

Privacy and Customer Segmentation in the Smart Grid

Lillian J. Ratliff, Roy Dong, Henrik Ohlsson, Alvaro A. Cárdenas and S. Shankar Sastry

Abstract—In the electricity grid, networked sensors which record and transmit increasingly high-granularity data are being deployed. In such a setting, privacy concerns are a natural consideration. We present an attack model for privacy breaches, and, using results from estimation theory, derive theoretical results ensuring that an adversary will fail to infer private information with a certain probability, independent of the algorithm used. We show utility companies would benefit from less noisy, higher frequency data, as it would improve various smart grid operations such as load prediction. We provide a method to quantify how smart grid operations improve as a function of higher frequency data. In order to obtain the consumer’s valuation of privacy, we design a screening mechanism consisting of a menu of contracts to the energy consumer with varying guarantees of privacy. The screening process is a means to segment customers. Finally, we design insurance contracts using the probability of a privacy breach to be offered by third-party insurance companies.

I. INTRODUCTION

Increasingly advanced metering infrastructure (AMI) is replacing older technology in the electricity grid. Smart meters send detailed information about consumer electricity usage over a network every half-hour, quarter-hour, or in some cases, every five minutes. This high-granularity data is needed to support energy efficiency efforts as well as demand-side management. However, improper handling of this information could also lead to unprecedented invasions of consumer privacy [1], [2].

Given that smart grid operations inherently have privacy and security risks [2], it would benefit the utility company, to know the answer to the following questions: How do consumers in the population value privacy? How can we quantify privacy? How do privacy-aware policies impact smart grid operations? In this paper we address these questions as well as expose new directions for future research on privacy and customer segmentation in the smart grid.

Using our results on the fundamental limits of non-intrusive load monitoring [3], we are able to come up with probabilities for the success of an attack by an adversarial agent independent of the algorithm. Then using these probabilities we can design a screening mechanism consisting of a menu of contracts to be offered to consumers. One set of contracts to be offered by the utility company assess how the consumer values privacy thereby revealing his preferences. Based on their valuation of privacy as a good,

L. J. Ratliff, R. Dong, H. Ohlsson, and S. S. Sastry are with Faculty of Electrical Engineering and Computer Sciences, University of California, Berkeley, Berkeley, CA, 94707, USA {ratliff1, roydong, ohlsson, sastry}@eecs.berkeley.edu

A. A. Cárdenas is with the Department of Computer Science, University of Texas, Dallas, Dallas, TX 75080, USA alvaro.cardenas@utdallas.edu

consumers can select the quality of the service contract with the utility company. Essentially, electricity service is offered as a product line differentiated according to privacy where consumers can select the level of privacy that fits their needs and wallet. The screening process is a way to do customer segmentation the result of which can lead to targeting.

In particular, using knowledge of consumer preferences, the utility company could then incentivize consumers based on their preferences to choose a low privacy setting which helps increase the granularity of data for use by the utility company for programs like demand response, direct load control, etc. In addition, third-party insurance companies can design insurance contracts. Insurance allows the consumer to protect himself in the event of a privacy breach, i.e. she will be compensated for any experienced loss.

The paper is organized as follows. In Section II we review the problem of non-intrusive load monitoring (NILM) and show how NILM leads to our novel privacy metric. We show in an example that the probability of an adversary successfully implementing a privacy breach decreases with a decrease in sampling rate. We discuss the impact of sampling rate on smart grid operations in Section III. In Section IV we use the privacy metric to design a screening mechanism that consists of privacy contracts between the consumer and the utility company. Similarly, in Section V we use the privacy metric to design insurance contracts. Finally, in Section VI we summarize the results and discuss future research directions.

II. PRIVACY GUARANTEES

In this section, we discuss our metric for privacy, and guarantees of privacy under this metric. For this paper, we restrict the scope of our analysis to data collection policies. Another important aspect of privacy is how data retention policies can alter privacy and smart grid performance. Such a topic is reserved for future research.

A. Nonintrusive Load Monitoring

Our formulation of privacy builds on recent research into nonintrusive load monitoring (NILM) algorithms, first proposed by Hart [4]. The goal of NILM is to use the aggregate power consumption signal, which can be measured by metering infrastructures without the placement of additional sensors inside the home, and make inferences on the load profile. For example, the problem of energy disaggregation is to recover the power consumption signals of individual devices [5]. Another example would be to detect when devices switch on and off, which is often referred to as event-based NILM [6].

There are many approaches to the design of NILM algorithms, including but not limited to hidden Markov models (HMMs) [7], [8], [9], sparse coding methods [10], and dynamical systems approaches [11]. These methods vary in the prior information required; some are completely unsupervised and nonparametric, while others require a large amount of disaggregated data to build a dictionary. However, in our recent work, we provided a unifying framework for modeling all of these algorithms [3].

To the best of our knowledge, every approach to energy disaggregation gives devices a certain kind of model, be it an HMM, dictionary, or dynamical system. Most of these approaches have an input space: for HMMs, the input is a sequence of latent state transitions; for the sparse approaches, it is a sparse vector corresponding to the most representative elements in the dictionary, and in dynamical systems, it is an input signal to the systems.

In all these approaches, a fixed input yields a probability distribution on the device’s power consumption; this distribution is often assumed to be Gaussian. Finally, note that in all the mentioned frameworks, recovering the input is almost the same as recovering the device’s power consumption signal.

Other NILM formulations also can fall into this framework. For example, in event-based NILM, if the task is to determine whether or not the air conditioner is in use, we can consider the probability distribution across aggregate power consumption signals when the air conditioner is on and when the air conditioner is off. In this case, whether or not the air conditioner is on serves as an input.

In our framework, NILM algorithms are abstracted as measurable functions operating on the aggregate power consumption signal, and we ask whether or not the algorithms can successfully distinguish between inputs. For a more comprehensive treatment of this topic, we refer the reader to [3].

B. Privacy Metric

Now, we are ready to introduce our metric of privacy.

