

Differential Privacy for Stackelberg Games

Ferdinando Fioretto^{1*}, Lesia Mitridati² and Pascal Van Hentenryck²

¹Syracuse University

²Georgia Institute of Technology

ffiorett@syr.edu, lmitridati3@gatech.edu, pvh@isye.gatech.edu

Abstract

This paper introduces a *differentially private* (DP) mechanism to protect the information exchanged during the coordination of *sequential* and *interdependent* markets. This coordination represents a classic Stackelberg game and relies on the exchange of *sensitive information* between the system agents. The paper is motivated by the observation that the perturbation introduced by traditional DP mechanisms fundamentally changes the underlying optimization problem and even leads to unsatisfiable instances. To remedy such limitation, the paper introduces the *Privacy-Preserving Stackelberg Mechanism* (PPSM), a framework that enforces the notions of *feasibility* and *fidelity* (i.e. near-optimality) of the privacy-preserving information to the original problem objective. PPSM complies with the notion of differential privacy and ensures that the outcomes of the privacy-preserving coordination mechanism are close-to-optimality for each agent. Experimental results on several gas and electricity market benchmarks based on a real case study demonstrate the effectiveness of the proposed approach. A full version of this paper [Fioretto *et al.*, 2020b] contains complete proofs and additional discussion on the motivating application.

1 Introduction

In the context of the liberalization of energy markets, the coordination of *sequential* and *interdependent* agents, such as gas and electricity market operators, has become central to achieve an efficient and sustainable operation of the energy system [Ordoudis, 2018]. Such coordination problems have traditionally been modeled as *Stackelberg games* [Simaan and Cruz, 1973], which require the exchange of proprietary information between the agents in order to achieve an optimal strategy. For instance, in the context of electricity and gas markets, relevant data may represent the costs of producers, the loads of consumers, or technical characteristics of the energy network. As has been observed in various works (e.g.,

[Zugno *et al.*, 2013; Baringo and Conejo, 2013]), such information may be sensitive: It can provide a competitive advantage over other strategic agents in the system, it may induce financial losses, and it may even benefit external attackers [Maharjan *et al.*, 2013].

To address this issue, several privacy-preserving frameworks have been proposed, with *Differential Privacy* (DP) [Dwork *et al.*, 2006] emerging as a robust privacy framework for many applications. DP allows to measure and bound the risk associated with an individual participation in an analysis task. DP algorithms rely on the injection of carefully calibrated noise to the output of a computation. They can thus be used to *obfuscate* the sensitive data exchanged by the system agents in the market. However, as shown in Section 7, when perturbed data are used as input to *Stackelberg games*, they may produce results that are fundamentally different from those obtained on the original data: They often transform the nature of the underlying optimization problem and even lead to *severe feasibility issues*.

This paper is a first step in addressing this challenge. It introduces the *Privacy-Preserving Stackelberg Mechanism* (PPSM) for the coordination of sequential and interdependent agents. PPSM is a two-stage protocol that allows the coordinating agents to exchange differentially private data of high fidelity. In particular, PPSM relies on an optimization-based fidelity phase (including *fidelity constraints* on the objectives and coordination variables of the agents) to redistribute the noise introduced on the privacy-preserving data exchanged between the agents, and ensure that it preserves the feasibility and near-optimality of the original Stackelberg game. PPSM has been analyzed both theoretically and experimentally. The theoretical guarantees ensure differential privacy and near optimality, while the experimental results validate the approach on a real test case for the coordination of electricity and natural gas markets in the Northeastern United States [Byeon and Van Hentenryck, 2019]. The case study shows that PPSM can bring up to two orders of magnitude error reduction over standard privacy-preserving mechanisms.

Although the paper was motivated by the coordination between natural gas and electricity markets, the proposed methods may apply to any type of coordination mechanism between sequential and interdependent agents where the agents exchange information and synchronize through price signals.

*Authors names listed alphabetically. All authors have equal contributions.




Figure 1: Stackelberg game in sequential interdependent markets.

2 Problem Definition and Privacy Goal

The strategies of two sequential and interdependent agents, such as energy market operators, represent a classic *Stackelberg game* [Simaan and Cruz, 1973]. In this framework, which is schematically illustrated in Figure 1, the *leader* (e.g., the first market operator) optimizes its decisions while anticipating the reaction of the *follower* (e.g., the second market operator). The leader actions impact the reaction of the follower, which in turn impacts the leader objective value. As a result, the leader strategy in Stackelberg games can be modeled as a bilevel optimization problem $\mathcal{P}^L(D^L, D_P^F, D_S^F)$:

$$\mathcal{O}^{L*} = \min_{\mathbf{x}^L, \mathbf{y}^F} \mathcal{O}^L(\mathbf{x}^L, D^L) \quad (1a)$$

$$\text{s.t. } (\mathbf{x}^L, \mathbf{y}^F) \in \mathcal{F}^L(D^L) \quad (1b)$$

$$\mathbf{y}^F = \text{dual sol. } \min_{\mathbf{x}^F} \mathcal{O}^F(\mathbf{x}^F, D_P^F, D_S^F) \quad (1c)$$

$$\text{s.t. } (\mathbf{x}^F, \mathbf{x}^L) \in \mathcal{F}^F(D_P^F, D_S^F), \quad (1d)$$

