Allocative and Strategic Effects of Privacy Enhancement in Smart Grids

S. Kessler[a,*], C. M. Flath[b,1], K. Böhm[a]

[a] *Karlsruhe Institute of Technology (KIT), Kaiserstraße 12, 76313 Karlsruhe, Germany*
[b] *Julius-Maximilians-Universität Würzburg, Josef-Stangl-Platz 2, 97070 Würzburg, Germany*

Abstract

Local energy markets are a promising approach for automatic and efficient matching of renewable energy with household demand in smart grids. Therefore, such markets can help to improve power system reliability and at the same reduce emissions. However, to participate in such markets, customers need to disclose private consumption data. A number of studies show that such data records may reveal a broad range of personal, sensitive information on the inhabitants. Privacy-enhancement mechanisms can be applied to preserve the privacy of individuals modify the data reported to the market. Yet, these mechanisms can lower allocative efficiency and alter theoretical properties of the market mechanism.

In this paper, we characterize both theoretically and numerically the effect of privacy mechanisms applied in a local energy market scenario. Our model considers demand side flexibility as well as energy storage systems. Furthermore, we allow for a free specification of the desired privacy enhancement level. We show that under certain natural assumptions market mechanisms retains in-expectation incentive compatibility despite the presence of privacy enhancement. Our numerical analysis based on real-world data shows that the welfare impact of privacy enhancement mechanisms is limited. Furthermore, energy storage can mitigate this efficiency loss to a large extent.

Keywords: privacy enhancement, local energy markets, market engineering, allocative effects, strategic effects

1. Introduction

Historically, the electricity grid is tailored to a centralized generation structure. At its core, there are few large power plants generating electricity for a large number of consumers [1]. However, reducing the CO_2 emissions of the

*Corresponding author
Email addresses: stephan.kessler@kit.edu (S. Kessler), christoph.flath@uni-wuerzburg.de (C. M. Flath), klemens.boehm@kit.edu (K. Böhm)
[1] Work originated while this author was at KIT.

energy production requires the integration of renewable sources, such as photovoltaic sites and micro combined heat and power plants. These sources are volatile and distributed. Compared to a large power plant, each of them produces only a small amount of power and cannot be controlled centrally. Due to their variability, integration of renewables remains a big challenge on today's power system.

Smart grids [2], the ICT-enabled electricity networks of the future, facilitate new operational paradigms [3, 4]. A case in point is the establishment of local energy markets. These are a means fpr matching regional energy demand and renewable supply [5, 6]. More 'local' (i.e. in spatial proximity of generation) energy consumption can help to improve integration of renewables and minimize transmission losses [7]. In order to work efficiently, local energy markets rely on truthful power-consumption information revealed by the participants, e.g., private households. In such markets, customers cover their energy needs by bidding for the required energy amounts over short time intervals. Consequently, a customer's consumption behavior is encapsulated in these bids. Yet a number of studies show that such fine-grained power consumption data can reveal significant amounts of private information [8, 9], including wealth, daily routines and household size. Consequently, electronic market systems in smart grids should strive to preserve privacy properties [10]. Privacy enhancement methods distort sensitive values, e.g., energy consumption levels. By doing so, they are able to retain a certain level of privacy despite personal data being revealed to the market.

However, distorted bids are likely to induce less efficient allocations. Depending on the nature of the distortion, more or less energy than actually needed may be allocated. Hence, privacy enhancement may lead to additional costs for consumers. In this paper, we quantify these privacy costs in a local energy market with demand side flexibility and storage. Due to the large number of possible influence factors (e.g., customer privacy preferences, demand side flexibility as well as supply and demand patterns), it is difficult to fully characterize this welfare loss in a general fashion. For instance, realistic supply and demand patterns are complex random processes, and the privacy enhancement methods applied add complexity as well. A general model covering all these details will lack expressiveness.

To understand the relationship of privacy enhancement and local energy markets, we model a smart grid marketplace with privacy enhancement methods together with a customer demand model. We characterize customer-bidding behavior and determine formal characteristics of the interplay between components of our model. This includes the definition of general properties of privacy enhancement methods. Furthermore, we show that a privacy-aware auction retains incentive compatibility with respect to valuations if the privacy enhancement method is monotonic and marginal utility is independent of the demand level.

Subsequently, we instantiate a numerical evaluation using empirical load and generation data. By means of simulations we quantify the costs of privacy enhancement. Specifically, we assess the economic effect of varying numbers and types of generators, demand properties and storage endowments. The experiments illustrate the relationship between distortion level and welfare loss incurred. Further, we can quantify the positive effect of storage in the presence of privacy enhancement methods. Small scale electricity storage can reduce privacy-induced welfare loss by almost two-thirds. In summary, this paper explores the economic effects of different privacy levels. As such, it offers a concrete application scenario to evaluate the actual impact of recent privacy enhancement methods.

The remainder of this paper is structured as follows. In Section 2 we provide an overview of related research on smart grids and privacy enhancement techniques. The market model is described in Section 3. Subsequently, we present theoretical results derived from the model in Section 4. The following Sections 5 and Section 6 introduce the parametrization and evaluation of a simulation-based instantiation of the model. Section 7 concludes.

2. Related Work

2.1. Privacy in Smart Grids

Renewable sources for electricity generation are distributed and volatile by nature. The efficient utilization of such sources is an important part of the smart grid vision. Local energy markets efficiently coordinate decentralized generation of electricity [5]. Generators of renewable energy as well as consumers participate in such local markets and trade energy over short time intervals, e.g., 30 minutes or less. Transparency obligations like the EUC 543/2013 mandate the publication of comprehensive market data. Market transparency is key to ensure market liquidity and hence market efficiency [11, 12].

Yet, fine-grained power consumption records contain sensitive personal information [9]. For example, appliances have a characteristic power consumption pattern over time called load signature [8]. This facilitates the detection of appliances [13] or even the currently selected TV channel [14]. Removing personal identifiers like name or address is not sufficient, since the data itself is identifying [15]. Consequently, power consumption data is subject to privacy legislation, e.g., European Directive 95/46/EC.

To mediate between the diametric goals of market transparency and customer privacy-protection, local energy markets need to incorporate appropriate privacy-enhancement techniques. To this end, Buchmann et al. [16] investigate the impact of privacy enhancing methods on the expenses of individuals and as a measure for the impact on data utility. While their work is related to our research, their approach is limited to simple strategies, bidding exactly one limit price which does not take different valuations of energy into account. In addition, they do not account for demand-side energy storage and do not provide any formal characterization of the relationship between privacy enhancement and incentive compatibility.

In contrast, Kalogridis et al. [17] as well as Varodayan and Khist [18] both rely on the use of energy storage for privacy protection. In both cases a stationary storage is used to completely mask the load signatures of the underlying household appliances. However, these results are mainly anecdotal and rely on an arbitrarily large storage system. We follow the general idea by investigating the economic interplay between privacy enhancement and a fixed storage system with limited capacity. This allows us to compare the previously orthogonal dimensions of storage costs and privacy.

2.2. Privacy Enhancement Techniques

There is a broad variety of approaches to protect the privacy of individuals in a data set, see [19] for a survey. In general, one can distinguish between approaches facilitating individual privacy preferences and approaches with a common privacy parameter for the entire data set. In the local energy market scenario, individual preferences let each household decide the degree of distortion and privacy independently of others.

The well-known k-Anonymity principle [20] is an instance of the latter case, jointly modifying the data of groups of individuals as a whole. An application to load data requires adaption of the principle to handle time series [21, 22]. As an example for distortion, this means that the anonymization algorithm takes k time series and replaces the values by averages over all k values at each point of time. Since a central instance governs the group size for the entire data set, individuals cannot specify their individual privacy preferences. Yet, in a smart grid scenario we need to acknowledge decentrality as well as heterogeneous privacy requirements. These are ultimately governed by customer context [23]. Differential Privacy [24] provides provable privacy guarantees in the presence of unlimited background knowledge. However, instantiations of Differential Privacy on time series [25] provide these guarantees only on aggregates for sets of such series. A local marketplace cannot operate with such aggregated values as it requires each participant to place individual bids.

