick McDaniel, and William Aiello. Protecting consumer privacy from electric load monitoring. In *ACM Conference on Computer and Communications Security*, pages 87–98, 2011.

[21] G. Kalogridis, C. Efthymiou, S.Z. Denic, T.A. Lewis, and R. Cepeda. Privacy for smart meters: Towards undetectable appliance load signatures. In *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on*, pages 232 –237, oct. 2010.

[22] D.C. Bergman, Dong Jin, J.P. Juen, N. Tanaka, C.A. Gunter, and A.K. Wright. Distributed non-intrusive load monitoring. In *Innovative Smart Grid Technologies (ISGT), 2011 IEEE PES*, pages 1 –8, jan. 2011.

[23] John G. Kassakian and Richard Schmalensee. The future of electric grid: An interdisciplinary MIT study. Technical report, Massachusetts Institute of Technology, 2011.

[24] Paul L. Joskow. Incentive regulation in theory and practice: Electricity distribution and transmission networks. In *Economic Regulation and Its Reform: What Have We*

Learned?, NBER Chapters. National Bureau of Economic Research, Inc, Julio Dic 2011.

[25] Ingo 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 Univ Pr, 2002.

[27] Patrick Bolton and Mathias Dewatripont. *Contract Theory*, volume 1 of *MIT Press Books*. The MIT Press, 2005.

[28] Mark Armstrong and David E.M. Sappington. *Recent Developments in the Theory of Regulation*, volume 3 of *Handbook of Industrial Organization*, chapter 27, pages 1557–1700. Elsevier, 2007.

Author Information

Saurabh Amin 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 (UC), Berkeley, MS from the University of Texas at Austin, and B.Tech. from IIT Roorkee.

Galina A. Schwartz is a Researcher in the Department of Electrical Engineering and Computer Sciences at UC 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 MS in mathematical physics from Moscow Institute of Engineering Physics and 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 Universidad de los Andes. Prior to joining UT Dallas he was a Researcher with Fujitsu Laboratories of America and a postdoctoral scholar at UC Berkeley. His research interests include information security, cyber-physical systems, the smart grid, and intrusion detection.

S. Shankar Sastry is dean of Engineering at 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 telesurgery, 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.

Corresponding address: Massachusetts Institute of Technology, 77 Massachusetts Avenue 1-241, Cambridge, MA 02139, U.S.A. Email: amins@mit.edu, Telephone: 617-253-8003.

Sidebar 1: 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 consumers (attackers) choose q_f^S such that the marginal payment for small increase in billed quantity of electricity is equal to the expected marginal fine for the small increase in q_f^S see FOCs 7. This holds for the case when consumer valuations for total electricity consumed are same for fraudulent and genuine consumers. Essentially, the choice of q_f^S determines how different the distribution of meter measurements sent by type-f consumers will be from the distribution of measurements sent by type-g consumers.

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 formulation forms the basis of recent work [18], where electricity traces for 108 residential consumers were obtained from a real AMI deployment.

The design of 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 he can steal.

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 total loss due to undetected attacks to the distributor, 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 dataset poisoning attacks (where the attacker can feed our profiling algorithm malicious data).

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 assumption on the adversary model, where the fraudulent consumers continue to use electricity as genuine consumers, but will try to send lower meter readings to the utility in order to minimize the amount of electricity they need to pay. 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 consumers [20], [21]. The idea of these schemes is to shape the electricity usage signal to prevent inferences that can be made with Non Intrusive Load Monitors (NILMs)

[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 consumer choices.

Sidebar 2: 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 the following objectives: 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-run cost and long-run investment by the distributor [24], [25].

Design of regulatory requirements would be an easy task if the regulator were perfectly informed about the distributor's costs and the consumer demand [26]. In reality, the distributor has an informational advantage over the regulator about both aggregate consumer demand and its own operational costs; see Fig. S2. In such cases, the regulatory design can become extremely subtle and fragile to changes of regulator's assumptions about distributor's efficiency and costs. A well-developed body of work dedicated to designing optimal regulatory policies of a monopolistic distributor who has

privileged information about his technological capabilities and customers' demand, and when the regulator has well-defined inter-temporal 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 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, i.e., it must perform "reasonably well" under broad conditions, although such a policy may be sub-optimal in each particular case [28]. Two main regulatory regimes have been adapted for distribution utilities: (i) rate of return (dominant regime in USA) and (ii) price cap (dominant regime in European Union and developing countries).