One of the common theoretical definitions for a privacy metric is the notion of differential privacy [12], [13]. While differential privacy has many attractive properties, it is most useful when we want to share data via a trusted third party aggregator, or by injecting noise in the original messages sent to a third party; however, for many practical, regulatory, dispute resolution, performance, or business reasons, there will always be several cases where we need to get access to the raw data, and in these cases differential privacy will not help us identify a good security mechanism to prevent raw data from being compromised.

In contrast, we fix a definition of a privacy breach where the user has a set of possible inputs, and he wishes to keep the true input private. For example, the definition of privacy breach might be whether or not an adversary knows if the user is doing dishes in the dishwasher, watching TV, or exercising on a treadmill in the evening. This notion of equivocation is related to recent work in privacy who measures

privacy not with differential privacy but with equivocation metrics [14].

We also assume a powerful adversary that knows the probability distribution of the energy consumption of digital signals, and their prior probability for being active in an aggregated signal. As an observation, we assume the adversary has access to the aggregate power consumption signal, but not the device-level signals, for a building.

More formally, suppose there are two inputs u_1 and u_2 . For each input, the aggregate power consumption signal follows distributions F_1 and F_2 , respectively. In this paper, we assume these distributions to be Gaussian; however, we consider the general case in [3].

Suppose the distributions have means μ_1 and μ_2 , and both distributions have the same covariance $\sigma^2 I$. Let $a = \sigma^{-2}(\mu_0 - \mu_1)$ and $b = \frac{1}{2\sigma^2}(\|\mu_1\|_2^2 - \|\mu_2\|_2^2)$. Suppose the adversary uses any estimator \hat{u} and suppose the events $\{u = u_1\}$ and $\{u = u_2\}$ are equally likely. Then, the probability of our adversary successfully distinguishing two inputs is bounded by

$$P(\hat{u} = u) \leq \frac{1}{2} \left(1 - \operatorname{erf} \left(\frac{-\frac{1}{\|a\|_2} (a^T \mu_0 + b)}{\sqrt{2\sigma^2}} \right) \right) \quad (1)$$

where erf is the Gauss error function. More details can be found in [3].

III. SMART GRID OPERATIONS

In Section II, we developed a metric for privacy. If privacy is the only thing of concern, a trivial solution is to record nothing transmit nothing. However, the utility company has other objectives than just preserving the privacy of its consumers. Hence, privacy issues arise because the sensitive data has other uses. Such polices as noise injection or varying the sampling rate can be employed to protect against privacy breaches while still allowing the utility company to operate.

In advanced metering infrastructures, the data is used to improve the performance of smart grid operations. How smart grid operations degrade under different metering policies is an active topic of research; for preliminary investigations, see [15], [16]. Intuitively, the performance will degrade as fewer samples are collected or more noise is added. We attempt to quantify this degradation. In this section, we develop an direct load control example to demonstrate how smart grid operations performance is affected by different sampling rates.

A. Direct Load Control

The problem of direct load control has recently been studied as viable option to improve smart grid operations [17], [18].

Generator output is generally determined by two processes: unit commitment and economic dispatch. *Unit commitment* is done in advance, and sets the generator ramping schemes. *Economic dispatch* is done online, and determines the output levels of generators that are already online to meet total demand.

When the demand exceeds the output capacity of all online generators, economic dispatch schemes will use generators with quick ramp-up times to ensure stability of the power grid. These generators are very inefficient. One goal of direct load control (DLC) as an economic dispatch scheme is to reduce the deviation of actual demand from the forecasted demand.

Consider the direct load control model:

$$x_{k+1} = x_k + u_k + \mu_k + d_k \quad (2)$$

Here, $x_k \in \mathbb{R}$ represents the power consumption of a unit at time k , where a unit can be a household, an HVAC system for a building, or a sector of the power grid. $u_k \in \mathbb{R}$ represents the direct load control signal at time k . $\mu_k \in \mathbb{R}$ represents the affine term which generates our nominal demands at time k ; if $u_k \equiv 0$ and $d_k \equiv 0$, then μ_k creates our forecasted demand. Finally, d_k represents the disturbance at time k . In this model, disturbances from the nominal demand persist, and DLC policies must be employed to return the power consumption to the nominal demand.

Now, consider different sampling rates. That is, we suppose our controller is only able to receive measurements every N time steps. However, it is still able to issue control commands at every time step. We wish to design a controller that makes use of the available measurements to optimally issue control commands to a sector of the power grid.

The subsampled system can be modeled in a Markov jump linear system (MJLS) framework. To define optimality, we consider the \mathcal{H}_∞ norm of MJLSs, as defined in [19], [20]. In our application, the \mathcal{H}_∞ norm represents a worst case estimate of how much the true power consumption will deviate from the power consumption used for unit commitment. This worst case estimate is a function of the uncertainty in the load forecast.

Recent results in the analysis of MJLS gives us the optimal \mathcal{H}_∞ controller for subsampled scalar systems [16]. Thus, we can analyze how the performance of direct load control is affected by different sampling schemes. For example, Figure 1 gives us the performance for equidistant sampling. Thus, the formulation allows us to quantify the value of a higher sampling rate to the utility company.

IV. PRIVACY CONTRACTS

In this section, we discuss how the utility company can design a screening mechanism in order to assess the consumer's unknown type. This is a mechanism design problem with asymmetric information. The utility company designs a screening mechanism whose contracting device is the privacy setting offered to the consumer. This screening process can be thought of as customer segmentation since it will extract each consumer's type after which the consumers can be grouped according to their type.

