Privacy by Design in Local Electricity Markets: A Differentially Private Market Mechanism

Milad Hoseinpour
Tarbiat Modares University – Iran
m.hoseinpour@modares.ac.ir

Mahmoud-Reza Haghifam
Tarbiat Modares University – Iran
haghifam@modares.ac.ir

ABSTRACT

Privacy-preserving electricity markets have a key role in steering customers towards participation in local electricity markets by guarantying to protect their sensitive information. Moreover, these markets make it possible to safely release and share the market outputs for social good. This paper aims to design a privacy-preserving market for local energy communities by implementing Differential Privacy (DP) that provably guarantees the privacy of market participants. Besides achieving a near-optimal solution and preserving the privacy of market participants, the proposed market-clearing mechanism maintains the utility of the market data for statistical releases. The required randomization for satisfying DP is embedded as Gaussian noise in each iteration of the gradient ascent algorithm underlying the optimization process of the market-clearing problem. In numerical studies, we investigate the impacts of the privacy loss parameter on the output distribution of the market-clearing quantities. In addition, we demonstrate the inherent trade-off between privacy protection and social welfare in the market under different privacy regimes.

Keywords*: local electricity markets, differential privacy, mechanism design, data privacy.*

INTRODUCTION

The emerging Local Electricity Markets (LEMs) are an abundant source of individuals' data, such as financial data and electricity transactions. Statistical releases of these data and giving access to researchers, business owners, policymakers, etc., brings a multitude of economic as well as technical and societal benefits. In addition, there are transparency acts, e.g., Commission Regulation No. 543/2013 in the European Union (EU), on submission and publication of data in wholesale electricity markets, which will be extended to LEMs for facilitating market integration and development of Renewable Energy Sources (RESs).

However, these rich and fine-grained datasets could expose individuals to a privacy breach and reveal sensitive information about them. Therefore, the privacy concern can motivate the market participants to behave strategically by misreporting their data or even opt out of the market. Moreover, data privacy laws and regulations, e.g., General Data Protection Regulation (GDPR) passed by the EU, impose legal obligations on local energy communities to safeguard the private information of the market participants [1]. The GDPR requires LEMs to implement a Privacy by Design (PbD) paradigm for preserving the privacy of individuals. That is, the privacy protection measures should be embedded in the design of the market, which is in contrast to the preventive and passive actions for maintaining data privacy. The main question motivating this paper is how to publicly release the market-clearing outputs while simultaneously maintain the privacy of market participants and utility of their data for social good. In this regard, this paper proposes a PbD paradigm for LEMs by applying the notion of Differential Privacy (DP). DP is a formal mathematical framework for preserving the privacy of individuals in a dataset while still allowing for useful data analysis. DP is the de-facto standard for privacy protection. In fact, other privacy-preserving approaches, e.g., data anonymization and k-anonymity, are fragile under appropriate side information, while DP makes no assumptions about an adversary's computational power or side information.

The cryptographic literature also addresses the privacy concerns in electricity markets. For instance, in [2], a decentralized privacy-preserving protocol based on secure Multi-Party Computation (MPC) is proposed for local electricity markets. The proposed model performs the bid selection and calculation of the market-clearing price in a data oblivious and secure manner. In [3], a secure double auction mechanism is proposed for smart grids. To protect the market participants anonymity and privacy, a pseudo-identity is assigned to each participant, and their bids are encrypted using a cryptosystem. A privacy-preserving Peer-to-Peer (P2P) energy trading platform is proposed in [4]. The private information of market participants, including sellers' price and buyers' demand, are encrypted based on homomorphic encryption cryptosystem. However, our paper fundamentally differs from these works. Indeed, the overarching goal of our paper is to safely promote public access to electricity markets outputs, whereas the goal of cryptographic realization of privacy-preserving electricity markets is to hide all information except for running the market-clearing mechanism. In other words, cryptography solves the data security problem by using encryption algorithms and preventing unauthorized access to individuals' sensitive data [5].