Cryptographic auctions encrypt the bids and provide verifiability of the correctness. However, they do not allow ex-post access of the information, which limits market transparency. Furthermore, they do not facilitate repeated and parallel market interactions as they are designed for a single seller [26] or preserve secrets only until the end of an auction [27].

Perturbation approaches add random noise to each time series in isolation. Thus, the algorithms handle each time series separately and are able to take individual preferences into account. While not offering formal privacy guarantees, evaluations [28, 29] show that privacy is preserved reasonably in many situations. In this paper, we apply the privacy enhancement method introduced by Papadimitriou et al. [28]. Furthermore, we present a modification that retains incentive compatibility on the market. Both methods allow individually defined privacy preferences. However, the ideas presented are not limited to a certain approach as long as they fulfill certain requirements as discussed in Section 3.3.

Allocation efficiency on a local energy market offers an application-specific measure for comparing different privacy enhancement methods. This is in contrast to general abstract measures that do not consider an actual real-world scenario [30].

3. Model

In this section, we specify our theoretical local energy market model as well as corresponding privacy enhancement methods.

3.1. Market Structure

A local energy market allocates electricity generated locally (distributed generation, imports) to local consumers. We want to study the effects of applying privacy enhancement techniques on such a local energy market with the help of a model. Following the market engineering paradigm [31, 6], an appropriate model of a marketplace requires specifying the participants and their behavior (agent behavior, Section 3.1.1), the transaction object (Section 3.1.2), the market mechanism (market microstructure, Section 3.1.4) as well as market performance measures (market outcome, Section 3.1.5). Given our focus on the cost of privacy, we also describe the privacy enhancement methods that are part of the bidding process. (Section 3.1.3).

3.1.1. Market Participants

$\mathcal{A}$ is the set of participants in our local energy market. Each participant (actor) $a \in \mathcal{A}$ is either a consumer $c \in \mathcal{C}$ or a producer (generator) $g \in \mathcal{G}$, i.e., we assume $\mathcal{C} \cap \mathcal{G} = \emptyset$.[2] Consumer energy demand is varying over time. We model the time domain as a sequence of time intervals $\mathcal{T}$. For each $t \in \mathcal{T}$, a consumer's maximum consumption level, referred to as the saturation level, is given by $\overline{x}_c(t)$. The trajectories of saturation levels form a set of time series: $\mathcal{S} = \{\overline{x}_c(t) | t \in \mathcal{T}, c \in \mathcal{C}\}$. The purchasing behavior of a consumer is governed by individual utility as specified in Section 3.2. Given temporally varying electricity needs $\overline{x}_c(t)$, optimal bidding requires customers to dynamically determine quantity-utility mappings.

The set of producers consists of local generation units (PV, CHP) and a balancing party. This party reflects energy imports from the superordinate grid. Producers participate in the market by selling electricity. Individual rationality requires them to at least cover their marginal generation costs. Their capacity limits their bid quantities.

3.1.2. Transaction Object

Our market instantiation follows traditional wholesale electricity markets in that electrical energy supply and demand commitments are traded. A bid contains the issuing market participant and its type (buy or sell order), the amount of energy procured (in kWh) as well as the reservation price p. In the case of a sell order the latter specifies the minimum price, in case of buy order it is the maximum price. Individual actors can submit several bids to reflect non-linear customer utility and generator cost functions.

[2]We do not consider prosumers (producer and consumer) as they give rise to new strategic considerations by acting on both sides of the market.

3.1.3. Privacy-enhanced Bidding

As noted above, a customer c will formulate her bid at time t to reflect her current energy demand saturation level $\overline{x}_c(t)$. Consumers will place a collection of bids reflecting their utility function under $\overline{x}_c(t)$. In the presence of privacy enhancement, the bidding process slightly changes: Consumers report $\overline{x}_c(t)$ to the privacy protection system which in turn determines a modified demand report $\widetilde{x}_c(t)$ which is communicated to the market.

Note that there are two distinct elements in the report that a consumer could strategize about, quantity and price. As private information is embedded only in the quantity component, we rule out that consumers will modify their demand report as this could open a side channel undermining the privacy enhancement technique. Consequently, we assume that the privacy protection system receives the true initial demand reports of the consumers. For the valuation, we do not make this assumption but show that in specific cases privacy-aware auctions are incentive compatible with respect to prices and consumers thus will optimally report their true valuation to the system. See Lemma 5 for further discussion. Bidding with and without a privacy enhancement method is illustrated in Figure 1.

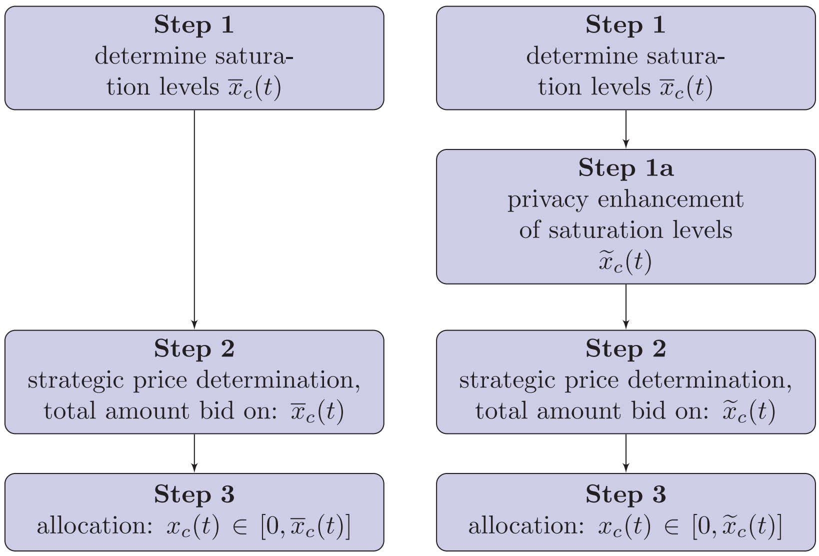

Figure 1: Auctions *without (left)* and *with (right)* privacy enhancement

3.1.4. Market Mechanism

Since electricity is a homogeneous good, double-sided auction formats can achieve a high level of market liquidity and efficiency [32, 33]. Following previous research on local energy markets, e.g., [6, 16], we select the discrete time double auction (also referred to as periodic call auction or call market) as the market mechanism. Here, market clearing is not continuous but occurs in repeated time slots $t \in \mathcal{T}$. For each time slot, the market mechanism determines the allocation and clearing price for the submitted bids and asks. It does so by first constructing demand and supply curves and subsequently determining the intersection of the two. For a more detailed treatise of the call market we refer to Parsons et al. [34].

3.1.5. Market Quality Measure

In order to assess the economic outcome of our local marketplace, we analyze the market's allocative efficiency. Consequently, we use social welfare as an application-specific measure for the effect of privacy-enhancement mechanisms.

Definition 1 (Social Welfare): Social welfare is the sum of consumer surplus (difference between willingness-to-pay and clearing price) and producer surplus (difference between clearing price and costs):

$$\mathcal{W} = \sum_{\forall c \in \mathcal{C}} CS_c + \sum_{\forall g \in \mathcal{G}} PS_g$$

To ease comparing different simulation runs we rely on the relative Welfare.