where \mathbf{x}^L represents the vector of decision variables of the leader, and \mathbf{x}^F and \mathbf{y}^F the vectors of *primal* and *dual* variables of the follower. Additionally, D^L and (D_P^F, D_S^F) are the inputs of the leader and follower problems, respectively. The follower inputs are either *public* (D_P^F) or *sensitive* (D_S^F). The upper-level problem minimizes the leader objective cost $\mathcal{O}^L(\mathbf{x}^L, D^L)$ (1a), constrained by its feasible decision space $\mathcal{F}^L(D^L)$ (1b), and the reaction of the follower in the lower-level problem (1c) and (1d). The follower problem, denoted by $\mathcal{P}^F(\mathbf{x}^L, D_P^F, D_S^F)$, minimizes the follower objective cost $\mathcal{O}^F(\mathbf{x}^F, D_P^F, D_S^F)$ in (1c), constrained by the feasible space $\mathcal{F}^F(D_P^F, D_S^F)$ of the follower decisions (1d).

Coordination Variables. The leader primal variables \mathbf{x}^L , appear as fixed parameters in the expression of the follower feasible space \mathcal{F}^F . In return, the lower-level problem provides feedback from the follower *dual variables* \mathbf{y}^F , to the upper-level problem through its feasible space \mathcal{F}^L . In energy markets, primal variables typically represent commitment and energy production decisions, while dual variables represent energy prices. These variables shared between the follower and the leader are called *coordination variables*.

Assumptions. Due to the sequential decision-making nature, the leader needs to anticipate the reaction of the follower. Therefore, this paper assumes that the leader has access to a prediction model $\mathcal{M}^L(D^L, D_P^F, D_S^F)$ that predicts the values $\bar{\mathbf{y}}^F$ of the follower dual variables. This is a natural assumption in energy markets applications that motivates this paper: Such forecasting models are used in practice since generators must predict energy prices in order to efficiently bid in the markets

and the market needs to ensure reliability of the overall system. Similarly, the follower has access to a forecasting model $\mathcal{M}^F(\mathbf{x}^L, D_P^F, D_S^F)$ that predicts its objective value $\bar{\mathcal{O}}^{F*}$ and the values of its dual variables $\bar{\mathbf{y}}^{F*}$. This is also a natural assumption in energy markets since the energy prices and costs, representing the dual variables values and objective cost of the follower problem $\mathcal{P}^F(\mathbf{x}^L, D_P^F, D_S^F)$, are public and can be used to train precise estimators. The PPSM mechanism in this paper applies these forecasting models on privacy-preserving versions of the sensitive parameters.

Motivation Problem. Byeon and Van Hentenryck [2019] recently showed that the coordination between electricity and natural gas markets can be modeled as a Stackelberg game between a leader, i.e. the *gas-aware electricity unit commitment* (GAUC), and two followers, i.e. the *electricity market* (EM) and *natural gas market* (GM). This game can alleviate the reliability issues that emerged in the recent polar vortex events. In this context, the leader coordination variables represent the *commitment* of Gas-Fired Power Plants (GFPPs), which impacts their participation in both EM and GM. The relevant follower coordination variables represent natural gas *prices* in the GM. These prices impact the GAUC decisions through *coordination constraints* representing the profitability of the bids of GFPPs.

Privacy Goal. The paper focuses on situations where the follower inputs D_S^F contain sensitive information that should not be revealed. In the case of electricity and natural gas markets, a FERC directive allows the gas and electricity operators to share network data, while the bids of the generators are public information. Therefore, the sensitive parameters typically represent the gas demand profile of consumers. As discussed in the introduction, if released, they can provide a competitive advantage to strategic agents in the energy system and may result in financial losses for the follower, as shown in [Zugno *et al.*, 2013; Baringo and Conejo, 2013]. Thus, the *privacy goal* is to ensure that the sensitive information D_S^F is not breached during the coordination process described in Problem (1). The next section introduces a formal notion that will be used to achieve this goal.

3 Background: Differential Privacy

Differential privacy [Dwork *et al.*, 2006] is a privacy notion used to protect disclosures of an individual's data in a computation. The paper considers datasets $D_S^F = (d_1, \dots, d_n)$ with each $d_i \in \mathbb{R}_+$ describing a sensitive quantity, such as the participants' demand value in the GM. DP relies on the notion of dataset adjacency, which captures the differential information to be protected. To protect the values in the dataset, the paper uses the following *adjacency relation*:

$$D \sim_\alpha D' \Leftrightarrow \exists i : |d_i - d'_i| \leq \alpha \wedge \forall j \neq i : d_j = d'_j,$$

where D and D' are two datasets and α is a positive real value. Such definition is useful to hide individual participation up to some quantity α [Chatzikokolakis *et al.*, 2013]. In the paper context, it allows customers to reveal gas demand profiles that hide the real consumption by a factor of α .