Below each regime is outlined briefly. This article considers 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 pre-specified a rate-of-return, and the tariff structures for the electricity which are adjusted as the distributor's cost change to ensure that he 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 the 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 distributor faces limited risks of expropriation of his sunk investments by the regulator, upgrades of distribution network can be sustained in this form of regulation. The investment incentives of regulated distributor are especially important, since the infrastructure upgrades (e.g., capacity expansion) and modernization (e.g., 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 pre-specified 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 pre-defined 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 price cap does not directly provide incentives for long-term investments, such as distribution network upgrades and reduction of non-technical losses. Similarly, price cap does not incentivize the distributor to choose optimal allocation of service quality. To remedy this, additional requirement 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 (e.g., consumer theft) at socially optimal levels.

When the pricing flexibility of price cap regulation is combined with the rewards (resp. punishments) for performance improvement (resp. deterioration) relative to the regulator's benchmark, the resulting regime is termed *performance-based* (or incentive) regulation. Indeed, in the face of 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 distributor's incentives by decoupling regulated price structure from the need to know the exact operating / maintenance costs.

Regulated distributor

This article presented distributor's optimal choices (p, ℓ) for the case of an unregulated monopolist and the case of perfect competition. This analysis can be extended to 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 $F^r(\cdot)$, and the investment level ℓ in fraud monitoring, 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 audit and regulatory process followed by the regulator.

The average revenue can be computed based on the total quantity Q_T provisioned by the distributor, or the quantity $(Q_T - Q_f^S)$ which excludes the stolen quantity of

electricity; see (17)-(18). 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:

$$\frac{\mathrm{R}_{\lambda}(\mathrm{p},\ell)}{\mathrm{Q}_{\mathrm{T}}(\mathrm{p})} \leq \bar{\mathrm{p}}, \tag{33}$$

$$\frac{\mathrm{R}_{\lambda}(\mathrm{p},\ell)}{\left(\mathrm{Q}_{\mathrm{T}}(\mathrm{p}) - \mathrm{Q}_{\mathrm{f}}^{\mathrm{S}}(\mathrm{p},\ell)\right)} \leq \bar{\mathrm{p}}, \tag{34}$$

where the total revenue $\mathrm{R}_{\lambda}(\mathrm{p},\ell)$ is given by (19). Clearly, (34) imposes a stricter regulatory imposition on the distributor. In the case of (34), the regulator does not account for the fraudulent consumers' surplus resulting from the successfully stolen electricity $\mathrm{Q}_{\mathrm{f}}^{\mathrm{S}}$. From the viewpoint of distributor (resp. regulator), the constraint (33) (resp. (34)) is more desirable because it eases (resp. tightens) the regulatory constraint. Using (17) and (19), and for special case $\mathrm{F}^{\mathrm{r}}(\cdot) \approx F$, constraints (33) and (34) can be re-written as:

$$-\mathrm{p}^2\mathrm{q}_{\mathrm{f}}^{\mathrm{S}} + F\mathrm{p}\rho_{\mathrm{D}}(\ell, \mathrm{q}_{\mathrm{f}}^{\mathrm{S}}) \leq \frac{1}{\lambda}\left(\frac{\bar{\mathrm{p}}}{\mathrm{p}} - \mathrm{A}\mathrm{p} - 1\right)$$

$$\mathrm{p}\mathrm{q}_{\mathrm{f}}^{\mathrm{S}}\left(-\mathrm{p} + \bar{\mathrm{p}}\rho_{\mathrm{D}}(\ell, \mathrm{q}_{\mathrm{f}}^{\mathrm{S}})\right) + F\mathrm{p}\rho_{\mathrm{D}}(\ell, \mathrm{q}_{\mathrm{f}}^{\mathrm{S}}) \leq \frac{1}{\lambda}\left(\frac{\bar{\mathrm{p}}}{\mathrm{p}} - \mathrm{A}\mathrm{p} - 1\right).$$

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 (21), subject to ROC curve (9), consumer responses (3) and (11)-(13), and the average revenue constraint (33) or (34).