We consider a model in which there are only two types and we utilize standard results from the theory of screening (see, e.g., [21]) to develop a framework for designing privacy contracts. We remark that as a result of the screening process the utility company will know how each consumer values

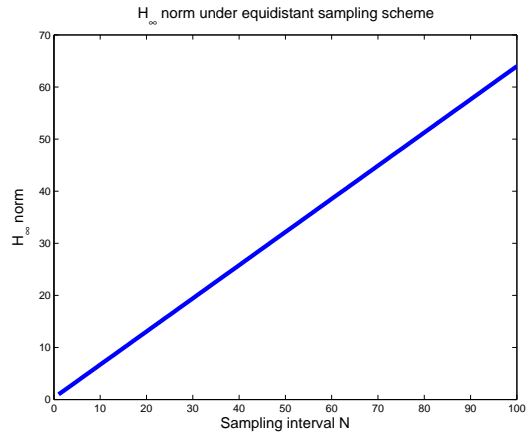


Fig. 1. The \mathcal{H}_∞ norm for the equidistant sampling scheme as a function of the sampling interval N . A higher sampling interval N corresponds to a lower sampling frequency, i.e. the utility company receives less data.

privacy and can leverage that in the design of incentives aimed at inducing the consumer to select a privacy setting more desirable from the perspective of the utility company.

A. Two Types: High-Privacy and Low-Privacy Settings

We model privacy-settings on smart meters as a good. The *quality* of the good is either a high-privacy setting q_H or a low-privacy setting q_L . The consumer can choose either a high privacy setting or a low privacy setting, i.e. the consumer selects $q \in \mathcal{Q} = \{q_H, q_L\} \subset \mathbb{R}$ where $q_L \leq q_H$ and $-\infty < q_L < q_H < \infty$. The consumer's valuation of privacy is his *type* which takes values $\theta \in \{\underline{\theta}, \bar{\theta}\} \subset \mathbb{R}$ where $\bar{\theta}$ represents how much the consumer values high-privacy over low-privacy and $\underline{\theta} < \bar{\theta}$. We assume that type θ is distinct from the private information itself; by this we mean that how much the consumer values privacy is not also private information. We note here that these types implicitly make use of the probabilities presented in Section II.

The consumer's type θ is related to his willingness to pay in the following way: if the utility company announces a price t for choosing q , the type-dependent consumer's utility is equal to zero if he does not select a privacy setting q , and it is

$$U(q, \theta) - t \geq 0 \quad (3)$$

if he does select a privacy setting. The case in which the consumer does not select a privacy setting is considered the *opt-out* case in which consumer exercises his right to not participate. The inequality in (3) is often called the *individual rationality* constraint and in the design of the privacy contracts we will enforce it in order to make sure that all consumers will opt-in. The function $U : \mathbb{R} \times \Theta \rightarrow \mathbb{R}$ is assumed to be strictly increasing in (q, θ) , concave in q , and represents the consumer's preferences.

Since we have only two types, the contracts offered will be indexed by the privacy settings q_L and q_H . Further, as we mentioned the consumer can opt-out by not selecting a privacy option at all. Hence, we need to constrain the

mechanism design problem by enforcing the inequality given in Equation (3) for each value of $\theta \in \{\underline{\theta}, \bar{\theta}\}$. In addition, we need to enforce *incentive-compatibility* constraints

$$U(q_H, \bar{\theta}) - t_H \geq U(q_L, \bar{\theta}) - t_L \quad (4)$$

and

$$U(q_L, \underline{\theta}) - t_L \geq U(q_H, \underline{\theta}) - t_H \quad (5)$$

where the first inequality says that given the price t_H a consumer of type $\bar{\theta}$ should prefer the high-privacy setting q_H and the second inequality says that given the price t_L a consumer of type $\underline{\theta}$ should prefer the low-privacy setting q_L .

The utility company has unit utility

$$v(q, t) = -g(q) + t \quad (6)$$

where we assume that the function $g : \mathcal{Q} \rightarrow \mathbb{R}$ is the unit cost to the utility company for the privacy setting q . We assume that it is a strictly increasing, continuous function which is reasonable because, as we have mentioned in Section III, a low-privacy setting q_L provides the utility company with the high-granularity data it needs to efficiently operate and maintain the smart grid. For example, recall Section III-A in which we show that the performance of DLC degrades with decreases in sampling rate.

The screening problem is to design the contracts, i.e. $\{(t_L, q_L), (t_H, q_H)\}$ where $t_L, t_H \in \mathbb{R}$, so that the utility companies expected profit is maximized where the expected profit is

$$\Pi(t_L, q_L, t_H, q_H) = (1-p)v(q_L, t_L) + pv(q_H, t_H) \quad (7)$$

where $p = P(\theta = \bar{\theta}) = 1 - P(\theta = \underline{\theta}) \in (0, 1)$ where $P(\cdot)$ denotes probability.

To find the optimal pair of contracts, we solve the following optimization problem:

$$\max_{\{(t_L, q_L), (t_H, q_H)\}} \Pi(t_L, q_L, t_H, q_H) \quad (\text{P-1})$$

$$\text{s.t.} \quad U(q_H, \bar{\theta}) - t_H \geq U(q_L, \bar{\theta}) - t_L \quad (\text{IC-1})$$

$$U(q_L, \underline{\theta}) - t_L \geq U(q_H, \underline{\theta}) - t_H \quad (\text{IC-2})$$

$$U(q_L, \underline{\theta}) - t_L \geq 0 \quad (\text{IR-1})$$

$$U(q_H, \bar{\theta}) - t_H \geq 0 \quad (\text{IR-2})$$

$$q_L \leq q_H$$

Depending on the form of $U(q, \theta)$ and $g(q)$ problem (P-1) can be difficult to solve. So, we examine the constraints and try to eliminate as many as we can.

First, we show that (IR-1) is active. Indeed, suppose not. Then, $U(q_L, \underline{\theta}) - t_L > 0$ so that, from the first incentive compatibility constraint (IC-1), we have

$$U(q_H, \bar{\theta}) - t_H \geq U(q_L, \bar{\theta}) - t_L \geq U(q_L, \underline{\theta}) - t_L > 0 \quad (8)$$

where the second to last inequality holds since $U(q, \theta)$ is increasing in θ by assumption. As a consequence, the utility company could increase the price for both types since neither incentive compatibility constraint would be active. This would lead to an increase in the utility company's payoff, i.e. a contradiction. Now, since $U(q_L, \underline{\theta}) = t_L$, the last

inequality in (8) is equal to zero. This implies that (IR-2) is redundant. Further, this argument implies that the constraint (IC-1) is active. Indeed, again suppose not. Then,

$$U(q_H, \bar{\theta}) - t_H > U(q_L, \bar{\theta}) - t_L \geq U(q_L, \underline{\theta}) - t_L = 0 \quad (9)$$

so that it would be possible for the utility company to decrease the incentive t_H without violating (IR-2).