A randomized mechanism $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{R}$ with domain \mathcal{D} and range \mathcal{R} is ϵ -DP if, for any output response $O \subseteq \mathcal{R}$ and any

two adjacent datasets $D \sim_{\alpha} D'$, for a fixed value $\alpha > 0$:

$$Pr[\mathcal{M}(D) \in O] \leq \exp(\epsilon) Pr[\mathcal{M}(D') \in O]. \quad (2)$$

The parameter $\epsilon \geq 0$ is the *privacy loss* of the mechanism, with values close to 0 denoting strong privacy. The level of *indistinguishability* is controlled by the parameter $\alpha > 0$. The definition above was introduced by Chatzikokolakis *et al.* [2013]. It is a *generalization* of the classical DP definition, that protects individuals participation, to arbitrary metric spaces. W.l.o.g. the paper fixes $\epsilon = 1.0$ and focuses in the indistinguishability parameter α . In the context of this work, Definition (2) is referred to as α -indistinguishability.

An important DP property is immunity to post-processing.

Theorem 1 (Post-Processing [Dwork and Roth, 2013])

Let \mathcal{M} be an α -indistinguishable mechanism and g be an arbitrary mapping from the set of possible outputs to an arbitrary set. Then, $g \circ \mathcal{M}$ is α -indistinguishable.

A function Q (also called *query*) from a data set $D \in \mathcal{D}$ to a result set $R \subseteq \mathbb{R}^n$ can be made differentially private by injecting random noise to its output. The amount of noise depends on the *sensitivity* of the query, denoted by Δ_Q and defined as $\Delta_Q = \max_{D \sim D'} \|Q(D) - Q(D')\|_1$. In other words, the sensitivity of a query is the maximum l_1 -distance between the query outputs of any two adjacent datasets D and D' .

4 Privacy-Preserving Stackelberg Problem

The *Privacy-Preserving Stackelberg* (PPS) problem establishes the fundamental desiderata to be delivered by the obfuscation mechanism. It operates on the follower sensitive parameters D_S^F exchanged to ensure coordination between the leader and the follower in the resolution of Problem (1). The goal is to produce a privacy-preserving version \hat{D}_S^F that is α -indistinguishable from D_S^F and which ensures that $\hat{O}^{L*} \approx O^{L*}$, where \hat{O}^{L*} is the leader objective value in the Stackelberg game (1) when the sensitive data D_S^F is replaced by its privacy-preserving version \hat{D}_S^F .

5 The PPSM Mechanism

The *Privacy-Preserving Stackelberg Mechanism* (PPSM) is described schematically in Figure 2 and consists of the following steps which will be described in detail subsequently:

- **[1] [Follower]:** Given the sensitive data D_S^F , the follower produces the obfuscated data \tilde{D}_S^F which is α -indistinguishable from D_S^F using the Laplace mechanism.
- **[2a] [Leader]:** Given the obfuscated data (\tilde{D}_S^F) computed in [1] and the public data (D^L, D_P^F), the leader uses model \mathcal{M}^L to (privately) estimate the value of the dual coordination variables \bar{y}^F .
- **[2b] [Leader]:** It then solves the leader problem \mathcal{P}^L (1a)-(1b) with the values of the follower variables \mathbf{y}^F fixed to the estimates \bar{y}^F to obtain the estimates of the coordination variables \bar{x}^L .
- **[3a] [Follower]:** Given estimates \bar{x}^L and the obfuscated data (D_P^F, \tilde{D}_S^F), the follower uses model \mathcal{M}^F to (privately) estimate the objective value \bar{O}^{F*} and the values for the dual variables \bar{y}^{F*} of the follower problem \mathcal{P}^F (1c) and (1d).

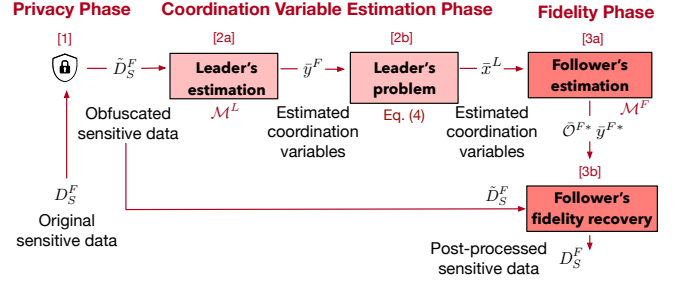


Figure 2: PPSM Illustration.

- **[3b] [Follower]:** Finally, using \tilde{D}_S^F and the estimates computed in [2b] and [3a], the follower produces a new privacy-preserving vector \hat{D}_S^F to achieve fidelity, i.e. near-optimality of the leader problem, such that $\hat{O}^{L*} \approx O^{L*}$. Once PPSM produces the privacy-preserving demand \hat{D}_S^F , the leader can solve its problem $\mathcal{P}^L(D^L, D_P^F, \hat{D}_S^F)$ (1a)–(1d) to produce x^{L*} which is communicated to the follower for solving its own problem $\mathcal{P}^F(x^{L*}, D_P^F, \hat{D}_S^F)$ (1c) and (1d), as depicted in Figure 1.

5.1 Privacy Phase

In Step [1], the follower takes as input its sensitive parameters D_S^F and constructs a privacy-preserving version \tilde{D}_S^F using the Laplace mechanism [Dwork *et al.*, 2006].

Theorem 2 (Laplace Mechanism) Let Q be a numeric query that maps datasets to \mathbb{R}^n . The Laplace mechanism that outputs $Q(D) + \xi$, where $\xi \in \mathbb{R}^n$ is drawn from the Laplace distribution $\text{Lap}(\Delta_Q/\epsilon)^n$, achieves α -indistinguishability.