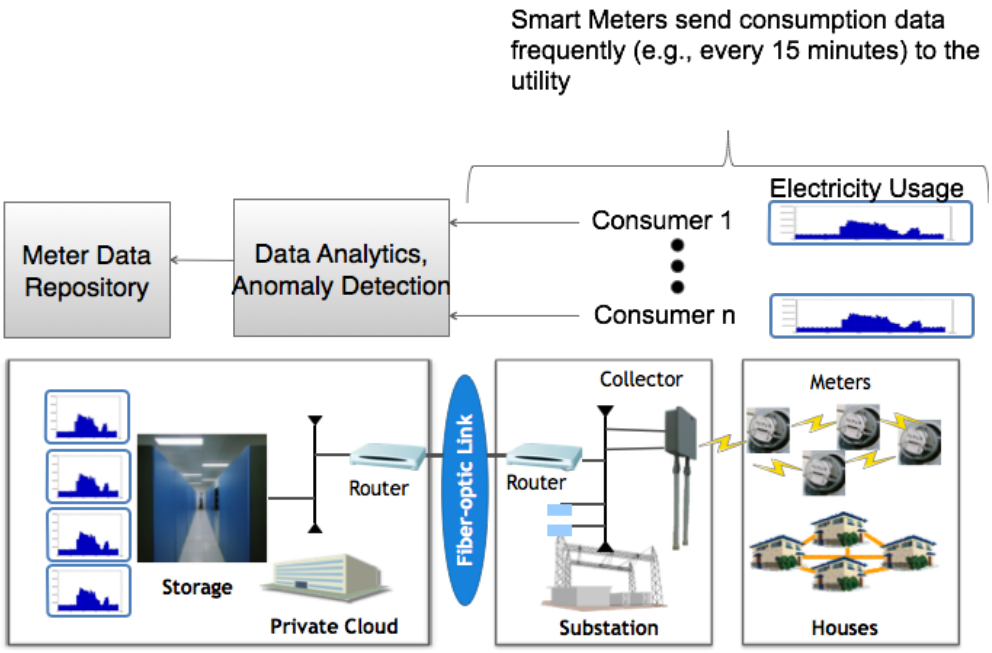

Figure 5. Centralized Meter Data Management (MDM). When electric utilities deploy Advanced Metering Infrastructures they also need to deploy a Meter Data Management (MDM) system in their backend to manage smart meter data storage and analytics (forecasting, anomaly detection etc.) One of the key services offer 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 physical personnel checking for tamper-evident seals. report reprogramming or tampering attempts.