Definition 2 (Relative Welfare): Let $\mathcal{W}$ be the welfare achieved in a local energy market without privacy enhancement. Further let $\widetilde{\mathcal{W}}$ be the welfare achieved in the same market (concerning supply and demand), but in the presence of a privacy enhancement method modifying the bids of consumers. Relative welfare $\mathcal{W}'$ is then given by

$$\mathcal{W}' = \frac{\mathcal{W}}{\widetilde{\mathcal{W}}}$$

We posit that higher relative welfare is an indication that a privacy-enhancement method retains a higher data quality in the application scenario under consideration.

3.2. Customer Model

The key element to model customer interactions (i.e., bidding behavior) with the market is the underlying utility model. While electricity traditionally is subject to billing and is considered a homogeneous good, the Smart Grid includes differentiated energy services [35], and we follow this notion. To this end, we propose an analytical customer model similar to the one presented by Bitar and Low [36].

3.2.1. Customer Utility

Denoting customer c's electricity allocation for time slot t as $x_c(t)$, utility $U_{c,t}(.)$: $\mathbb{R}_+ \mapsto \mathbb{R}_+$ is assumed to be a non-decreasing, concave function. Furthermore, we assume an demand saturation level $\overline{x}_c(t)$ beyond which customers no longer obtain any utility from additional electricity consumption. In other words

$$U_{c,t}(x_c(t)) = U_{c,t}(\overline{x}_c(t)) = \overline{U}_{c,t} \forall x_c(t) \geq \overline{x}_c(t). \tag{1}$$

If the saturation level does (not) affect marginal utility, we refer to the utility function as saturation-level-dependent (independent).

This customer model reflects the smart grid rationale of customers adapting consumption to current system conditions. Following standard economic theory, marginal utility of additional consumption is assumed to decrease – most valuable usage forms are activated first. At some point the customer will not be able to put additional energy allocations to any meaningful use. In our analysis, we make the following non-restrictive assumptions: Market quantities are discretized with granularity D. The admissible $x_c(t)$-values are discretized as well: $x_c(t) \in \{n \cdot D | n \in \mathbb{N}\}$. Similarly, we discretize the $\overline{x}_c(t)$ values. To model the temporal pattern of the energy-usage behavior of customers, the saturation levels $\overline{x}_c(t)$ fluctuate over time in tune with a representative energy-demand profile. When considering families of utility functions, we interpret the concavity of each function as differing levels of load flexibility: In case of a linear utility function, marginal utility from consumption is constant and hence load shedding has constant cost. Conversely, for a very concave function shedding utility losses at high load levels are limited. In the following, f denotes the demand flexibility. Higher f values indicate more flexible demand.

3.2.2. Bidding Behavior

Since $U_{c,t}$ provides a mapping from allocation to utility space, we can express a customer's optimal bidding behavior under this utility function using the marginal utility $U'_{c,t}$: Instead of placing a single price-quantity bid, a rational customer will rather place a continuum of bids with infinitesimal quantity and decreasing bid price to match her marginal utility function. Under our market discretization scheme, customers will place $n = \frac{\overline{x}_c(t)}{D}$ bids with quantity D each. The corresponding optimal bid prices then are $U^{'}_{c,t}(i \cdot D)$ with $i = 1...n$.

We distinguish between utility functions that are dependent or independent of the saturation level $\overline{x}_c(t)$. In the first case, the saturation level affects a customer's (marginal) utility value over the complete range of allocation quantities. In contrast, for saturation-level independency the (marginal) utility is independent of the saturation level over the interval $[0, \overline{x}_c(t)]$. Figure 2 illustrates the utility and marginal utility function for both cases In practice, the utility of a household is a combination of both. The analysis of the polar cases allows us to better structure our results.

3.3. Privacy Enhancement Methods

A privacy enhancement method $\mathcal{M}$ takes a set of time series $\mathcal{S}$ and parameters p and returns a privacy enhanced representation, i.e., $\widetilde{\mathcal{S}} = \mathcal{M}_p(\mathcal{S})$. In our smart grid scenario, the time series are given by consumers' saturation levels: $\mathcal{S} = \{\overline{x}_c(t) | t \in \mathcal{T}, c \in \mathcal{C}\}$. The parameter p is a method specific parameter that determines the level of privacy achieved. The privacy enhancement method

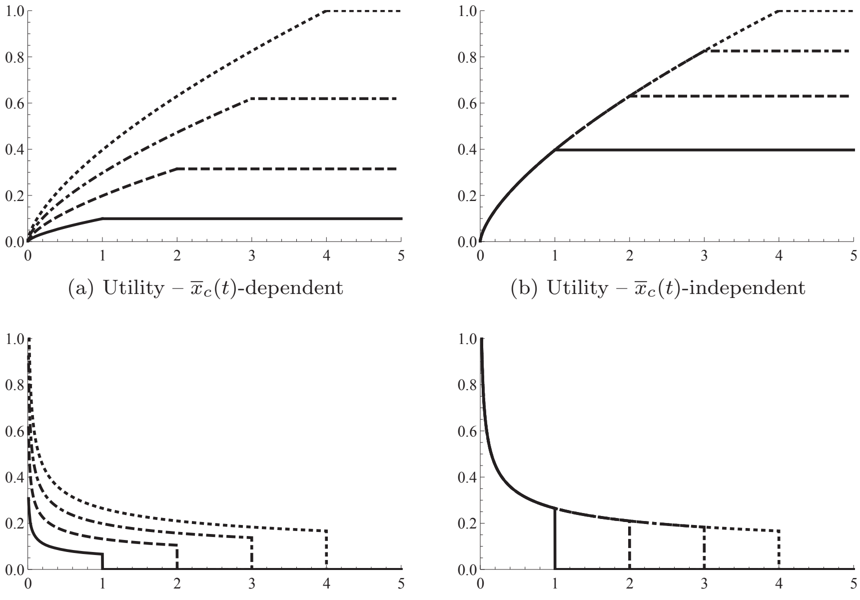

(a) Utility – $\overline{x}_c(t)$-dependent

(b) Utility – $\overline{x}_c(t)$-independent

(c) Marginal utility – $\overline{x}_c(t)$-dependent

(d) Marginal utility – $\overline{x}_c(t)$-independent

Figure 2: Illustration of saturation level dependency

modifies the values of the time series, in our case the saturation level values: $\widetilde{\mathcal{S}} = \mathcal{M}(\mathcal{S}) = \{\widetilde{x}_c(t) | t \in \mathcal{T}, c \in \mathcal{C}\}$. There are various methods with different approaches for preserving privacy in a set of time series. To derive theoretical results independent of an actual method, we need to come up with general properties of privacy enhancing methods.

A central distinction is the one between deterministic and randomized privacy enhancement methods.

Definition 3 (Deterministic): A privacy enhancement method is deterministic if the results of several runs are the same with the same input: $\widetilde{\mathcal{S}}_1 = \mathcal{M}_p(\mathcal{S}) \wedge \widetilde{\mathcal{S}}_2 = \mathcal{M}_p(\mathcal{S}) \Rightarrow \widetilde{\mathcal{S}}_1 = \widetilde{\mathcal{S}}_2$.

Definition 4 (Randomized): The privacy enhancement method depends on random calculations, and this may lead to different results if the method is run several times. The probability that the method returns a certain privacy-enhanced set $\widetilde{\mathcal{S}}$, $P(\mathcal{M}_p(\mathcal{S}) = \widetilde{\mathcal{S}})$, is the same for each run.

The rationale behind the following notions is to further characterize the effect of different privacy enhancement methods.

Definition 5 (Balanced Modifier): Let $\widetilde{\mathcal{S}} = \mathcal{M}(\mathcal{S})$, if $\mathcal{M}$ is a balanced modifier the following holds for all consumers $c \in \mathcal{C}$:

$$\sum_{\forall t \in \mathcal{T}} \widetilde{x}_c(t) - \overline{x}_c(t) = 0$$

A randomized privacy enhancement method is a balanced modifier if the expected value of the sum of these differences equals zero.