Now, let us assume that the marginal gain from raising the value of the privacy setting q is greater for type $\bar{\theta}$, i.e. $U(q, \bar{\theta}) - U(q, \underline{\theta})$ is increasing in q . Then, since (IC-1) is active, we have

$$t_H - t_L = U(q_H, \bar{\theta}) - U(q_L, \bar{\theta}) \geq U(q_H, \underline{\theta}) - U(q_L, \underline{\theta}). \quad (10)$$

This inequality implies that we can ignore (IC-2). Further, since U is increasing in (q, θ) and we have assumed that $\bar{\theta} > \underline{\theta}$, we can remove the constraint $q_L \leq q_H$. We have reduced the constraint set to

$$t_H - t_L = U(q_H, \bar{\theta}) - U(q_L, \bar{\theta}) \quad (11)$$

$$t_L = U(q_L, \underline{\theta}) \quad (12)$$

Thus, the optimization problem becomes

$$\max_{(q_L, q_H)} \left\{ p(U(q_H, \bar{\theta}) - g(q_H) - U(q_L, \bar{\theta}) + U(q_L, \underline{\theta})) \right. \\ \left. (1-p)(U(q_L, \underline{\theta}) - g(q_L)) \right\} \quad (\text{P-2})$$

This reduces further to two independent optimization problems:

$$\max_{q_H} \{U(q_H, \bar{\theta}) - g(q_H)\} \quad (\text{P-3a})$$

$$\max_{q_L} \{-p(U(q_L, \bar{\theta}) - U(q_L, \underline{\theta})) + \\ (1-p)(U(q_L, \underline{\theta}) - g(q_L))\} \quad (\text{P-3b})$$

B. Direct Load Control Example

Recall that the unit gain the utility company gets out of the privacy setting q is a function $g : \mathcal{Q} \rightarrow \mathbb{R}$. In this section, we discuss a particular example in which g is a metric for how access to high-granularity data affects direct load control. In Section III-A, Figure 1 shows how that as you decrease the sampling rate (increase the sampling interval) the performance degrades, i.e. the \mathcal{H}_∞ norm increases, and it degrades in a linear way. Hence, this motivates a choice for g such that $g(q_L) > g(q_H)$ and decreases in a linear way. Hence, for this example, let

$$g(q) = \zeta q \quad (13)$$

where $0 < \zeta < \infty$. Note that a decreased sampling rate corresponds to a higher privacy setting. The function g as defined is increasing in q so that $g(q_L) > g(q_H)$.

Assume that the consumer's utility is given by

$$U(q, \theta) = \frac{1}{2}(\bar{q}^2 - (q - \bar{q})^2) \theta \quad (14)$$

where $q \in [0, \bar{q}]$ so that it is proportional to how close they are to the maximum privacy setting \bar{q} , and their type θ . Suppose

that $\theta \in \{\underline{\theta}, \bar{\theta}\}$ where $0 < \underline{\theta} < \bar{\theta}$. Note that at $\bar{\theta}$ the utility is $\frac{1}{2}\bar{q}^2$ and at $\underline{\theta}$ the utility is zero. U satisfies the assumption that it is increasing. Let $p = P(\theta = \bar{\theta})$. Then, the optimal solutions to the screening problem are

$$(q_H^*, q_L^*) = \left(\bar{q} - \frac{\zeta}{\bar{\theta}}, \left[\bar{q} + \frac{(1-p)\zeta}{(p\bar{\theta} - \underline{\theta})} \right]_+ \right) \quad (15)$$

where $q_L^* = 0$ when the probability p is greater than the critical point $p^* = \bar{\theta}/\underline{\theta}$. The optimal prices t_H^*, t_L^* can be found by plugging (q_H^*, q_L^*) into (11) and (12). If the utility company knew the types, then the optimal solution would be

$$(q_H^\dagger, q_L^\dagger) = \left(\bar{q} - \frac{\zeta}{\bar{\theta}}, \bar{q} - \frac{\zeta}{\underline{\theta}} \right) \quad (16)$$

In Figure 2, we show that as the probability of the high-




Fig. 2. Comparison between full information and asymmetric information solutions as a function of p the probability of the high-type in the population.

type being drawn from the population increases, q_L^* decreases away from the optimal full information solution q_L^\dagger . This occurs until $p = p_0 = \underline{\theta}/\bar{\theta}$ which is the critical probability. After p_0 , $q_L^* = 0$ and remains there until $p = 1$.

The social welfare is defined to be sum of the pay-off to the utility company and to the consumer. The social welfare is given by

$$\Psi^*(p) = \Pi(t_L^*, q_L^*, t_H^*, q_H^*) + p(U(q_H^*, \bar{\theta}) - t_H^*). \quad (17)$$

Let

$$\begin{aligned} \varphi = (1-p) & \left(\frac{1}{2} \left(\bar{q}^2 - \left(\frac{(1-p)\zeta}{p\bar{\theta} - \underline{\theta}} \right)^2 \right) \underline{\theta} \right. \\ & \left. - \zeta \left(\bar{q} - \frac{(1-p)\zeta}{p\bar{\theta} - \underline{\theta}} \right) \right). \end{aligned} \quad (18)$$

Then,

$$\begin{aligned} \Psi^*(p) = & -p\zeta \left(\bar{q} - \frac{\zeta}{\bar{\theta}} \right) \\ & + \begin{cases} \varphi, & p \leq p_0 \\ p \left(\frac{1}{2} \left(\bar{q}^2 - \left(\frac{(1-p)\zeta}{p\bar{\theta} - \underline{\theta}} \right)^2 \right) (\bar{\theta} - \underline{\theta}) \right), & p > p_0 \end{cases} \end{aligned} \quad (19)$$

The social welfare reaches a critical point at p_0 beyond which the utility company will exclude the low-type from the market and only provide privacy contracts to the high-type. This is called the *shutdown solution*. It is reasonable

that as soon as the probability of the utility company facing a consumer of high-type reaches a critical point, they will focus all their efforts on this type of consumer since a high-type desires a higher privacy setting which results in a degradation of the DLC scheme.