In the above, $\text{Lap}(\lambda)^n$ denotes the i.i.d. Laplace distribution with 0 mean and scale λ over n dimensions. As a result, the privacy-preserving parameters \tilde{D}_S^F are obtained as follows:

$$\tilde{D}_S^F = D_S^F + \text{Lap}(\alpha)^n. \quad (3)$$

Importantly, the Laplace mechanism has been shown to be *optimal*, i.e., it minimizes the mean-squared error for identity queries w.r.t. the L1-norm [Koufogiannis *et al.*, 2015].

While (3) ensures α -indistinguishability, the obfuscated data may not achieve strong *fidelity* w.r.t. the original problem. Crucially, in energy systems, the inputs generated by this mechanism often fail to produce a feasible solution to the problem of interest, as highlighted in Section 7. To remedy this limitation, the proposed PPSM introduces an *optimization-based* approach that aims at producing a new privacy-preserving dataset \hat{D}_S^F establishing *feasibility* and *fidelity* (i.e. near-optimality) w.r.t. the constraints and objectives of the leader and follower problems.

5.2 Estimating the Coordination Variables

After having received the Laplace-obfuscated data \tilde{D}_S^F from the follower, the leader uses model $\mathcal{M}^L(D^L, D_P^F, \tilde{D}_S^F)$ to estimate the value of the coordination variables \mathbf{y}^F of the follower (Step [2a]). Next, in Step [2b], the leader solves the optimization problem

$$\bar{x}^L = \arg \min_{x^L} O^L(x^L, D^L) \quad \text{s.t.} \quad (x^L, \bar{y}^F) \in \mathcal{F}^L(D^L). \quad (4)$$

Problem (4) describes, in fact, the leader *subproblem* (Equations (1a) and (1b)), in which the values of variables \mathbf{y}^F have been fixed to the estimates \bar{y}^F .

The leader then communicates the estimates \bar{x}^L to the follower. In turn, the follower uses model $\mathcal{M}^F(\bar{x}^L, D_P^F, \hat{D}_S^F)$ to estimate the values of its subproblem objective value $\bar{\mathcal{O}}^{F*}$ and dual variables \bar{y}^{F*} (Step [3a]).

5.3 Fidelity Phase

Finally, given the obfuscated parameters \tilde{D}_S^F , computed in Step [1], the estimated values \bar{x}^L , computed in Step [2b], and the follower objective value $\bar{\mathcal{O}}^{F*}$ and dual variables \bar{y}^{F*} , computed in Step [3a], the follower executes the following bilevel optimization problem:

$$\min_{\hat{D}_S^F, \hat{\mathbf{x}}^F, \hat{\mathbf{y}}^F} \|\hat{D}_S^F - \tilde{D}_S^F\|_2^2 \quad (5a)$$

$$\text{s.t. } |\hat{\mathcal{O}}^F - \bar{\mathcal{O}}^{F*}| \leq \eta_p \quad (5b)$$

$$|\hat{\mathbf{y}}^F - \bar{y}^{F*}| \leq \eta_d \quad (5c)$$

$$\hat{\mathbf{y}}^F = \text{dual sol. of } \mathcal{P}^F(\bar{x}^L, D_P^F, \hat{D}_S^F), \quad (5d)$$

where η_p and η_d are parameters specifying the desired fidelity for the value of the objective and dual variables of the follower. Its objective is to find new values \hat{D}_S^F that minimize the distance to the Laplace-obfuscated \tilde{D}_S^F (5a), while ensuring (component-wise) fidelity w.r.t. the estimated objective value $\bar{\mathcal{O}}^{F*}$ (5b) and dual variables \bar{y}^{F*} (5c). The follower objective function $\hat{\mathcal{O}}^F$ and dual variables $\hat{\mathbf{y}}^F$ are defined as the solutions to the lower-level problem (5d), which represents the follower subproblem $\mathcal{P}^F(\bar{x}^L, D_P^F, \hat{D}_S^F)$ with the new privacy-preserving parameters \hat{D}_S^F and the values of the coordination variables \mathbf{x}^L fixed to the estimates \bar{x}^L . Additionally, this lower-level problem (5d) enforces feasibility of the follower problem w.r.t. the estimates \bar{x}^L . Using the equivalent Karush-Kuhn-Tucker (KKT) conditions of the linear lower-level problem (5d) and the Fortuny-Amat linearization, this bilevel problem can be recast as a mixed-integer second-order cone program (MISOCP) [Gabriel *et al.*, 2012].

The obfuscated data \hat{D}_S^F returned by the PPSM mechanism is α -indistinguishable from D_S^F , since Steps [2a]–[3b] can be seen as a post-processing transformation. Indeed, they use exclusively public information and data-independent models.

6 Error Analysis

This section analyzes the impact of the data perturbation induced by PPSM on the optimal objective value of both agents in the privacy-preserving problem. The results below hold under the assumption that the estimated values \bar{x}^L , $\bar{\mathcal{O}}^{F*}$, and \bar{y}^{F*} are accurate.

Theorem 3 (Error) *After the fidelity phase, the expected error induced by PPSM on the original, sensitive, parameters D_S^F is bounded by the inequality:*

$$\mathbb{E}[\|\hat{D}_S^{F*} - D_S^F\|] \leq 4\alpha^2,$$

where \hat{D}_S^{F*} is the solution to Problem (5).