Definition 6 ($\cup$-homomorphism): Assume that $\mathcal{S}_1$ and $\mathcal{S}_2$ is an arbitrary partitioning of the time series in $\mathcal{S}$:

$$\mathcal{S} = \mathcal{S}_1 \cup \mathcal{S}_2 \wedge \mathcal{S}_1 \cap \mathcal{S}_2 = \emptyset$$

A deterministic privacy enhancement method is a homomorphism of $\cup$ if the following holds:

$$\mathcal{M}(\mathcal{S}_1) \cup \mathcal{M}(\mathcal{S}_2) = \mathcal{M}(\mathcal{S})$$

A randomized privacy enhancement method is a homomorphism of $\cup$ if the following holds:

$$P(\mathcal{M}(\mathcal{S}_1) \cup \mathcal{M}(\mathcal{S}_2) = \widetilde{\mathcal{S}}) = P(\mathcal{M}(\mathcal{S}) = \widetilde{\mathcal{S}})$$

In the following, we show that a privacy enhancement method, which is a homomorphism of $\cup$, will modify the time series independently of each other.

Lemma 1: *Let $\widetilde{\mathcal{S}} = \mathcal{M}(\mathcal{S})$. If $\mathcal{M}$ is a $\cup$-homomorphism, the modifications of time series $\widetilde{x}_c(t) \in \widetilde{\mathcal{S}}$ are independent of the time series of any other consumer $c' \neq c : \overline{x}_{c'}(t)$*

Proof: Let $\mathcal{S}_1 = \{\overline{x}_c(t)\}$ and $\mathcal{S}_2 = \mathcal{S} \backslash \mathcal{S}_1$. By definition $\mathcal{S}_1$ and $\mathcal{S}_2$ are partitions of $\mathcal{S}$. Since $\mathcal{M}$ is a $\cup$-homomorphism $\mathcal{M}(\mathcal{S}_1) \cup \mathcal{M}(\mathcal{S}_2)$ equals $\mathcal{M}(\mathcal{S})$. In particular, the resulting $\widetilde{x}_c(t)$ is independent of possible other time series in $\mathcal{S}$. $\square$

For instance, k-anonymity [37] usually does not have this property: The output of most implementations depends on the groups created. In turn, adding symmetric random noise is a $\cup$-homomorphism.

The privacy parameters p influence the privacy enhancement. In the following we define an order.

Definition 7 (Order of privacy parameters ($p_1 > p_2$): Let p_1 and p_2 be different parameters for privacy method $\mathcal{M}_p$. p_1 is greater than p_2 if $\mathcal{M}_{p_1}(\mathcal{S})$ provides a better privacy protection than $\mathcal{M}_{p_2}(\mathcal{S})$ in terms of the definition of $\mathcal{M}_p$.

Commonly known distance metrics, e.g., the L1-Norm, quantify the distance between two time series. Choosing a greater privacy parameter may lead to a larger distance if the privacy method is monotonically increasing, as defined in the following. Let $dist(\overline{x}_c(t), \widetilde{x}_c(t))$ be such a distance metric.

Definition 8 (Monotonically increasing): Let p_1, p_2 be two privacy parameter choices for $\mathcal{M}_p$ with p_1 having greater order than p_2, that is $p_1 > p_2$.

Furthermore, $\widetilde{x}_c^1(t) \in \mathcal{M}_{p_1}(\mathcal{S})$ and $\widetilde{x}_c^2(t) \in \mathcal{M}_{p_2}(\mathcal{S})$ are time series obtained by applying $\mathcal{M}_p$ on the same time series $\overline{x}_c(t) \in \mathcal{S}$.

A *deterministic* privacy enhancement method is monotonically increasing with respect to a metric $dist(\cdot)$, if the following holds for:

$$dist(\overline{x}_c(t), \widetilde{x}_c^1(t)) \geq dist(\overline{x}_c(t), \widetilde{x}_c^2(t)).$$

A *random* privacy enhancement method is monotonically increasing with respect to a metric $dist(\cdot)$ if in expectation the following holds:

$$\mathbb{E}\left[dist(\overline{x}_c(t), \widetilde{x}_c^1(t)) - dist(\overline{x}_c(t), \widetilde{x}_c^2(t)\right] > 0.$$

Intuitively, a privacy enhancement method is monotonically increasing if greater privacy parameter choices give rise to greater changes to the original time series values.

4. Theoretical Results

We now derive formal results on the impact of privacy enhancement on local energy markets. In the following we assume that the time series are of infinite length. We also assume non-triviality of the privacy enhancement methods, i.e., we exclude the case that $\widetilde{x}_c(t) = \overline{x}_c(t), \forall t \in \mathcal{T}$.

Lemma 2: *The welfare loss is monotonically increasing for greater privacy parameters if the privacy enhancement method is monotonically increasing and a balanced modifier.*

Proof: Assume $\widetilde{\mathcal{S}} = \mathcal{M}_p(\mathcal{S})$. Further, let d be the difference between the saturation level and the privacy enhanced saturation level of consumer c on time slot t: $d = \overline{x}_c(t) - \widetilde{x}_c(t)$. If $d > 0$ the higher d the lower the $\widetilde{x}_c(t)$ and potentially the higher the welfare loss. c may not get electricity allocated even if the marginal utility is greater than zero because there are no bids exceeding $\widetilde{x}_c(t)$. Similar result holds for $d < 0$: the smaller d the higher the potential welfare loss, because c may allocate energy at a price greater than zero, even if the marginal utility is zero. Depending on the actual utility functions, the welfare loss is higher for $d > 0$ or $d < 0$. However, if the privacy enhancement method is a balanced modifier, the sum of the welfare loss of all time slots remains the same. Let the privacy enhancement method $\mathcal{M}_1(p)$ be monotonically increasing, then the welfare loss for more restrictive privacy requirements is higher, since the distance between $\overline{x}_c(t)$ and $\widetilde{x}_c(t)$ also increases for $p_1 > p$. $\square$

Lemma 3: *In the presence of storage the welfare loss is equal or smaller than without storage if the privacy enhancement method is a balanced modifier.*

Proof: Let $\mathcal{M}_p$ be a balanced modifier and $\widetilde{\mathcal{S}} = \mathcal{M}_p(\mathcal{S})$. Assume that there exists a $t_1 \in \mathcal{T}$ where $\widetilde{x}_c(t_1) > \overline{x}_c(t_1)$. Since $\mathcal{M}_p$ is a balanced modifier, we assume that there exists a t_2 where $\widetilde{x}_c(t_2) < \overline{x}_c(t_2)$. If the allocation at t_1 is also greater than $\overline{x}_c(t_1)$ the additional electricity bought is stored and used in times of undersupply, or at t_2. If the storage did not exist, the additional electricity bought at t_1 would not return in utility. Only in the case that after t_1 there is no time-slot with undersupply, storage cannot reduce the welfare loss. $\square$

Lemma 4: *Privacy induced welfare loss is weakly decreasing in demand flexibility if the utility function is saturation level dependent.*

Proof: Assume a privacy enhanced market allocation for a given flexibility level f. If the demand flexibility is raised to $f' > f$, the assumed concavity of the utility functions leads to the following effect: The marginal utility $U^{'}_{c,t}(i \cdot D)$ with $i = 1...n$ for small i's is higher for f' than for f and drops faster for greater i's. Formally, let $\left[U^{'}_{c,t}(i \cdot D)\right]_f$ be the marginal utility for flexibility level f, then there exists a threshold $\widehat{i}$ where

$$\left[U^{'}_{c,t}(\widehat{i} \cdot D)\right]_f \leq \left[U^{'}_{c,t}(\widehat{i} \cdot D)\right]_{f'}$$

and

$$\left[U^{'}_{c,t}((\widehat{i}+1) \cdot D)\right]_f > \left[U^{'}_{c,t}((\widehat{i}+1) \cdot D)\right]_{f'}$$

holds. Since the higher valued units have a higher probability of being allocated, and a lower probability of being omitted if the privacy enhancement method changes the saturation level, the welfare loss is weakly lower for f'. $\square$

Note that if the utility is independent of the saturation level, the actual saturation $\overline{x}_c(t)$ respectively $\widetilde{x}_c(t)$ does not necessarily reach the threshold $\widehat{i}$. Thus, Lemma 4 does not hold for saturation level independent utility.