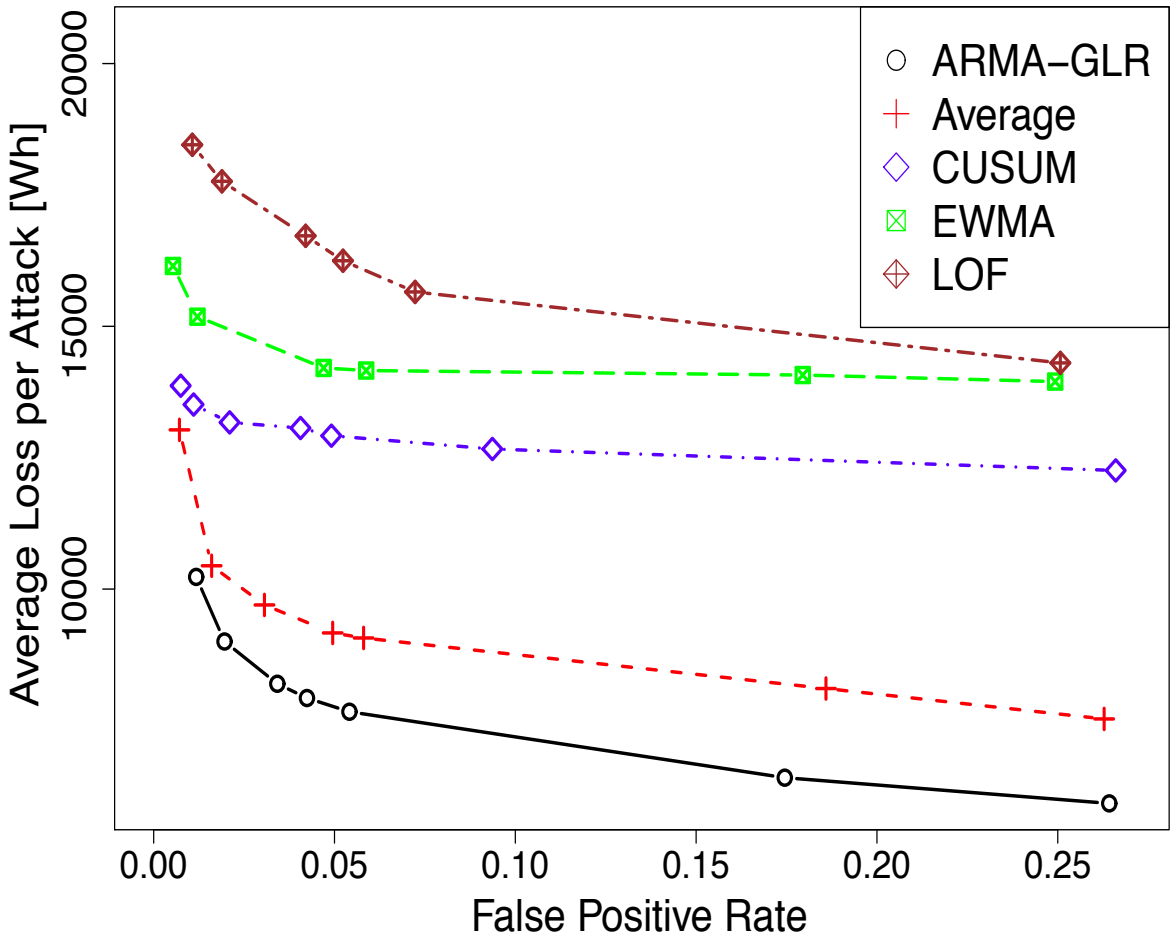

Figure S1. Performance evaluation of detection schemes: i) Auto-Regressive Moving Average with a Generalized Likelihood Ratio (ARMA-GLR) test; ii) simple average consumption test; iii) Non-parametric CUmulative SUm (CUSUM) algorithm; iv) Exponential Weighted Moving Average (EWMA) detector; and v) Outlier detection algorithm called Local Outlier Factor (LOF).

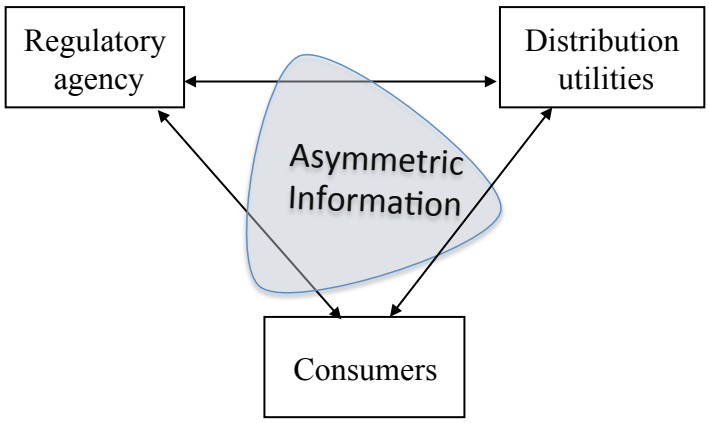

Figure S2. Players in regulated electricity distribution. A central issue in regulation of distribution utilities is the presence of asymmetric information between the three entities affected the electricity distribution system: the regulator, distributors, and end consumers.