PROBLEM SETUP

We consider a centralized market-clearing problem in a local energy community with a finite set of market participants denoted by Ω, consisting of producers Ω^p and consumers Ω^c. There is trustworthy market operator who centrally collects the private information of the market

participants and runs the market. Moreover, the market operator should publicly release the market outputs and, meanwhile, has obligation to preserve the privacy of market participants due to data privacy laws. For notational brevity, the following notations are based on a general agent without distinguishing producers and consumers. There is a set of potential social decisions $S = \prod_{i=1}^{n} S_i$, where $S_i \subset R^{|S_i|}$ is the domain of agent i's potential local decisions. Then, the local decisions $s_i \in S_i$ of consumer i and producer i are characterized by demand $d_i \in \left[\underline{d_i}, \overline{d_i}\right]$ and active power $g_i \in \left[\underline{g_i}, \overline{g_i}\right]$, respectively.

Each agent $i \in \Omega$ is endowed with a private information $\theta_i \in \Theta_i$, called type, that represents its preferences over the set of potential decisions S. Given a type, agent i's preferences can be evaluated by a valuation function $v_i: S \times \theta_i \to R$, where $v_i(s, \theta_i)$ denotes the value of alternative $s \in S$ for agent i with type θ_i. Moreover, since we assume that agent i's valuation depends only on its local decisions, we can say $v_i(s, \theta_i) = v_i(s_i, \theta_i)$. The valuation function of consumer i reflects the utility of using demand d_i, denoted by $v_i(d_i, \theta_i) = U_{i,\theta_i}(d_i)$. Also, for producer i, the valuation function reflects the negative of the cost of generating active power g_i, denoted as $v_i(g_i, \theta_i) = -C_{i,\theta_i}(g_i)$. Note that $U_{i,\theta_i}(\cdot)$ and $C_{i,\theta_i}(\cdot)$ are the utility and cost functions of consumer i and producer i respectively, parameterized by type θ_i. Furthermore, for mathematical convenience, we normalize the valuation functions so that their range is [0,1].

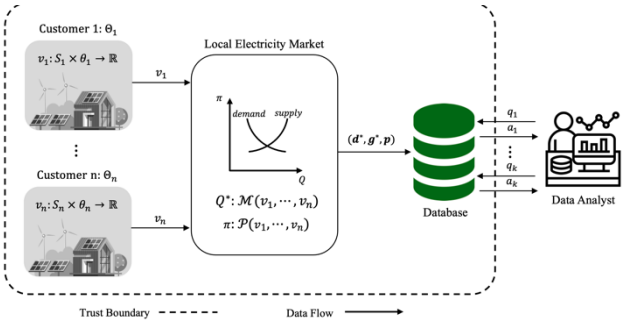

Figure 1: A high-level view of the problem setup.

Figure 1 depicts a high-level view of the problem setup. As can be seen, customer i, $\forall i \in \Omega$, reports the valuation v_i to the market. Given the valuation profile $v = (v_i)_{i \in \Omega}$, a trusted market operator defines an allocation rule $\mathcal{M}(v)$ for determining the market-clearing quantities, $d^* = (d_i^*)_{i \in \Omega^c}$ and $g^* = (g_i^*)_{i \in \Omega^p}$, and a payment rule $\mathcal{P}(v)$ for determining customers' payment $p = (p_i)_{i \in \Omega}$. In electricity markets, the allocation rule is determined by maximizing the social welfare function $\mathrm{sw}(v, s) = \sum_{i \in \Omega} v_i(s_i)$ subject to technical constraints of the market participants and the market-clearing equation. By substituting the valuation functions of consumers and producers in $\mathrm{sw}(v, s)$, the allocation rule $\mathcal{M}(v)$ is based on:

$$(d^*, g^*) \in \arg\max_{d,g} \sum_{i \in \Omega^c} U_{i,\theta_i}(d_i) - \sum_{i \in \Omega^p} C_{i,\theta_i}(g_i) \quad (1a)$$