We remark that people who value high privacy more need to be compensated more to participate in the smart grid. If there are two contracts, then even consumers who do not value privacy much will have an incentive to lie. Through the screening mechanism, the consumer will report his type truthfully.

V. PRIVACY INSURANCE CONTRACTS

In this section, we will design an insurance contract, to be offered by a third-party company to the consumer, that uses the probability η that an adversary will fail to infer private information about the consumer. In the previous section we designed contracts to get consumers to allow for lower privacy settings; in this section we design insurance contracts that allow consumers to purchase protection against attacks given they know the probability of a successful attack occurring. The analysis that follows is well known in the economics literature (see, e.g., [22], [23], [24]). Using the theory of insurance contracts when there is asymmetric information and the probability that an adversary can gain access to consumers' private information, we analyze both the consumer's choice on how much insurance to invest in as well as the insurer's decision about which contracts to offer to a population with both high- and low-risk consumers.

A. Analysis of Consumer's Decision

Let us start by analyzing the decision the consumer would make about selecting an amount of insurance given knowledge of η . Let the consumer's utility function be denoted by $\tilde{U} : \mathbb{R} \rightarrow \mathbb{R}$ and assume that \tilde{U} is increasing, twice differentiable and strictly concave. Let us suppose that the consumer is *risk-averse* which means that the consumer, who makes a decision under uncertainty, will try to minimize the impact of the uncertainty on her decision.

In addition, suppose the consumer has initial wealth y , runs the risk of loss ℓ with probability $1 - \eta$. In the context of our problem, wealth represents *private information* that can be gained through analysis of consumer energy consumption data and loss represents exposure of this private information. Recall that $1 - \eta$ is the probability with which an adversarial agent could gain access to private information about a consumer through access to their consumption data.

The consumer must decide how much insurance to buy. Let the cost of one unit of insurance be c and suppose that the insurer pays the consumer β in the event that an adversary attacks them resulting in an exposure of private information where β is the amount of insurance the consumer agrees to buy. Then the consumer wants to solve the following optimization problem:

$$\max_{\beta \geq 0} \{ \eta \tilde{U}(y - \beta c) + (1 - \eta) \tilde{U}(y + (1 - c)\beta - \ell) \} \quad (\text{P-4})$$

Suppose that β^* is a local optimum, then there exists a Lagrange multiplier λ such that

$$\begin{cases} 0 = -\lambda - \eta c \tilde{U}'(y - \beta^* c) \\ \quad + (1 - \eta)(1 - c) \tilde{U}'(y + (1 - c)\beta^* - \ell) \\ 0 = \lambda \beta^* \\ 0 \geq \beta^* \\ 0 \geq \lambda \end{cases} \quad (20)$$

These conditions are the Karush-Khun-Tucker (KKT) necessary conditions. Combining the first and the last condition, we get

$$0 \geq -\eta c \tilde{U}'(y - \beta^* c) + (1 - \eta)(1 - c) \tilde{U}'(y + (1 - c)\beta^* - \ell) \quad (21)$$

We analyze the consumer's decision by considering two cases and we present the results in the following propositions.

Proposition 1: Suppose that the consumer is offered privacy insurance at the rate $c = 1 - \eta$, i.e. at a rate equal to the probability of a successful attack. Then the consumer will choose to purchase an amount of insurance equal to the loss, i.e. $\beta^* = \ell$.

Proof: Since $c = 1 - \eta$, (21) reduces to

$$0 \geq \eta(1 - \eta)(\tilde{U}'(y + \eta\beta^* - \ell) - \tilde{U}'(y - \beta^*(1 - \eta))) \quad (22)$$

and since $(1 - \eta)\eta \geq 0$ this again reduces to

$$0 \geq \tilde{U}'(y + \eta\beta^* - \ell) - \tilde{U}'(y - \beta^*(1 - \eta)) \quad (23)$$

Recall that we assumed \tilde{U} to be a concave function and that a function is strictly concave if and only if its derivative \tilde{U}' is decreasing. Hence,

$$\tilde{U}'(z) < \tilde{U}'(z - \ell) \quad (24)$$

since $\ell > 0$. This fact along with (23) implies that $\beta^* > 0$. Now, we claim that $\beta^* = \ell$. Indeed, suppose that $0 < \beta^* < \ell$, then from (23) we have

$$\tilde{U}'(\bar{y} - \ell) \leq \tilde{U}'(\bar{y} - \beta^*) \quad (25)$$

where $\bar{y} = y + \eta\beta^*$. This inequality violates (24). On the other hand, suppose that $0 \leq \ell \leq \beta^*$, then from (24) we have

$$\tilde{U}'(\bar{y} - \beta^*) > \tilde{U}'(\bar{y} - \ell) \quad (26)$$

but this violates the KKT inequality (23). Hence, $\beta^* = \ell$ which is to say that the consumer will purchase an amount of insurance equal to the loss of privacy she would endure under an attack. ■

Proposition 2: Suppose that the consumer is offered insurance at the rate $c > 1 - \eta$, i.e. at a rate higher than the probability of a successful attack. Then the consumer will not purchase the full insurance, i.e. $\beta^* < \ell$.