A privacy enhancement method leads to a distortion of saturation levels, this has the following effect: Replacing the saturation level $\overline{x}_c(t)$ with a distorted value $\widetilde{x}_c(t)$ could naturally have a *quantity* effect on the resulting bidding behavior. This becomes evident from Equation (1): Inflated values, i.e., $\widetilde{x}_c(t) > \overline{x}_c(t)$, lead to positive marginal utility assessments when the marginal utility was zero in the non-distorted case. Discounted values in turn, i.e., $\widetilde{x}_c(t) < \overline{x}_c(t)$, yield premature zero-marginal-utility assessments. However, remember that we ruled out untrue saturation level reports to the privacy enhancement method $\mathcal{M}$, as this potentially leads to a privacy breach.[3]

While we rule out quantity misreports, we are interested in characterizing privacy-aware markets that induce consumers to reveal their true valuation.

Definition 9 (Incentive Compatibility): A privacy-aware market mechanism is (in expectation) incentive compatible with respect to valuation if consumers cannot (in expectation) profitably deviate from placing bids that reflect their true valuation.

[3] The semantic of $\mathcal{M}$ is defined on sensitive and true personal data, the effects if applied to untrue data are unknown. Furthermore, a deviation from the true saturation level will not have influence on the bidding quantities of others if the privacy enhancement method is a $\cup$-homomorphism.

Thus, consumers will bid according to their true valuation in the presence of an incentive compatible privacy enhancement method.

Lemma 5: *An incentive compatible market mechanism retains this property in the presence of privacy enhancement, if the following holds: The privacy enhancement method is an $\cup$-homomorphism and the utility function is saturation level independent.*

Proof: While the distortion of the saturation level always affects the optimal quantity, it does not necessarily have an effect on the optimal bid price. The occurrence of a price effect hinges on the structure of the customer-utility function: If $U'_{c,t}$ is independent of $\overline{x}_c(t)$, the bid price will always reflect the customer's true valuation for all demand increments $x \in [0, \min\{\overline{x}_c(t), \widetilde{x}_c(t)\}]$. On the other hand, if the utility function is saturation level dependent, this leads to a price effect, and the consumers may strategize and report prices differing from their true valuation. The $\cup$-homomorphism property excludes incentives from true reports, since other consumers are not influenced. Consequently, incentive compatibility is preserved if the utility is independent from the saturation level, and the privacy enhancement method is a $\cup$-homomorphism. $\square$

5. Model Implementation

The actual privacy costs depend on a large number of possible influence factors. This includes different privacy preferences as well as fluctuating supply and demand patterns. To derive meaningful results we conduct simulations based on real-world data. Additionally simulations require an instantiations of all model components theoretically described in Section 3. In this section, we describe all the details for conducting simulations.

5.1. Demand Model

To perform a numerical evaluation of our scenario, we need a concrete instantiation of the utility model. We study two alternatives — one featuring dependency of marginal utility on the saturation level and one with independence. In both cases we have a parameter f that represents load flexibility (i.e., concavity). To improve comparability, we normalize the utility functions with a scalar representing a maximum saturation level A.

Denoting the allocation by $x_c(t)$, the saturation level by $\overline{x}_c(t)$ and the flexibility level by f, the function with saturation level *dependent* marginal utility (superscript D) is given by

$$U^D(x_c(t), \overline{x}_c(t), f) = \frac{\overline{x}_c(t)}{A} \sqrt[f]{\frac{\min\{x_c(t), \overline{x}_c(t)\}}{A}}.$$

For saturation-level-*independent* marginal utility (superscript I), we have

$$U^I(x_c(t), \overline{x}_c(t), f) = \sqrt[f]{\frac{\min\{x_c(t), \overline{x}_c(t)\}}{A}}.$$

By taking the first derivative with respect to $x_c(t)$ dependence and independence are easily verified: Leaving the minimum function aside, because it only reflects the upper border of the utility, the derivative of U^D is dependent on $\overline{x}_c(t)$, and the derivative U^I is not.

Consumption in the simulation is based on the CER Smart Metering data set [38]. This data set consists of roughly 5000 Irish homes with different numbers of inhabitants, measuring electricity consumption every 30 minutes over more than one year.

5.2. Market supply

We assume that three types of generators in the local market, namely PV sites, CHP units and conventional backup generation (balancing party). The PV and CHP model follows the one introduced by Buchmann et al. [16]. In the following we summarize the key points.

Photovoltaic Sites. The energy output of a PV site depends on its peak capacity, the sun intensity and the mounting angle on the roof. The peak capacity is the maximum power output technically possible, mainly determined by the size of the site and the quality of the modules.[4] The electricity output from PV sites has no marginal generation costs and consequently a zero ask price is quoted.

Combined Heat and Power Units. We consider CHP units in heat-led operation mode where power operation is governed by heating requirements. A typical household CHP unit will generate approximately $1kW$ of electricity output. Generation availability is driven by heating demand and thus depends heat storage or insulation but not on market parameters. To reflect this 'randomness', we simulate start and stop times based on empirical CHP unit data. Under heat-led operation, all operational costs can be attributed to heating demand with electricity output arising as a byproduct. Consequently, CHP output is also bid into the local market at a zero limit price.

Balancing Party. A standard economic assumption for modeling conventional backup generation is a convex cost function [39]. This reflects the technological heterogeneity on the supply side (merit order dispatch). A quadratic cost function is a simple example of such a supply curve [40]. We follow this rationale and assume that the balancing party quotes a bid price of $p(x) = \alpha \cdot x^2$ for the x-th unit of output.

5.3. Energy storage

Due to increased uncertainty on the supply side, energy storage is expected to play a more important role in future smart grids. Storage operators can

[4] In our simulation PV panel sizes are distributed according to recent data German installation data censored at a maximum of $11kWp$. Then, empirical generation data is used to simulate output. We apply random shifts to mimic heterogeneous panel orientations.

capitalize on the expected price fluctuations. This energy arbitrage motive has been investigated in the recent literature [41]. Our analysis adds a new economic perspective of active storage management. We investigate to what extent energy storage can mitigate the welfare loss due to privacy enhancing methods in local energy markets. For the sake of generality, we assume that each customer owns a generic energy store with capacity $\overline{\mathcal{B}}_c$ (in kWh), fill level $\mathcal{B}_c(t) \in [0, ..\overline{\mathcal{B}}_c]$ and efficiency level $\mathcal{L} < 1$.[5]

Departing from an economic storage operation paradigm, we posit a simple strategy. Denoting the deviation from the current saturation level $\overline{x}_c(t)$ by ξ, the following cases are possible:

1. $\xi > 0$ — Whenever privacy enhancement results in an upward distortion, the amount $\min\{\mathcal{L} \cdot \xi, \frac{\overline{\mathcal{B}}_c - \mathcal{B}_c(t)}{\mathcal{L}}\}$ is transferred to the storage unit.

2. $\xi < 0$ — In case of an allocation shortfall, due to high market prices or a downward distortion, customers withdraw the amount $\min\{\xi, \mathcal{B}_c(t)\}$ from the energy store.

This policy could be improved, e.g., by adopting dynamic threshold levels. However, by focusing on this rather naïve policy, we can isolate interactions between privacy enhancement and the presence of storage capacities.