Proof Sketch. The proof relates the distance between \hat{D}_S^{F*} and D_S^F to that between the noisy \tilde{D}_S^F and D_S^F , and relies on triangular inequality on norms, optimality of \hat{D}_S^{F*} , and the fact that the Laplace mechanism is an unbiased estimator. \square

The next theorem bounds the difference in objective value for the leader problem when the leader problem is convex and the coordination constraints are linear.

Theorem 4 (Cost of privacy) *Consider a convex leader problem whose coordination constraints (6c) are linear;*

$$\hat{\mathcal{O}}^{L*} = \min_{\mathbf{x}^L, \mathbf{y}^F} \mathcal{O}^L(\mathbf{x}^L, D^L) \quad (6a)$$

$$\text{s.t. } \mathbf{x}^L \in \mathcal{F}^L(D^L) \quad (6b)$$

$$A\mathbf{x}^L + B\mathbf{y}^F \geq b \quad (6c)$$

$$\mathbf{y}^F = \text{dual sol. of } \mathcal{P}^F(\mathbf{x}^L, D_P^F, \hat{D}_S^F), \quad (6d)$$

where $\mathcal{P}^F(\mathbf{x}^L, D_P^F, \hat{D}_S^F)$ uses the privacy-preserving \hat{D}_S^F . After the fidelity phase, the error induced by PPSM on the objective value $\hat{\mathcal{O}}^{L*}$ of the leader problem (6) is bounded by:

$$\|\hat{\mathcal{O}}^{L*} - \mathcal{O}^{L*}\| \leq \eta_d \|B^T \mathbf{y}^{L*}\|_1, \quad (7)$$

where \mathbf{y}^{L*} are the dual variables values associated with the constraints (6c) of the original leader problem \mathcal{P}^L .

Proof Sketch. The proof relies on the best- and worst-case counterparts of the linear coordination constraints (6c). The fidelity criteria (5c) can be reformulated as a perturbation on the dual variables $\mathbf{y}^F = \bar{y}^{F*} + \eta_d \boldsymbol{\epsilon}$, where each component of the random vector $\boldsymbol{\epsilon}$ is such that $|\epsilon_i| \leq 1$. Therefore, the difference $(\hat{\mathcal{O}}^{L*} - \mathcal{O}^{L*})$ can be upper- and lower-bounded by the objective value of the *perturbed* leader problem, in which the right-hand sides of the coordination constraints (6c) are perturbed by the small quantity $\pm \eta_d \|B^T\|_1$. \square

While fidelity w.r.t. the follower objective value is *explicitly* enforced by Constraint (5b), the result above ensures fidelity w.r.t. the leader objective value. This fidelity is *implicitly* enforced by Constraint (5c) on the follower coordination variables \mathbf{y}^F , and their impact on the leader objective value via the coordination constraints (6c). This sub-optimality in the leader cost represents the so-called *cost of privacy*.

7 Experimental Evaluation

The performance of the proposed PPSM is illustrated on the motivation problem introduced in Section 2. The leader represents the GAUC problem, and the two followers represent the EM and GM problems. The PPSM aims at preserving privacy on the sensitive gas demand profiles (D_S^g) in the GM, and fidelity w.r.t. the original Stackelberg game. Fidelity constraints are *explicitly* enforced on the original objective value of the GM (\mathcal{O}^{g*}) and gas prices (\mathbf{y}^{g*}). Additionally, the *implicit* impact of the PPSM on the original objective values of the GAUC (\mathcal{O}^{uc*}) and the EM (\mathcal{O}^{e*}) is analyzed.

Case Study Setup. The PPSM is evaluated on a test system representing the joint natural gas and electricity systems in the Northeastern US [Byeon and Van Hentenryck, 2019]. The system is composed of the IEEE 36-bus power system

\mathcal{M}	α	sat.(%)	$\Delta_{D_S^g}$ (L1)	$\Delta_{\mathcal{O}^{uc}}$ (%)	$\Delta_{\mathcal{O}^e}$ (%)	$\Delta_{\mathcal{O}^g}$ (%)
Laplace	0.1	71.49	5.08	0.1237	0.3222	0.3335
	1.0	18.13	50.85	1.2959	3.5538	3.5540
	10.0	4.47	508.55	22.940	52.414	52.414
PPSM _p	0.1	98.45	4.45	0.0631	0.1503	0.1503
	1.0	91.30	21.87	0.1216	0.1764	0.1761
	10.0	80.71	24.31	0.2143	0.3851	0.3853
PPSM	0.1	99.10	3.89	0.0192	0.1056	0.1057
	1.0	95.09	12.71	0.0698	0.1465	0.1465
	10.0	91.35	14.16	0.1367	0.2330	0.2331

Table 1: Left: Satisfactory instances (%) and L1 errors (MWh) on the gas demands ($\Delta_{D_S^g}$) for varying indistinguishability parameters α , and $\eta_p = \eta_d = 0.1\%$ of the leaders objective value and the dual variables values, respectively. Right: Errors (%) on the leader objective ($\Delta_{\mathcal{O}^{uc}}$) and followers’ objectives ($\Delta_{\mathcal{O}^e}$ and $\Delta_{\mathcal{O}^g}$).