$s.t.$

$$\underline{d_i} \le d_i \le \overline{d_i}, \forall i \in \Omega^c \quad (1b)$$

$$\underline{g_i} \le g_i \le \overline{g_i}, \forall i \in \Omega^p \quad (1c)$$

$$\sum_{i \in \Omega^p} g_i - \sum_{i \in \Omega^c} d_i = 0, \quad (1d)$$

where constraints (1b) and (1c) reflect the demand and supply limits of the consumers and producers, respectively. Also, constraint (1d) relates to the market-clearing equation. Hence, the output of the market-clearing problem is the tuple (d^*, g^*, p), which is stored in a database. The data analyst represents all the third parties, e.g., energy efficiency service providers, policymakers, and insurance companies, requesting access to the market-clearing outputs for purposes that are not expected by the market participants. The goal of a well-intentioned data analyst is to learn useful statistics about the population of the market participants by making queries $\{q_j\}_{j=1}^{k}$ and receiving answers $\{a_j\}_{j=1}^{k}$ from the database. But, the statistical release of the market-clearing outputs exposes the market participants to the risk of a privacy breach by malicious third parties. Thus, the privacy-preserving market-clearing mechanism should ensure that whoever outside of the trust boundary is not able to learn anything at the individual level and violate the privacy of individuals.

METHODOLOGY OVERVIEW

To achieve our privacy goal in this paper, we implement DP. The main idea in DP is to mask the contribution of any individual's data in the computation by adding a calibrated amount of noise. In the following, we provide the formal definition of DP and, then, introduce Gaussian mechanism as the perturbation technique for satisfying DP.

Definition 1 (Differential Privacy). For $\epsilon \ge 0$ and $\delta \in [0,1]$, a randomized algorithm $\mathcal{M}$ with domain $\mathcal{X}^n$ is (ϵ, δ)-differentially private if for every pair of neighboring datasets $x \sim x' \in \mathcal{X}^n$ (i.e., x and x' differ in one element) and for all $S \subseteq \mathcal{R}$ [5]:

$$\Pr[\mathcal{M}(x) \in S] \le e^{\epsilon}.\Pr[\mathcal{M}(x') \in S] + \delta, \quad (2)$$

where the probability space is over the internal randomness of $\mathcal{M}$. The aforementioned definition is about the behavior of $\mathcal{M}$ and promises that no individual's data has a large impact in the output of the algorithm. More formally, when a (ϵ, δ)-DP[1] algorithm runs on two neighboring datasets, the resulting distributions over the

1 For simplicity of expressions, we replace "(ϵ, δ)-differentially private" with "(ϵ, δ)-DP".

output space will be very similar, and this similarity is captured by a multiplicative factor e^{ϵ} and an additive factor δ. The required noise for satisfying DP is calibrated based on the global sensitivity of the computation. We formalize the mathematical definition of the global sensitivity in the following.

Definition 2 (Global Sensitivity). For function $f: \mathcal{X}^n \rightarrow \mathbb{R}^k$, the ℓ_2 −sensitivity over all pairs of neighboring datasets $x \sim x' \in \mathcal{X}^n$ is [6]

$$\Delta(f) = \max_{x \sim x' \in \mathcal{X}^n} \|f(x) - f(x')\|_2. \quad (3)$$

Definition 3 (Gaussian Mechanism). If $f: \mathcal{X}^n \rightarrow \mathbb{R}^k$, then Gaussian mechanism directly add Gaussian noise to the output of the computation in the following sense:

$$\mathcal{M}(x) = f(x) + (Y_1, \cdots, Y_k), \quad (4)$$

where $(Y_i)_{i=1}^k$ are independent random variables drawn from $\mathcal{N}(0, 2\ln(1.25/\delta)\Delta(f)^2/\epsilon^2)$ [6].

THE MODEL

For achieving a differentially private market-clearing mechanism, we need to bound the contribution of any market participant's data on the computation of market-clearing quantities and payments. In this regard, the privacy mechanism and pricing mechanism of the market along with their synthesis approach are the key modelling choices we should make. In the next section, we demonstrate that naively applying additive noise approaches to privatize the market-clearing quantities and price is problematic and doesn't work well. Then, we propose our differentially private market-clearing mechanism to overcome the existing challenges.