Proof: Suppose that the consumer is offered privacy insurance at a rate $c > 1 - \eta$ and that the optimal choice for the consumer is $\beta^* = \ell \geq 0$. Then, first-order optimality conditions imply that

$$-\eta \tilde{U}'(y - \ell)c + (1 - \eta) \tilde{U}'(y - \ell c)(1 - c) = 0 \quad (27)$$

However, since $c > 1 - \eta$ and \tilde{U} is increasing, from (21) we have

$$(-\eta c + (1 - \eta)(1 - c)) \tilde{U}'(y - \ell c) < 0 \quad (28)$$

so that, in fact, the optimal β has to be less than the loss experienced, i.e. $\beta^* < \ell$. ■

B. Analysis of the Insurer's Decision

Let us now consider the case of a third-party insurance company offering offering privacy insurance to the consumer which protects them against losses due to attacks. In a way, insurance allows the consumer to *hedge their bet* against selecting a contract with a low-privacy setting.

We consider a similar setup as before: the consumer's utility function \tilde{U} is strictly concave, increasing and twice differentiable and for the sake of analysis we assume that $\tilde{U}(0) = 0$. We consider a scenario in which the insurer faces two types: high-risk consumer θ_h and low-risk consumer θ_l . That is to say we are assuming that there is a portion of the population that is more likely to be attacked, i.e. the risky consumers, possibly because they engage in high-risk behavior or due to the fact that they selected a low-privacy setting contract with the utility company. The consumer again has an initial amount of private information with value y and with probability $1 - \eta_j$ some of her private information is exposed resulting in a loss ℓ where $j = h, l$ indicates the consumer's type. We assume that $1 - \eta_l < 1 - \eta_h$.

We will assume that the insurer has a prior over the distribution of types. In particular, we assume that the risky type θ_h occurs in the population with probability p and that $p > 0$.

Suppose we are given an insurance contract (α_a, α_n) where α_a is the compensation to the consumer given that a successful attack occurred and α_n is the neutral case (no attack). Let X be a random variable representing the consumer's wealth such that with probability $1 - \eta_i$ it takes value $y - \ell + \alpha_a$ and with probability η_i it takes value $y - \alpha_n$. Then, the consumers expected utility is

$$E[\tilde{U}(X)] = (1 - \eta_i) \tilde{U}(y - \ell + \alpha_a) + \eta_i \tilde{U}(y - \alpha_n) \quad (29)$$

Note that in the previous subsection we analyzed the consumer's decision given an insurance contract of the form

$$\alpha_a = (1 - c)\beta, \quad \alpha_n = \beta c \quad (30)$$

The insurer is a monopolist whose expected cost is

$$\begin{aligned} \Pi(\alpha_a^h, \alpha_n^h, \alpha_a^l, \alpha_n^l) = & p \left(-(1 - \eta_h) \alpha_a^h + \eta_h \alpha_n^h \right) \\ & + (1 - p) \left(-(1 - \eta_l) \alpha_a^l + \eta_l \alpha_n^l \right) \end{aligned} \quad (31)$$

In the case of asymmetric information, i.e. the insurer does not know the consumer's type, the optimization problem he

must solve is

$$\begin{aligned}
& \max_{\{\alpha_a^j, \alpha_n^j\}_{j=h,l}} \Pi(\alpha_a^h, \alpha_n^h, \alpha_a^l, \alpha_n^l) & (P-5) \\
& \text{s.t.} \quad (1 - \eta_i)\tilde{U}(y - \ell + \alpha_a^i) + \eta_i\tilde{U}(y - \alpha_n^i) \\
& \quad \geq (1 - \eta_i)\tilde{U}(y - \ell + \alpha_a^j) + \eta_i\tilde{U}(y - \alpha_n^j), \\
& \quad \quad i, j \in \{h, l\}, i \neq j & (IC) \\
& \quad (1 - \eta_i)\tilde{U}(y - \ell + \alpha_a^i) + \eta_i\tilde{U}(y - \alpha_n^i) \\
& \quad \geq (1 - \eta_i)\tilde{U}(y - \ell) + \eta_i\tilde{U}(y), \quad i \in \{h, l\} & (IR)
\end{aligned}$$

Constraints labeled (IC) are the incentive compatibility constraints and constraints (IR) are the individual rationality constraints. Both are similar to those presented in Section IV-A. Incentive compatibility ensures that the consumer will report their type truthfully and the individual rationality constraint ensures that the consumer will participate.

Following a similar reasoning as in Section IV-A, we can reduce the optimization problem (P-5) by reasoning about the constraint set defined by (IC) and (IR). In particular, we argued that the high-privacy type's incentive compatibility constraint was active and that the low-privacy type's individual rationality constraint was active. In addition, we showed the other two constraints (IC-2) and (IR-1) could be removed. Now, in the insurance case, since $1 - \eta_l < 1 - \eta_h$, the incentive compatibility constraint for the risk type is active and the individual rationality constraint for the safe type is active, i.e. the constraint set for (P-5) becomes

$$\begin{aligned}
& (1 - \eta_h)\tilde{U}(y - \ell + \alpha_a^h) + \eta_h\tilde{U}(y - \alpha_n^h) \\
& = (1 - \eta_h)\tilde{U}(y - \ell + \alpha_a^l) + \eta_h\tilde{U}(y - \alpha_n^l) & (IC-h) \\
& (1 - \eta_l)\tilde{U}(y - \ell + \alpha_a^l) + \eta_l\tilde{U}(y - \alpha_n^l) \\
& = (1 - \eta_l)\tilde{U}(y - \ell) + \eta_l\tilde{U}(y) & (IR-1)
\end{aligned}$$

Let us try to restate the problem in a way which allows us to characterize the solutions. Since we have assumed that \tilde{U} is strictly concave, increasing and twice differentiable, we can define W be its inverse, where $W' > 0$ and $W'' > 0$. Further, define

$$\tilde{U}_a^i = \tilde{U}(y - \ell + \alpha_a^i) \quad \text{and} \quad \tilde{U}_n^i = \tilde{U}(y - \alpha_n^i). \quad (32)$$