[Allen *et al.*, 2008] and a gas transmission network covering the Pennsylvania-to-northeast New England area.

This case study analyzes the performance of PPSM under various operating conditions of the gas and electricity systems. The electricity demand profile is uniformly increased by a stress factor ranging from 30% to 60%, and the gas demand profile is increased by a stress factor ranging from 10% to 130%, producing increasingly stressed and difficult operating conditions. The experiments compare the proposed PPSM to a version (PPSM_p) that omits the fidelity constraint on the dual variables (5c). Both versions are compared with the standard Laplace mechanism for varying values of the privacy parameter $\alpha \in \{0.1, 1, 10\} \times 10^2$ MWh, and the fidelity parameters $\eta_p = \eta_d \in \{0.01, 0.1, 10.0\}\%$ of the original objective value of the GM (\mathcal{O}^{g*}) and gas prices (\mathbf{y}^{g*}), respectively. In the GAUC problem, the original, sensitive, gas demand vector is denoted D_S^g (in lieu of D_S^F). Notice that, in the highest stress factor adopted, this demand vector has minimum, median, and maximum values: 0, 3.38×10^2 , 98.31×10^2 , respectively. Therefore, the privacy parameters adopted ensure a very low privacy risk.

The uncertainty resulting by predicting the gas cost estimates \bar{y}^g (in lieu of \bar{y}^F) is simulated by using a noisy version obtained by perturbing the original quantities \mathbf{y}^{g*} using Normal noise with standard deviation of 10% the average value of \mathbf{y}^{g*} . The experiments also evaluated benchmarks that use precise cost estimates and present consistent trends.

We generate 30 repetitions for each test case and report average results in all experiments and impose a 1-hour wall-clock limit. A wall-clock limit of 1 hour is given to generate and solve each of the instances. The resolution of the privacy-preserving demand profiles (phases [1] and [2] of PPSM) takes less than 30s for any of the instances.

Limits of the Laplace Mechanism

This section studies the applicability of the Laplace mechanism to our context of interest. Table 1 (left) reports the percentage of the feasible solutions (over 1260 instances) across different values of the privacy parameter α . It compares the Laplace mechanism with PPSM_p and PPSM. When $\alpha > 0.1$ the Laplace-obfuscated gas demands rarely produce a feasible solution to the GAUC problem. *These results justify the need for more advanced privacy-preserving mechanisms for Stackelberg games, and hence the proposed PPSM.* In contrast, the PPSMs result in a much higher number of feasible solutions.

\mathcal{M}	η	$\Delta_{\mathcal{O}^{uc}}$ (%)	$\Delta_{\mathcal{O}^e}$ (%)	$\Delta_{\mathcal{O}^g}$ (%)
Laplace	NA	22.9400	52.4141	52.4141
	0.1%	0.1915	0.3851	0.3851
	1.0%	0.1946	0.4102	0.4102
PPSM _p	10.0%	0.2224	0.4543	0.4543
	0.1%	0.1367	0.2330	0.2330
	1.0%	0.1242	0.2060	0.2060
PPSM	10.0%	0.2086	0.3601	0.3605

Table 2: Cost objective differences (%) at varying fidelity parameters $\eta = \eta_p = \eta_d$ %, and indistinguishability parameter $\alpha = 10$.

Indeed, all “unsolved” PPSM cases are due to timeout and not as a direct result of infeasibility. Additionally, we verified that the two PPSMs are always able to find a feasible solution to the GM follower problem with the gas demand profile \hat{D}_S^g .

Table 1 (left) also reports the L1 distance Δ_d between the original gas demand D_S^g and the privacy-preserving versions obtained with each of the mechanisms analyzed. Unsurprisingly, the L1 errors increase with the increasing of parameter α , as larger values of α induce more noise. However, the L1 errors introduced by the PPSM are much more contained than the Laplace ones, producing more than an order of magnitude more accurate results for the larger privacy parameters. *These results indicate that the highly-perturbed demand profiles induced by the Laplace mechanism lead to infeasibility in the GAUC problem, whereas both PPSM and PPSM_p manage to restore feasibility of the post-processed demand profiles.*

Leader and Follower Objectives

The next results evaluate the ability of PPSM to preserve the optimal objective values of the leader and the follower problems. Table 1 (right) tabulates the errors, in percentage, on the objective costs of the GAUC problem (leader), the EM problem and the GM problem (followers) at varying indistinguishability parameters α , and for a fixed fidelity parameter $\eta = \eta_p = \eta_d = 0.1\%$. The errors $\Delta_{\mathcal{O}}$ are defined as $\frac{|\mathcal{O}^* - \bar{\mathcal{O}}^*|}{\mathcal{O}^*} \%$, where $\mathcal{O} \in \{\mathcal{O}^{uc}, \mathcal{O}^e, \mathcal{O}^g\}$ and are computed in expectation over the feasible instances only. Parameter α controls the amount of noise being added to the gas demand profiles, therefore, the objective costs are closer to their original values when α is small. *Observe that the PPSMs induce objective costs differences that are from one to two orders of magnitude more accurate than those induced by the Laplace mechanism, and that are at most 0.4% of the original objective costs.* Additionally, PPSM is consistently more accurate than PPSM_p; By enforcing fidelity of the coordination variables \mathbf{y}^g , PPSM better limits the impact of the noise on the leader objective (GAUC), which in turn results in more faithful solutions for the followers’ problems.