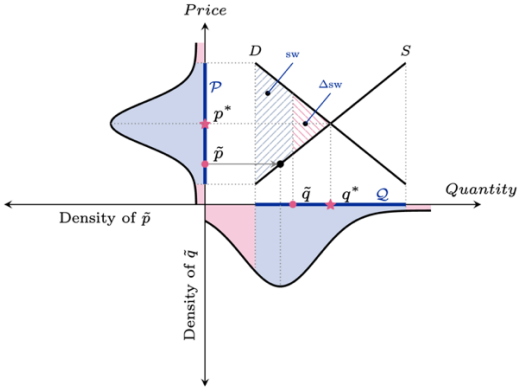

Figure 2: A trivial approach for achieving a differentially private market-clearing mechanism

A Trivial Approach

The supply curve S and demand curve D of a local electricity market based on the Uniform Price Double Auction (UPDA) mechanism are depicted in Figure 2, where $\mathcal{Q}$ and $\mathcal{P}$ are the feasible regions for the market-clearing quantities and price, respectively. The intersection of the supply and demand curves shows $q^* \in \mathcal{Q}$ as the optimal market-clearing quantity and $p^* \in \mathcal{P}$ as the optimal market-clearing price. Due to the privacy concern of the market participants, we cannot publicly release these quantities. In this regard, for unlocking the benefits of the data that local electricity markets can offer, we should compute the market-clearing outcomes in a DP framework. Thus, we implement Gaussian mechanism for privatizing the market-clearing outcomes via output perturbation. Due to Figure 2, in the first step, we add a calibrated amount of Gaussian noise to $p^* \in \mathcal{P}$ that results in a Gaussian Probability Density Function (PDF) over the price axis. So, the market-clearing price is determined by a random realization $\tilde{p}$ of this PDF, which its mean is centred on p^*. It should be mentioned that the private computation of the market-clearing price may lead to an infeasible price $\tilde{p} \notin \mathcal{P}$, indicated by the shaded areas in red.

Besides the market-clearing price, we should also privately compute the market-clearing quantities. Thus, we map $\tilde{p} \in \mathcal{P}$ onto a corresponding market-clearing quantity via the supply curve S. Since this mapping touches the sensitive information of the market participants encapsulated in the supply curve S, we should do this mapping privately. In this regard, we add a calibrated amount of Gaussian noise to the mapped point on the quantity axis, which results in a Gaussian PDF centred on this point. However, simply adding noise to the market-clearing quantities may lead to an infeasible solution $\tilde{q} \notin \mathcal{Q}$ that entails corrective mechanisms by the market operator. Besides, in this technique, we have no measure to guarantee a near-optimal solution, and the social welfare may suffer drastically. As you can see in Figure 2, the social welfare is reduced by Δsw quantity under the privacy constraint in compare to the non-private solution. To tackle these challenges, we should embed the required randomization in the market-clearing mechanism in such a way to ensure the feasibility and quality of the market-clearing outcomes.

Proposed Approach

In our model, we embed the required randomization in the optimization process of the market-clearing problem, which is based on the Gradient Ascent (GA) algorithm. That is, we add a calibrated amount of Gaussian noise to each iteration of the GA updating rule. Also, for ensuring the feasibility of the market-clearing outputs, we apply a projection method in the updating rule of GA. Figure 3 represents the computational block in the GA algorithm that we should compute privately, where $v = (v_i)_{i \in \Omega}$ is the valuation profile of the market participants, s is the vector of market-clearing quantities, $g_t(\cdot)$ is the gradient of the social welfare function $\mathrm{sw}(v, s) = \sum_{i=1}^n v_i(s_i)$, and η is the updating rate.