The transformed utility is

$$\begin{aligned}
\tilde{\Pi}(\tilde{U}_a^h, \tilde{U}_n^h, \tilde{U}_a^l, \tilde{U}_n^l) &= p(-\eta_h W(\tilde{U}_n^h) - (1 - \eta_h)W(\tilde{U}_a^h)) \\
& \quad + x - (1 - \eta_h)\ell + (1 - p)(-\eta_l W(\tilde{U}_n^l) \\
& \quad - (1 - \eta_l)W(\tilde{U}_a^l) + x - (1 - \eta_l)\ell) & (33)
\end{aligned}$$

Then problem (P-5) becomes

$$\begin{aligned}
& \max_{\{\tilde{U}_a^j, \tilde{U}_n^j\}_{j=h,l}} \tilde{\Pi}(\tilde{U}_a^h, \tilde{U}_n^h, \tilde{U}_a^l, \tilde{U}_n^l) & (P-6) \\
& \text{s.t.} \quad (1 - \eta_h)\tilde{U}_a^h + \eta_h\tilde{U}_n^h = (1 - \eta_h)\tilde{U}_a^l + \eta_h\tilde{U}_n^l \\
& \quad (1 - \eta_l)\tilde{U}_a^l + \eta_l\tilde{U}_n^l = (1 - \eta_l)\tilde{U}(y - \ell) + \eta_l\tilde{U}(y)
\end{aligned}$$

The Lagrangian of the optimization problem is

$$\begin{aligned}
L(\tilde{U}_a^h, \tilde{U}_n^h, \tilde{U}_a^l, \tilde{U}_n^l, \lambda_1, \lambda_2) &= \tilde{\Pi}(\tilde{U}_a^h, \tilde{U}_n^h, \tilde{U}_a^l, \tilde{U}_n^l) \\
& \quad + \lambda_1((1 - \eta_h)\tilde{U}_a^h + \eta_h\tilde{U}_n^h - (1 - \eta_h)\tilde{U}_a^l - \eta_h\tilde{U}_n^l) \\
& \quad + \lambda_2((1 - \eta_l)\tilde{U}_a^l + \eta_l\tilde{U}_n^l - (1 - \eta_l)\tilde{U}(y - \ell)). & (34)
\end{aligned}$$

Proposition 3: Given the probabilities $1 - \eta_j$, $j = h, l$ that the consumer of type j will experience a privacy breach, if the insurer solves the optimization problem (P-6), then the high-risk consumer will be fully insured and the low-risk consumer will not be fully insured.

Proof: We first show that the risky type will be fully insured. Taking the derivative of the Lagrangian with respect to \tilde{U}_a^h and \tilde{U}_n^h we get the following two equations:

$$0 = -p(1 - \eta_h)W'(\tilde{U}_a^h) + \lambda_1(1 - \eta_h) \quad (35)$$

$$0 = -p\eta_h W'(\tilde{U}_n^h) + \lambda_1\eta_h \quad (36)$$

Solving for λ_1 in the first equation and plugging it into the second, we get $\tilde{U}_a^h = \tilde{U}_n^h$ so that $\ell - \alpha_a^h = \alpha_n^h$, i.e. the amount the high-risk type pays for insurance is equal to the compensation minus the loss in the event of a privacy breach. Thus, the high-risk type will be fully insured.

Now, we show that the low-risk type will not be fully insured. Taking the derivative of the Lagrangian with respect to \tilde{U}_a^l and \tilde{U}_n^l , we get

$$0 = -(1 - \eta_l)(1 - p)W'(\tilde{U}_a^l) - \lambda_1(1 - \eta_h) + \lambda_2(1 - \eta_l) \quad (37)$$

$$0 = -(1 - p)\eta_l W'(\tilde{U}_n^l) - \lambda_1\eta_h + \lambda_2\eta_l \quad (38)$$

From (35), we solved for $\lambda_1 = pW'(\tilde{U}_a^h)$. By plugging in λ_1 into (37), solving for λ_2 and plugging both λ_1 and λ_2 into (38), we get the following expression:

$$\begin{aligned}
0 &= W'(\tilde{U}_a^h)p \left(-\eta_h + \eta_l \frac{1 - \eta_h}{1 - \eta_l} \right) \\
& \quad + \eta_l(1 - p)(W'(\tilde{U}_a^l) - W'(\tilde{U}_n^l)) & (39)
\end{aligned}$$

Since $\eta_l > \eta_h$ and W' is increasing by assumption, the above equation implies that

$$\tilde{U}_n^l - \tilde{U}_a^l > 0 \quad (40)$$

and hence the low-risk type does not fully insure. ■

The above proposition tells us that in order to keep the high-risk type from masking as a low-risk type, the insurer must make the contract for the low-risk type unappealing to the high-risk type.

We remark that the analysis in this section can be applied to the case where the utility company is purchasing insurance as well. In particular, if the utility company has not invested in a lot of security or tjeu are not following the best practices recommendations, e.g. NIST-IR 7628 [25], then they are engaging in *risky* behavior. The insurance company will not know a propri whether or not the utility company is high-risk type. Through the design of insurance contracts the insurance company can asses the utility's type while offering contracts that maximize their own utility.

VI. CONCLUSION

Utilizing our results on the fundamental limits of non-intrusive load monitoring, we provide a novel upper bound on the probability of a successful privacy breach. Under the privacy metering policy in which sampling rate variation is used, we study how the performance of direct load control degrades using the \mathcal{H}_∞ norm. This provided us with a metric for understanding how sampling rate affects the quality of direct load control. Using this metric along with the upper bound on the probability for a successful privacy breach, we design a screening mechanism for the problem of obtaining the consumer's type when there is asymmetric information. Further, we design insurance contracts using the probability of successful privacy breach given that in the population of consumers there is both high-risk and low-risk consumers.