These results are further illustrated in Table 2, that analyzes the difference in objective costs at varying fidelity parameters η_p and η_d , for a fixed privacy parameter $\alpha = 10$ (i.e., the largest privacy level attainable in our experimental setting). Once again, the results of the PPSM mechanisms are at least two orders of magnitude more precise than those obtained by the Laplace mechanism. Additionally, notice that the fidelity parameters η_p and η_d impact the accuracy of the privacy-preserving objective costs. Indeed, they indirectly




Figure 3: GAUC objective cost difference (%) at varying gas (g) and electricity (e) stress levels for \tilde{D}_S^g obtained via *Laplace* (top), *PPSM_p* (middle), *PPSM* (bottom). Privacy: $\alpha = 0.1$ (left), 1.0 (center), 10.0 (right). Fidelity: $\eta_p = \eta_d = 0.1\%$ of respective quantities.

control the deviation of the privacy-preserving GAUC and GM objectives w.r.t. the original counterparts. While the results differences are small, in percentage, their impact on the objective functions (in the 10^6 order) is non-negligible.

Stress Levels Analysis

Figure 3 reports heatmaps of the total (GAUC) objective cost difference, in percentage, at varying electricity (e) and gas (g) stress levels for the privacy-preserving data obtained via the Laplace mechanism (top), *PPSM_p* (middle), and *PPSM* (bottom). Each square represents the objective cost difference averaged over 30 instances for a particular electricity and gas stress level. The darker the color, the more pronounced are the errors committed by the mechanisms, as reported in the legends on top of each subfigure. Gray squares represent the set of instances for which no feasible solution of the GAUC problem was found or when a timeout is reached. The illustration reports the cost differences for privacy parameters $\alpha = 0.1$ (left) $\alpha = 1.0$ (middle), and $\alpha = 10.0$ (right).

These results illustrate three trends: Firstly, for every mechanism, the objective differences become more pronounced as the electricity and gas stress levels increase. This can be explained by the increased impact of the Laplace perturbations on higher values of gas demand profiles. Secondly, they remark that the *PPSMs* produce privacy-preserving Stackelberg problems that are consistently more faithful to the original problems w.r.t. those produced by the Laplace mechanism. Finally, they show that *PPSM* is consistently more accurate than *PPSM_p* across all stress levels. *These results are significant, as they show the robustness of the proposed PPSMs over different electricity and natural gas demand profiles. They indicate that PPSM can provide a realistic and efficient solution for the coordination of these markets.*

8 Related Work

The obfuscation of data values under the lens of differential privacy is a challenging task that has been studied from several angles. Often, the released data is generated from a data synopsis in the form of a noisy histogram [Li *et al.*, 2010; Hay *et al.*, 2016; Qardaji *et al.*, 2014; Xiao *et al.*, 2010]. These methods are typically adopted in the context of statistical queries. The design of markets for private data has also received considerable attention, see for instance [Ghosh and

Roth, 2010; Niu *et al.*, 2018].

However, all the proposals above, do not involve data used as input to a complex optimization problem, as in the case of this work. The closest work related to the proposal in this paper can be considered [Fioretto and Van Hentenryck, 2018; Fioretto *et al.*, 2020a], which, in the context energy networks, propose a privacy-preserving mechanism for releasing datasets that can be used as input to an *optimal power flow* problem. A similar line of work uses hierarchical (bilevel) optimization for obfuscating the energy network parameters or locations while ensuring high utility for the problem of interest [Fioretto *et al.*, 2019; Mak *et al.*, 2020].

In contrast to these studies, the proposed *PPSM* focuses on solving Stackelberg games in which the follower parameters are sensitive. *PPSM* also enforces the notion of fidelity of the privacy-preserving information w.r.t. the leader and follower objectives. Finally, to the best of our knowledge, this is the first DP mechanism that is applied to the coordination of sequential electricity and natural gas markets.

9 Conclusions

This paper introduced a differentially private (DP) mechanism to protect the *sensitive information* exchanged during the coordination of the sequential electricity and natural gas market clearings. The *proposed Privacy-Preserving Stackelberg Mechanism (PPSM)* obfuscates the gas demand profile exchanged by the gas market, while also ensuring that the resulting problem preserves the fundamental properties of the original Stackelberg game. The *PPSM* was shown to enjoy strong properties: It complies with the notion of DP and ensures that the outcomes of the privacy-preserving Stackelberg mechanism are close-to-optimality for each agent. Experimental results on several gas and electricity market benchmarks based on a real case study demonstrated the effectiveness of the approach: *The PPSM was shown to obtain up to two orders of magnitude improvement on the cost of the agents when compared to a traditional Laplace mechanism.*

Future work will focus on several avenues, including extended theoretical bounds on the cost of privacy, studying the game-theoretic properties of this privacy-preserving Stackelberg game, accounting for uncertainty on the public data, and, studying the applicability of the *PPSM* to other domains.