5.4. Privacy Enhancement Methods

For our experiments we apply the wavelet privacy enhancement algorithm proposed by Papadimitriou et al. [28]. Additionally, we consider a slightly modified version of this algorithm which retains incentive compatibility. The concept of privacy in this modified version remains unchanged.

Wavelet Privacy Enhancement Algorithm. Uncorrelated noise applied to a time series is easily filtered out by means of wavelet based filtering [42]. To circumvent this, we need to apply noise dependent on the wavelet representation of the actual time series: Let K be the number of wavelet coefficients exceeding σ and N the total number of wavelet coefficients. Then noise with the standard deviation of $\sigma \cdot \sqrt{N/K}$ and the mean value is the current coefficient if it is greater than or equal to σ. This method is obviously *randomized*. Since all time series are treated independently it is also a $\cup$-*homomorphism*. Finally, the symmetry of the noise distribution ensures that the method is a *balanced modifier*.

Yet, the method is not monotonically increasing: Applying a higher threshold $\sigma_1 > \sigma_2$ most likely results in noise with a higher variance, but is only applied to fewer coefficients. In general, we cannot assess whether or not $\mathcal{M}_{\sigma_1}(s)$ will distort a single data point to a larger extent than $\mathcal{M}_{\sigma_2}(s)$.

[5]For sake of exposition, we only account for losses during charge.

Algorithm 1: Privacy enhancement method

Input: Privacy Parameter σ
Input: Noise parameter ϕ
Input: Set of time series $\mathcal{S}$
Result: Privacy enhanced time series $\widetilde{\mathcal{S}}$
foreach $\overline{x}_c(t) \in \mathcal{S}$ **do**
    $\widetilde{\overline{x}_c(l,t)} = DWT(\overline{x}_c(t))$ //Wavelet transform;
    $I_l = \left\{t : \left|\widetilde{\overline{x}_c(l,t)}\right| \geq \sigma\right\}$;
    foreach *level* l **do**
        $K = \sum_l K_l$ //coeffs exceeding l;
        $p = |N|/K$ //Noise 'density', N is number of coefficients;
        foreach *detail* $\widetilde{\overline{x}_c(l,t)}$ **do**
            if $t \in I_l$ **then**
                $\widetilde{\overline{x}_c(l,t)} += GaussRnd(0, \phi\sqrt{p})$;
            end
        end
    end
    $\widetilde{x}_c(t) = InvDWT(\widetilde{\overline{x}_c(l,t)})$;
    $\widetilde{\mathcal{S}} = \widetilde{\mathcal{S}} \cup \{\widetilde{x}_c(t)\}$;
end
return $\widetilde{\mathcal{S}}$;

Incentive Compatible Wavelet Privacy Enhancement Algorithm. In what follows, we propose a modification of the wavelet privacy enhancement algorithm referred to as *incentive compatible wavelet privacy enhancement algorithm (IC-wavelet privacy)*. As shown in Lemma 5, a privacy enhancement algorithm needs to be monotonously increasing to retain in-expectation incentive compatibility with respect to valuation. Our modification achieves monotonicity by decoupling the threshold for coefficients and the noise variance. To this end we introduce a parameter ϕ that determines the standard deviation of the applied noise. The detailed implementation is given in Algorithm 1. Note that by setting $\phi = \sigma$ the modified algorithm is identical to the unmodified wavelet privacy enhancement algorithm.

Choosing σ and ϕ. The same (absolute) σ may have very different effects on two different time series $\overline{x}_{c1}(t)$ and $\overline{x}_{c2}(t)$. σ may lead to a lot of modified coefficients in $\widetilde{x}_{c1}(t)$, since they exceed σ, while $\widetilde{x}_{c2}(t)$ remains unmodified. In order to keep the parameters comparable we choose σ and ϕ relative to the standard deviation of the time series currently modified. Let $\sigma, \phi \in [0, 1]$. Then the actual parameters σ_c and ϕ_c for time series $\overline{x}_c(t)$ are the product of σ, ϕ and the standard deviation of the time series $\overline{x}_c(t)$. Figure 3 illustrates the effect of the privacy enhancing techniques. The upper panel illustrates that the wavelet privacy enhancement method is not monotonically increasing: At many points of time, the perturbed time series with $\sigma = 80\%$ has a higher distance to the non-perturbed time series than the one perturbed with $\sigma = 100\%$. In contrast, the lower panel shows the corresponding results from the monotonically increasing and incentive compatible algorithm. The time series perturbed with the higher $\phi = 100\%$ usually has a higher distance to the one with a lower $\phi = 80\%$.

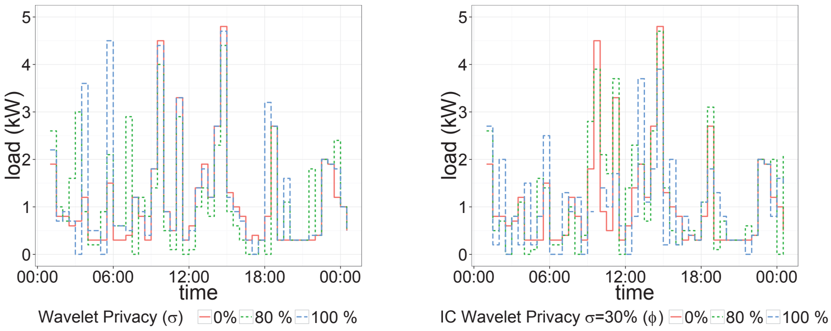

Figure 3: Examples of privacy enhancement methods realizations

6. Evaluation

We use a custom-built event-based simulator written in Java. The simulator architecture follows the MVC design principle with a decoupled GUI to allow headless simulation. The implementation is available from the authors upon request. The simulator supports the arbitrary combination of different utilities, privacy enhancement methods and data sets for households as well as energy suppliers.

6.1. Simulation Setup and Parameters

For the evaluation we choose 300 persons in total living in 191 randomly chosen households of the CER data set [38]. We assume homogeneous utility functions ($f = 2$, $A = 11kW$) across consumers. Additionally, to quantify the effect of demand side flexibility, we consider $f = 3$ and $f = 1$. Supply side is modeled as combinations of 25 or 150 PV and CHP sites with the balancing party parametrized with $\alpha = 2$. In scenarios with storage systems, we consider storage sizes of $\overline{\mathcal{B}} \in \{2.5kWh, 5kWh\}$, in line with currently marketed products. We assume storage efficiency of 80%. We apply the wavelet privacy enhancement method with $\sigma \in \{0, 0.2, 0.4, 0.6, 0.8, 1\}$. The time span for each simulation run is one day. We have tested longer simulation horizons as well, but these results did not exhibit any substantial differences to the ones described in the following. Each experiment is repeated ten times.

6.2. Effect of Market Structure

First, we investigate if and to what extent the number of PV and CHP sites influences the impact of different privacy enhancement levels. This sheds light on how privacy enhancement interacts with different market structures. The results are shown in Figure 4.

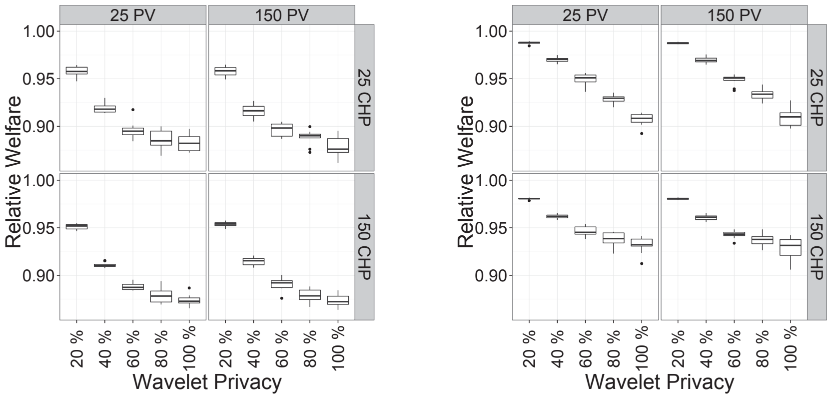

(a) saturation level independent utility (b) saturation level dependent utility

Figure 4: Impact of market configuration

Result 1: Market configuration has little impact on the welfare loss. While the variance of the results (size of boxes in Figure 4) is naturally higher in smaller markets with few generators, the median results are hardly affected.