This work opens up a number of questions in the area of privacy metrics as well as customer segmentation and targeting. We considered only two-type models in both the design of contracts. We are currently looking at the theory for a continuum of types. The screening problem with a continuum of types results in a problem that resembles a partial differential equation constrained optimal control problem. We are developing numerical techniques to solve this problem. We also assumed that the utility company and private insurer knew the distribution of types in the population. We are currently developing algorithms for learning these probabilities using data-drive techniques. In addition, we considered that the utility would offer a contract solely based on privacy settings whereas in reality the contract would normally contain additional items such as maximum power consumption, rate, etc. Consumers in the population may value these goods differently. In this setting, the screening problem would be come multi-dimensional [26]. We are exploring this in the context of privacy-aware incentive design for behavior modification.

REFERENCES

- [1] E. L. Quinn, "Smart metering and privacy: Existing laws and competing policies," Colorado Public Utilities Commission, Tech. Rep., 2009.
- [2] M. Salehie, L. Pasquale, I. Omoronyia, and B. Nuseibeh, "Adaptive security and privacy in smart grids: A software engineering vision," in *International Workshop on Software Engineering for the Smart Grid*, June 2012, pp. 46–49.
- [3] R. Dong, L. Ratliff, H. Ohlsson, and S. S. Sastry, "Fundamental limits of nonintrusive load monitoring," *Proceedings of the 3rd ACM International Conference on High Confidence Networked Systems*, 2013.
- [4] G. Hart, "Nonintrusive appliance load monitoring," *Proceedings of the IEEE*, vol. 80, no. 12, pp. 1870–1891, 1992.
- [5] J. Z. Kolter and M. J. Johnson, "Redd: A public data set for energy disaggregation research," in *Workshop on Data Mining Applications in Sustainability (SIGKDD)*, San Diego, CA, 2011.
- [6] K. Anderson, A. Ocneanu, D. Benitez, D. Carlson, A. Rowe, and M. Berges, "BLUED: a fully labeled public dataset for Event-Based Non-Intrusive load monitoring research," in *Proceedings of the 2nd KDD Workshop on Data Mining Applications in Sustainability*, Beijing, China, Aug. 2012.
- [7] J. Z. Kolter and T. Jaakkola, "Approximate inference in additive factorial HMMs with application to energy disaggregation," in *Proceedings of the International Conference on Artificial Intelligence and Statistics*, 2012, pp. 1472–1482.

- [8] M. J. Johnson and A. S. Willsky, "Bayesian nonparametric hidden semi-Markov models," *The Journal of Machine Learning Research*, vol. 14, no. 1, pp. 673–701, 2013.
- [9] O. Parson, S. Ghosh, M. Weal, and A. Rogers, "Nonintrusive load monitoring using prior models of general appliance types," in *26th AAAI Conference on Artificial Intelligence*, 2012.
- [10] J. Z. Kolter, S. Batra, and A. Y. Ng, "Energy disaggregation via discriminative sparse coding," in *Neural Information Processing Systems*, 2010.
- [11] R. Dong, L. J. Ratliff, H. Ohlsson, and S. Sastry, "Energy disaggregation via adaptive filtering," in *Proceedings of the 51st Annual Allerton Conference on Communication, Control, and Computing*, Oct 2013, pp. 173–180.
- [12] C. Dwork, "Differential privacy," in *Proceedings of the International Colloquium on Automata, Languages and Programming*. Springer, 2006, pp. 1–12.
- [13] J. Le Ny and G. Pappas, "Differentially private filtering," *IEEE Transactions on Automatic Control*, vol. 59, no. 2, pp. 341–354, 2014.
- [14] R. Shokri, G. Theodorakopoulos, J. Le Boudec, and J. Hubaux, "Quantifying location privacy," in *Security and Privacy (SP), 2011 IEEE Symposium on*. IEEE, 2011, pp. 247–262.
- [15] A. A. Cárdenas, S. Amin, G. Schwartz, R. Dong, and S. S. Sastry, "A game theory model for electricity theft detection and privacy-aware control in AMI systems," in *Proceedings of the 50th Allerton Conference on Communication, Control, and Computing*, 2012, pp. 1830–1837.
- [16] R. Dong, A. A. Cárdenas, H. Ohlsson, and S. S. Sastry, "Privacy-aware sampling policies for advanced metering infrastructures," *ArXiv*, 2014.
- [17] D. Callaway and I. Hiskens, "Achieving controllability of electric loads," *Proceedings of the IEEE*, vol. 99, no. 1, pp. 184–199, Jan 2011.
- [18] J. Mathieu, S. Koch, and D. Callaway, "State estimation and control of electric loads to manage real-time energy imbalance," *IEEE Transactions on Power Systems*, vol. 28, no. 1, pp. 430–440, 2013.
- [19] P. Seiler and R. Sengupta, "An H-infinity approach to networked control," *IEEE Transactions on Automatic Control*, vol. 50, no. 3, pp. 356–364, March 2005.
- [20] H. Ishii, "H-infinity control with limited communication and message losses," *Systems & Control Letters*, vol. 57, no. 4, pp. 322 – 331, 2008.
- [21] T. A. Weber, "Optimal control theory with applications in economics," *MIT Press Books*, vol. 1, 2011.
- [22] M. Rothschild and J. Stiglitz, "Equilibrium in competitive insurance markets: An essay on the economics of imperfect information," *The Quarterly Journal of Economics*, vol. 90, no. 4, pp. 629–645, 1976.
- [23] G. D. Jaynes, "Equilibria in monopolistically competitive insurance markets," *Journal of Economic Theory*, vol. 19, no. 2, pp. 394 – 422, 1978.
- [24] M. Mussa and S. Rosen, "Monopoly and product quality," *Journal of Economic Theory*, vol. 18, no. 2, pp. 301 – 317, 1978.
- [25] Smart Grid Interoperability Panel Cyber Security Working Group and others, "Introduction to NISTIR 7628 guidelines for smart grid cyber security," Tech. Rep., 2010.
- [26] S. Basov, "Multidimensional screening," in *Studies in Economic Theory*. Springer, 2005, vol. 22.