References

- [Allen *et al.*, 2008] Eric H Allen, Jeffrey H Lang, and Marija D Ilic. A combined equivalenced-electric, economic, and market representation of the northeastern power coordinating council us electric power system. *IEEE Transactions on Power Systems*, 23(3):896–907, 2008.
- [Baringo and Conejo, 2013] Luis Baringo and Antonio J Conejo. Strategic offering for a wind power producer. *IEEE Transactions on Power Systems*, 28(4):4645–4654, 2013.
- [Byeon and Van Hentenryck, 2019] Geunyeong Byeon and Pascal Van Hentenryck. Unit commitment with gas network awareness. *arXiv preprint arXiv:1902.03236*, 2019.
- [Chatzikokolakis *et al.*, 2013] Konstantinos Chatzikokolakis, Miguel E Andrés, Nicolás Emilio Bordenabe, and Catuscia Palamidessi. Broadening the scope of differential privacy using metrics. In *International Symposium on Privacy Enhancing Technologies Symposium*, pages 82–102. Springer, 2013.
- [Dwork and Roth, 2013] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Theoretical Computer Science*, 9(3-4):211–407, 2013.
- [Dwork *et al.*, 2006] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference (TCC)*, volume 3876, pages 265–284. Springer, 2006.
- [Fioretto and Van Hentenryck, 2018] Ferdinando Fioretto and Pascal Van Hentenryck. Constrained-based differential privacy: Releasing optimal power flow benchmarks privately. In *Proceedings of Integration of Constraint Programming, Artificial Intelligence, and Operations Research (CPAIOR)*, pages 215–231, 2018.
- [Fioretto *et al.*, 2019] Ferdinando Fioretto, Terrence W. K. Mak, and Pascal Van Hentenryck. Privacy-preserving obfuscation of critical infrastructure networks. pages 1086–1092, 2019.
- [Fioretto *et al.*, 2020a] F. Fioretto, T.W.K. Mak, and P. Van Hentenryck. Differential privacy for power grid obfuscation. *IEEE Transactions on Smart Grid*, 11(2):1356–1366, 2020.
- [Fioretto *et al.*, 2020b] Ferdinando Fioretto, Lesia Mitridati, and Pascal Van Hentenryck. Differential privacy for stackelberg games, 2020.
- [Gabriel *et al.*, 2012] Steven A Gabriel, Antonio J Conejo, J David Fuller, Benjamin F Hobbs, and Carlos Ruiz. *Complementarity modeling in energy markets*, volume 180. Springer Science & Business Media, 2012.
- [Ghosh and Roth, 2010] Arpita Ghosh and Aaron Roth. Selling privacy at auction, 2010.
- [Hay *et al.*, 2016] Michael Hay, Ashwin Machanavajjhala, Gerome Miklau, Yan Chen, and Dan Zhang. Principled evaluation of differentially private algorithms using dp-bench. In *Proceedings of the 2016 International Conference on Management of Data*, pages 139–154. ACM, 2016.
- [Koufogiannis *et al.*, 2015] Fragkiskos Koufogiannis, Shuo Han, and George J Pappas. Optimality of the laplace mechanism in differential privacy. *arXiv preprint arXiv:1504.00065*, 2015.
- [Li *et al.*, 2010] Chao Li, Michael Hay, Vibhor Rastogi, Gerome Miklau, and Andrew McGregor. Optimizing linear counting queries under differential privacy. In *Proceedings of the twenty-ninth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 123–134. ACM, 2010.
- [Maharjan *et al.*, 2013] Sabita Maharjan, Quanyan Zhu, Yan Zhang, Stein Gjessing, and Tamer Basar. Dependable demand response management in the smart grid: A stackelberg game approach. *IEEE Transactions on Smart Grid*, 4(1):120–132, 2013.
- [Mak *et al.*, 2020] Terrence WK Mak, Ferdinando Fioretto, Lyndon Shi, and Pascal Van Hentenryck. Privacy-preserving power system obfuscation: A bilevel optimization approach. *IEEE Transactions on Power Systems*, 35(2):1627–1637, 2020.
- [Niu *et al.*, 2018] Chaoyue Niu, Zhenzhe Zheng, Fan Wu, Shaojie Tang, Xiaofeng Gao, and Guihai Chen. Unlocking the value of privacy: Trading aggregate statistics over private correlated data. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '18*, pages 2031–2040, New York, NY, USA, 2018. ACM.
- [Ordoudis, 2018] Christos Ordoudis. Market-based approaches for the coordinated operation of electricity and natural gas systems. 2018.
- [Qardaji *et al.*, 2014] Wahbeh Qardaji, Weining Yang, and Ninghui Li. Privview: practical differentially private release of marginal contingency tables. In *Proceedings of the 2014 ACM SIGMOD international conference on Management of data*, pages 1435–1446. ACM, 2014.
- [Simaan and Cruz, 1973] Marwaan Simaan and Jose B Cruz. On the stackelberg strategy in nonzero-sum games. *Journal of Optimization Theory and Applications*, 11(5):533–555, 1973.
- [Xiao *et al.*, 2010] Yonghui Xiao, Li Xiong, and Chun Yuan. Differentially private data release through multidimensional partitioning. In *Workshop on Secure Data Management*, pages 150–168. Springer, 2010.
- [Zugno *et al.*, 2013] Marco Zugno, Juan M Morales, Pierre Pinson, and Henrik Madsen. Pool strategy of a price-maker wind power producer. *IEEE Transactions on Power Systems*, 28(3):3440–3450, 2013.