Result 2: Saturation-level independent utility exhibits decreasing marginal welfare cost of the privacy level. While the welfare loss is strictly increasing in the privacy level for both utility specifications, saturation independent utility exhibits decreasing marginal losses in our results. For saturation-level dependent utility we observe almost linear behavior.

6.3. Effect of Storage

In theory, storage can help to reduce the welfare loss of privacy enhancement (see Lemma 3). This is because it is capable to store electricity bought that exceeds the saturation level. Thus, it is not 'wasted' but can be used in times of undersupply. The results of the simulations quantify the actual impact on welfare in a real-world scenario. The results show that storage can reduce the privacy-induced welfare loss by 70% (Figure 5). Next to this result that has been expected, at least in qualitative terms, we make the following observations.

Result 3: Small storage systems are sufficient to mitigate privacy costs. Under our naïve privacy-driven storage operation strategy, the 2.5 kWh system is almost as efficient as the 5 kWh system. This suggests that privacy costs may be mitigated at comparably low costs.

Result 4: The value of storage is increasing in the privacy level. Higher privacy level choices induce more frequent quantity mismatches which are mitigated by the storage system.

Result 5: Privacy enhancement may increase welfare in the presence of storage. With storage, relative welfare is not monotonically decreasing with the privacy

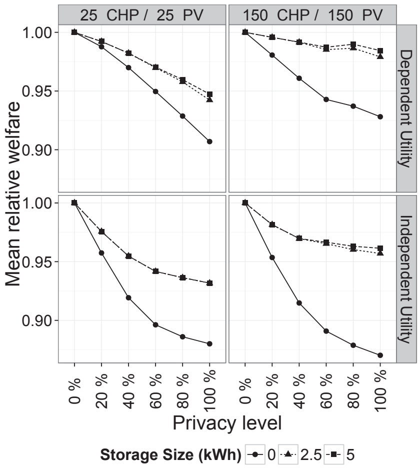

Figure 5: Effect of energy storage capacities

level. This is because privacy enhancement can induce economic dispatching of the storage system: Electricity bought above the saturation level $\overline{x}_c(t)$ usually is relatively 'cheap' due to the concave utility function. In times of undersupply the stored electricity is only used if less than $\overline{x}_c(t)$ is allocated on the market. Thus, the actual resulting utility of the stored electricity is much higher than the bid price in the former time slot.

Result 6: The value of storage is increasing in decentral generation capacity. In case of high decentral generation capacity, surplus energy stored will more often originate from these low cost sources. Consequently, the usage of the balancing party will decrease. This has a positive impact on social welfare.

This simulation results show that even small energy storage systems are very effective at mitigating the welfare loss due to privacy enhancement.

6.4. Impact of Demand Side Flexibility

We know from Lemma 4 that higher demand side flexibility curbs the influence of the privacy enhancement method on the welfare loss. However, the numerical results indicate that demand side flexibility can mitigate the welfare loss only to a very limited extent (Figure 6). For saturation-level-independent utility functions the influence of demand side flexibility is unpredictable (Lemma 4). Thus, we only cover the dependent utility function.

Both demand flexibility and storage can mitigate the welfare loss of privacy enhancement. However, our results suggest that storage has a much larger potential. Load flexibility only helps to reduce the welfare loss if $\widetilde{x}_c(t) < \overline{x}_c(t)$. Storage in turn also helps in cases of over-allocation: When the privacy enhancement method upward-adjusts the saturation level, the additional allocated

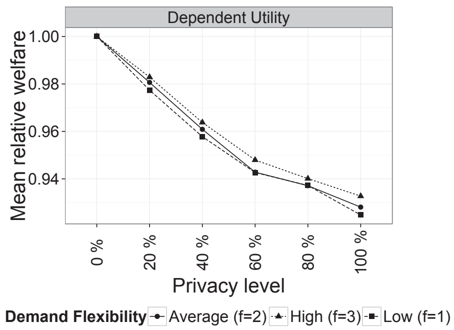

Figure 6: Effect of demand side flexibility

electricity is not lost. If the saturation level is downward-adjusted, storage can help to mitigate the welfare loss if $\mathcal{B}_c(t) > 0$.

6.5. Comparison of Privacy-Enhancement Methods

As noted before, the standard wavelet approach may induce strategic bidding on behalf of the consumers. This is because this privacy mechanism is not monotonously increasing. Here we want to analyze the variant of this algorithm for saturation level independent utility. The wavelet privacy parameter σ varies between 0% and 100%. To ensure that the second privacy enhancement method actually is monotonically increasing we choose a fixed $\sigma = 30\%$ and vary the noise ϕ only. We find that the welfare loss is more pronounced under our modified algorithm (Figure 7)

Result 7: The IC wavelet privacy method leads to a greater welfare loss. In the standard wavelet privacy-enhancement method, σ influences both the choice of coefficients perturbed and the standard deviation of the noise. For the incentive compatible method, the perturbed coefficients are always the same. This is because σ is fixed to 30%. For a higher privacy level the wavelet privacy method perturbs fewer coefficients, resulting in a higher welfare.The additional welfare loss can be interpreted as the cost of establishing incentive compatibility. While these costs remain negligible for privacy levels of up to 60%, they become more significant at higher privacy levels as the noise level is monotonically increasing.

7. Conclusion

Privacy-aware local energy markets are a promising approach for matching renewable supply and demand of private households. However, the potential effects of privacy-enhancement on the market outcome have so far remained vague. Our analysis provides the following novel insights on this subject matter:

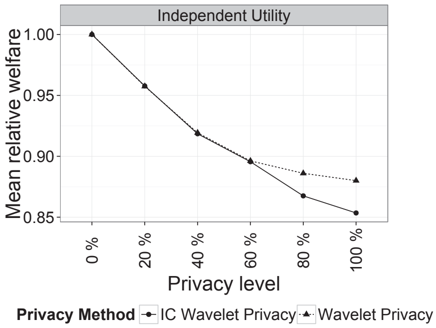

Figure 7: Cost of Incentive Compatibility

We provide a characterization of relevant privacy enhancement properties when applied in a market scenario. Under certain assumptions a market mechanism can retain incentive compatibility in the presence of privacy enhancement. Furthermore, we find that demand side flexibility and storage systems will mitigate the welfare loss due to privacy enhancement. Our numerical study shows that market mechanism remain a fairly efficient means for matching supply and demand in a local energy market even in the presence of privacy enhancement: We find that the negative impact of privacy enhancement is rather low, even for strong privacy requirements the welfare loss remains below 15%. Energy storage, even in small sizes, reduces the welfare loss. For strong privacy guarantees, the welfare loss is only approximately 5% in the presence of storage.

Our analysis only considers *direct* customer utility from electricity consumption. Privacy enhancement in turn is only accounted for as an exogenous parameter. However, it is conceivable that customers obtain a benefit from protection of personal information. If this *indirect* utility could be quantified, customers would face a trade-off between energy costs and the level of privacy protection. In such a setting, our model would allow determining an 'optimal' level of privacy protection.

The overall conclusion is that privacy enhancement methods are applicable in local energy markets including private households. From an economic perspective, the negative allocative effects are low and controllable while privacy enhancement significantly increases the privacy protection of participating individuals. From a computer science perspective, these markets are a meaningful performance indicator for the utility of privacy enhancement methods.