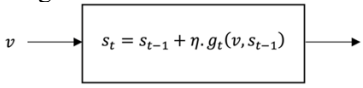

Figure 3:Updating rule of the market participants' decisions

For calibrating the required amount of Gaussian noise, we need to compute the global sensitivity of the aforementioned computational block in Figure 3. In fact, we should compute the global sensitivity of g_t, $\Delta(g_t) = \nabla_s \mathrm{sw}(v, s_{t-1})$. In most of the cases, there is no upper bound for $\Delta(g_t)$. However, due to the robustness of the GA algorithm, we can clip the gradient $g_t(v, s_{t-1})$ in each iteration of the algorithm with an arbitrary bound C. Thus, the gradient vector should be replaced with

$g/max(1, \|g\|_2/C)$, where C is the clipping threshold. In this regard, $\Delta(g_t)$ is equal to

$$\Delta(\boldsymbol{g_t}) = \max_{v \sim v'} \|\nabla \mathrm{sw}(\boldsymbol{v}, \boldsymbol{s_{t-1}}) - \nabla \mathrm{sw}(\boldsymbol{v'}, \boldsymbol{s_{t-1}})\|_2 \leq \max_{v \sim v'} (\|\nabla \mathrm{sw}(\boldsymbol{v}, \boldsymbol{s_{t-1}})\|_2 - \|\nabla \mathrm{sw}(\boldsymbol{v'}, \boldsymbol{s_{t-1}})\|_2) = \boldsymbol{2C}. \quad (5)$$

So, for achieving $(\epsilon', \delta') - DP$ in each iteration of GA, we should add Gaussian noise with parameter $\sigma \geq \frac{2C}{n\epsilon'}\sqrt{2\ln\left(\frac{1.25}{\delta'}\right)}$ to $g_t(v, s_{t-1})$ in each iteration. Algorithm 1 summarizes the proposed differentially private market-clearing mechanism.

Algorithm 1: Differentially private market-clearing mechanism

Inputs: set of valuation functions $v = (v_i)_{i\in\Omega}$, social welfare function $\mathrm{sw}(v, s) = \sum_{i=1}^{n} v_i(s_i)$, set of feasible solutions $\mathcal{O} \subseteq \mathbb{R}^n$, number of iterations T, updating rate η, noise parameter σ, clipping threshold C.

Outputs: market-clearing quantities at step T, s_T.

1: Initialization of s_0 with an arbitrary point in $\mathcal{O}$

2: **for** $t \in [T]$:

3: compute the gradient of the social welfare function $g_t = \nabla_s \mathrm{sw}(v, s_{t-1})$

4: clip the gradient:

$$g_t^{clip} = \frac{g_t}{max(1, \|g_t\|_2/C)}$$

5: add noise:

$$\tilde{g}_t = g_t^{clip} + \mathcal{N}(0, \sigma^2 I_n)$$

6: updating rule:

$$u_t = s_{t-1} + \eta \tilde{g}_t$$

7: projection onto the feasible region:

$$s_t = \Pi_{\mathcal{O}}(u_t)$$

8: **end for**

9: **return** s_T.

Differentially private computation of payments

Besides the market-clearing quantities, the payments of the market participants will be published publicly. Hence, an adversary who tries to learn about the private valuations of the market participants has access to all the payments. Thus, we should make the payment profile $p = (p_1(v), \cdots, p_n(v))$ of the market indistinguishable via DP. In this regard, for every pair of the neighboring valuation profiles $v \sim v' \in V^n$ and any possible payment $p \in \mathcal{P}$, the following privacy constraint should hold:

$$\Pr[p_1(v), \cdots, p_n(v) \in \mathcal{P}] \leq e^{\epsilon}.\Pr[p_1(v'), \cdots, p_n(v') \in \mathcal{P}]. \quad (6)$$

In this paper, the payments of the market participants are determined based on the Vickerly-Clarke-Groves (VCG) mechanism. In this mechanism, market participants report their valuation functions $v = (v_i)_{i\in\Omega}$ to the market, and the VCG mechanism selects an outcome s^* that maximizes the social welfare function. Then, the VCG mechanism charges each agent i with its social cost p_i, which is the difference between the social welfare of others in the absence and presence of agent i:

$$p_i(v_i, v_{-i}) = \max_{s\in S} \sum_{j\neq i\in\Omega} v_j(s_j) - \sum_{j\neq i\in\Omega} v_j(s^*). \quad (7)$$

As we can see, computing the social welfare is the building block of $p_i(v_i, v_{-i})$ in (7). So, there is no need for an additional differentially private mechanism, and it suffices to embed Algorithm 1 in the non-private computation of the VCG payments, which results in Algorithm 2.

Algorithm 2: Private computation of the VCG payments

Inputs: set of valuation functions $v = (v_i)_{i\in\Omega}$, privacy loss parameter ϵ.

Outputs: expected value of the market participants' payments p.

1: **call** Algorithm 1:
Inputs: $v = (v_i)_{i\in\Omega}, \epsilon$
Outputs: s$\sim \mathcal{D}$

2: **for all** agents $i \in \Omega$ **do**

3: **call** Algorithm 1:
Inputs: $v = (v_j)_{j\in\Omega, j\neq i}, \epsilon$
Outputs: s$\sim \mathcal{D}_{-i}$.

4: $\mathrm{sw}_{-i}(\mathcal{D}_{-i}) = \mathbb{E}_{s\sim\mathcal{D}_{-i}}[\sum_{j\neq i} v_j(s)]$

5: $\mathrm{sw}_{-i}(\mathcal{D}) = \mathbb{E}_{s\sim\mathcal{D}}[\sum_{j\neq i} v_j(s)]$

6: $p_i = \mathrm{sw}_{-i}(\mathcal{D}_{-i}) - \mathrm{sw}_{-i}(\mathcal{D})$

7: end for

8: return p.

NUMERICAL RESULTS

This section presents numerical case studies to reflect the theoretical properties of the proposed differentially private mechanism for local electricity markets. Towards that, we provide a local energy community consisting of three producers and three consumers. The cost function of producer i and utility function of consumer i are in the quadratic format $C_{i,\theta_i}(\cdot) \coloneqq a_i^g g_i^2 + b_i^g g_i + c_i^g$ and $U_{i,\theta_i}(\cdot) \coloneqq a_i^u d_i^2 + b_i^u d_i + c_i^u$, respectively. The parameters for producers and consumers are given in Table 1 and Table 2, respectively.

Table 1: Characteristics of producers

Producers	a_i^g ($/kWh^2)	b_i^g ($/kWh)	c_i^g ($)	$\underline{g_i}$ (kW)	$\overline{g_i}$ (kW)
1	0.015	0.038	0	0	20
2	0.008	0.047	0	0	25
3	0.011	0.056	0	0	30

Table 2: Characteristics of consumers

Consumers	a_i^u ($/kWh^2)	b_i^u ($/kWh)	c_i^u ($)	$\underline{d_i}$ (kW)	$\overline{d_i}$ (kW)
1	- 0.008	0.8	0	5	15
2	- 0.014	0.5	0	5	18
3	- 0.009	0.4	0	10	25

Market-clearing quantities

As we expect, the outputs of a differentially private market-clearing mechanism will be in the form of PDFs. The PDFs of market-clearing quantities for consumers and producers are depicted for $\epsilon = 0.05$, $\epsilon = 5$, and $\epsilon = 100$ in Figure 4 and Figure 5, respectively. Also, the mean μ and variance σ^2 of each PDF $\mathcal{N}(\mu, \sigma^2)$ is included. By decreasing the privacy loss parameter ϵ, which implies a higher level of privacy protection, the variance of the market-clearing quantities' PDF increases. In fact, for providing a higher level of privacy protection, we need to add more randomization to the market-clearing quantities and increase the entropy of the market. In an asymptotic analysis, we can say the highest degree of the privacy protection corresponds to a uniform distribution over the market-clearing quantities, when the market selects a pure random solution as the output of the market-clearing problem.