References

[1] H. Farhangi, The path of the smart grid, IEEE Power and Energy Magazine 8 (1) (2010) 18–28.

[2] S. M. Amin, B. F. Wollenberg, Toward a smart grid: power delivery for the 21st century, IEEE Power and Energy Magazine 3 (5) (2005) 34–41.

[3] H. Gharavi, R. Ghafurian, Smart grid: The electric energy system of the future, Proceedings of the IEEE 99 (6) (2011) 917–921.

[4] S. D. Ramchurn, P. Vytelingum, A. Rogers, N. R. Jennings, Putting the 'smarts' into the smart grid: a grand challenge for artificial intelligence, Communications of the ACM 55 (4) (2012) 86–97.

[5] F. Hvelplund, Renewable energy and the need for local energy markets, Energy 31 (13) (2006) 2293–2302.

[6] S. Lamparter, S. Becher, J.-G. Fischer, An agent-based market platform for smart grids, in: Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems, 2010, pp. 1689–1696.

[7] T. Ackermann, G. Andersson, L. Söder, Distributed generation: a definition, Electric power systems research 57 (3) (2001) 195–204.

[8] G. Hart, Nonintrusive appliance load monitoring, Proceedings of the IEEE 80 (12) (1992) 1870–1891.

[9] A. Molina-Markham, P. Shenoy, Private memoirs of a smart meter, Proceedings of the BuildSys (2010) 61–66.

[10] P. McDaniel, S. McLaughlin, Security and privacy challenges in the smart grid, IEEE Security and Privacy 7 (3) (2009) 75–77.

[11] M. Pagano, A. Roell, Transparency and Liquidity: A Comparison of Auction and Dealer Markets with Informed Trading, The Journal of Finance 51 (2) (1996) 579–611.

[12] A. Madhavan, Security prices and market transparency, Journal of Financial Intermediation 5 (3) (1996) 255–283.

[13] J. Z. Kolter, M. J. Johnson, REDD: A public data set for energy disaggregation research, in: Workshop on Data Mining Applications in Sustainability, San Diego, CA, 2011.

[14] U. Greveler, B. Justus, D. Loehr, Multimedia content identification through smart meter power usage profiles, in: Proceedings of the 5th International Conference on Computers, Privacy and Data Protection, 2012.

[15] E. Buchmann, K. Böhm, T. Burghardt, S. Kessler, Re-identification of smart meter data, Personal and Ubiquitous Computing 17 (4) (2013) 653–662.

[16] E. Buchmann, S. Kessler, P. Jochem, K. Böhm, The costs of privacy in local energy markets, in: Proceedings of the 15th IEEE Conference on Business Informatics, 2013, pp. 198–207.

[17] G. Kalogridis, C. Efthymiou, S. Denic, T. Lewis, et al., Privacy for smart meters: Towards undetectable appliance load signatures, in: Proceedings of the 1st IEEE International Conference on Smart Grid Communications, 2010, pp. 232–237.

[18] D. Varodayan, A. Khisti, Smart meter privacy using a rechargeable battery: Minimizing the rate of information leakage, Proceedings of the IEEE Conference on Acoustics, Speech and Signal Processing (2011) 1932–1935.

[19] B. C. M. Fung, K. Wang, R. Chen, P. S. Yu, Privacy-preserving data publishing, ACM Computing Surveys 42 (4) (2010) 14.

[20] R. Trujillo-Rasua, J. Domingo-Ferrer, On the privacy offered by (k, δ)-anonymity, Information Systems 38 (2013) 491–494.

[21] M. E. Nergiz, M. Atzori, Y. Saygin, Towards trajectory anonymization: A generalization-based approach, in: Proceedings of the SIGSPATIAL ACM GIS 2008 International Workshop on Security and Privacy, Vol. 2, 2008, pp. 52–61.

[22] O. Abul, F. Bonchi, M. Nanni, Never walk alone: Uncertainty for anonymity in moving objects databases, in: Proceedings of the 24th IEEE International Conference on Data Engineering, 2008, pp. 376–385.

[23] D. N. Jutla, P. Bodorik, Y. Zhang, PeCAN: An architecture for users' privacy-aware electronic commerce contexts on the semantic web, Information Systems 31 (4-5) (2006) 295–320.

[24] C. Dwork, Differential privacy, in: Automata, Languages and Programming, Vol. 4052 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2006, pp. 1–12.

[25] G. Ács, C. Castelluccia, I have a dream! (differentially private smart metering), in: Information Hiding, Vol. 6958 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2011, pp. 118–132.

[26] F. Brandt, Fully private auctions in a constant number of rounds, in: Financial Cryptography, Vol. 2742 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2003, pp. 223–238.

[27] D. C. Parkes, M. O. Rabin, S. M. Shieber, C. Thorpe, Practical secrecy-preserving, verifiably correct and trustworthy auctions, Electronic Commerce Research and Applications 7 (3) (2008) 294–312.

[28] S. Papadimitriou, F. Li, G. Kollios, P. S. Yu, Time series compressibility and privacy, in: Proceedings of the 33rd international conference on Very large data bases, 2007, pp. 459–470.

[29] Y.-S. Moon, H.-S. Kim, S.-P. Kim, E. Bertino, Publishing time-series data under preservation of privacy and distance orders, in: Database and Expert Systems Applications, Vol. 6262 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2010, pp. 17–31.

[30] X. Jin, N. Zhang, G. Das, ASAP: Eliminating algorithm-based disclosure in privacy-preserving data publishing, Information Systems 36 (5) (2011) 859–880.

[31] C. Weinhardt, C. Holtmann, D. Neumann, Market-engineering, Wirtschaftsinformatik 45 (6) (2003) 635–640.

[32] D. Ilic, P. G. Da Silva, S. Karnouskos, M. Griesemer, An energy market for trading electricity in smart grid neighbourhoods, in: Proceedings of the 6th IEEE International Conference on Digital Ecosystems Technologies, 2012, pp. 1–6.

[33] P. Vytelingum, S. D. Ramchurn, T. D. Voice, A. e. Rogers, Trading agents for the smart electricity grid, in: Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems, 2010, pp. 897–904.

[34] S. Parsons, J. A. Rodriguez-Aguilar, M. Klein, Auctions and bidding: A guide for computer scientists, ACM Computing Surveys 43 (2) (2011) 10.

[35] F. C. Schweppe, M. C. Caramanis, R. D. Tabors, Evaluation of spot price based electricity rates, IEEE Transactions on Power Apparatus and Systems (7) (1985) 1644–1655.

[36] E. Bitar, S. Low, Deadline differentiated pricing of deferrable electric power service, in: Proceedings of the 51st Annual Conference on Decision and Control (CDC), IEEE, 2012.

[37] L. Sweeney, k-Anonymity: A model for protecting privacy, IEEE Security And Privacy 10 (5) (2002) 557–570.

[38] Irish Social Science Data Archive, Electricity customer behaviour trial, `http://www.ucd.ie/issda/` (2012).

[39] S. D. Ramchurn, P. Vytelingum, A. Rogers, N. Jennings, Agent-based control for decentralised demand side management in the smart grid, in: Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 1, 2011.

[40] P. Scott, S. Thiébaux, M. van den Briel, P. Van Hentenryck, Residential demand response under uncertainty, in: Principles and Practice of Constraint Programming, Springer, 2013.

[41] P. Vytelingum, T. D. Voice, S. D. Ramchurn, A. Rogers, et al., Agent-based micro-storage management for the smart grid, in: Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems, 2010.

[42] D. L. Donoho, I. M. Johnstone, Adapting to unknown smoothness via wavelet shrinkage, Journal of the American Statistical Association 90 (432) (1995) 1200–1224.