Cost of privacy

The focus of this section is on the trade-off between the social welfare of the market and the level of privacy guarantee provided for the market participants. We don't get the privacy protection for free, and it decreases the social welfare. That is, the embedded randomization for satisfying DP deviates the market from the optimal point. Figure 6 shows the PDF of social welfare for different values of ϵ. You can see that for $\epsilon = 0.05$, there is a significant gap between the optimal social welfare $\text{sw}^* = 10.97$ \$ and the expected value of the corresponding PDF, $\mathbb{E}[\text{sw}] = 7.63$ \$. In contrast, when $\epsilon = 100$ the PDF of social welfare is more leaned towards the optimal value, and with a high guarantee we achieve a near-optimal solution.

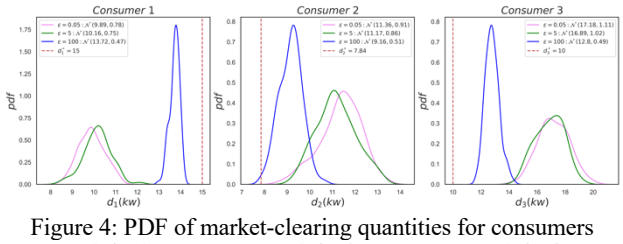

Figure 4: PDF of market-clearing quantities for consumers

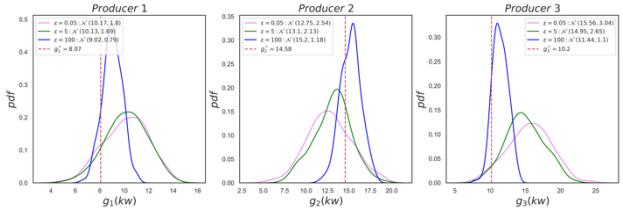

Figure 5: PDF of the market-clearing quantities for producers

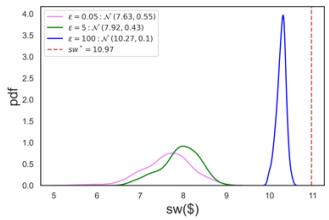

Figure 6: PDF of the market-clearing social welfare

CONCLUSION

In this paper, we proposed a differentially private market-clearing mechanism for local electricity markets. We implemented Gaussian mechanism for privatizing the underlying allocation rule and payment rule of the market, which is based on VCG mechanism. In the numerical results, we demonstrated the inherent trade-off between the model performance and privacy guarantee. We showed that higher level of privacy protection comes with considerable deviation of the market-clearing quantities from the optimal values, which degrades the fidelity of the market outcomes. In contrast, when the market participants are not overly concerned with their privacy the market-clearing quantities are more biased towards their optimal values. In addition, we saw that the privacy cost reflected as the gap between the private and non-private social welfare is directly related to the level of privacy protection.

REFERENCES

[1] D. Lee and D. J. Hess, "Data privacy and residential smart meters: Comparative analysis and harmonization potential," Utilities Policy, vol. 70, p. 101188, 2021.

[2] A. Abidin, A. Aly, S. Cleemput, and M. A. Mustafa, "An mpc-based privacy-preserving protocol for a local electricity trading market," in In- ternational Conference on Cryptology and Network Security. Springer, 2016, pp. 615–625.

[3] R. Sarenche, M. Salmasizadeh, M. H. Ameri, and M. R. Aref, "A secure and privacy-preserving protocol for holding double auctions in smart grid," Information Sciences, vol. 557, pp. 108–129, 2021.

[4] K. Erdayandi, A. Paudel, L. Cordeiro, and M. A. Mustafa, "Privacy- friendly peer-to-peer energy trading: A game theoretical approach," arXiv preprint arXiv:2201.01810, 2022.

[5] S. Vadhan, "The complexity of differential privacy," in Tutorials on the Foundations of Cryptography. Springer, 2017, pp. 347–450.

[6] C. Dwork, A. Roth et al., "The algorithmic foundations of differential privacy." Found. Trends Theor. Comput. Sci., vol. 9, no. 3-4, pp. 211– 407, 2014.