



HAL
open science

Learning Market Equilibria Preserving Statistical Privacy Using Performative Prediction

Hélène Le Cadre, Mandar Datar, Mathis Guckert, Eitan Altman

► **To cite this version:**

Hélène Le Cadre, Mandar Datar, Mathis Guckert, Eitan Altman. Learning Market Equilibria Preserving Statistical Privacy Using Performative Prediction. 2023. hal-04343535v3

HAL Id: hal-04343535

<https://inria.hal.science/hal-04343535v3>

Preprint submitted on 21 Oct 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Learning Market Equilibria Preserving Statistical Privacy Using Performative Prediction

Hélène Le Cadre¹, Mandar Datar², Mathis Guckert¹, and Eitan Altman³, Fellow, IEEE

Abstract—We consider a peer-to-peer electricity market modeled as a private network game, where End Users (EUs) minimize their cost by computing their demand and controllable generation. Their nominal demand constitutes sensitive information, that they might want to keep private. We prove that the private network game admits a unique Variational Equilibrium, which depends on the private information of all the EUs. Thus, to update their strategy, EUs rely on readings from the other EUs. However, to preserve their privacy, the EUs might report randomized readings. A Data Aggregator (DA) is introduced, which aims to learn the EUs’ private information, while remunerating them depending on the quality of their readings. Using performative prediction, we define a decision-dependent game $\mathcal{G}^{\text{stoch}}$, to explicitly take into account the distribution shift caused by the EUs’ hidden ability. We focus on three solution concepts: (\star) Nash Equilibrium (NE), (\star) Performatively Stable Equilibrium (PSE), and (\star) Stackelberg Equilibrium (SE). Relying on the strong monotonicity of the game, we prove that there exist unique NE and PSE solutions of $\mathcal{G}^{\text{stoch}}$. Both PSE and SE account for the feedback-loop effect of the EUs’ hidden ability. The related market robustness can be evaluated by comparing the market efficiency in the PSE and SE with respect to the social optimum. We show that under mild assumptions, the PSE can be found by distributed and sequential variants of the repeated (stochastic) gradient method while we propose a two-timescale stochastic approximation method to learn SE. Finally, we formulate the DA’s optimal contract design as a bilevel optimization problem that we cast as a more tractable nonlinear nonconvex optimization problem which can be solved using simulated annealing. Simulations on small and large scale problem instances illustrate the results.

Index Terms — Learning Game, Statistical Privacy, Performatively Stable Equilibrium, Mechanism Design.

I. INTRODUCTION

In distributed control system literature, there is a usual trade-off between cost minimization and privacy preservation: to obtain a better outcome, each End User (EU) in the system relies on its own information and the information the others have reported. However, the EUs might be reluctant to release sensitive information because of privacy limitations or because, by acting strategically on their readings – therefore providing fake information, they might decrease their cost. The former situation might give rise to moral

hazard, which could generate inefficiencies at the system level. To penalize free-rider behaviors and incentivize the EUs to minimize the variance of their readings, we aim to define contracts which remunerate the EUs as increasing functions of the data quality while discriminating between the EUs depending on their local privacy breach.

A. Problem Statement

In energy systems, data – and the information it contains – is becoming an asset that the agents might want to acquire to learn market equilibria [2], [3]. For example, in many demand-response programs, the aggregator uses the EUs’ historical energy consumption to determine their nominal demand profiles, and design incentives to flatten peak times. The nominal demand profiles are then used to compute incentives depending on how much they curtail relative to their nominal profiles. However, EUs can act strategically on the data they share. For example, energy consumers may artificially inflate their nominal demand prior to a demand-response program event, in order to receive larger payments under the program (see [4] and references therein).

In this paper, we want to quantify the EUs’ trade-off between statistical privacy preservation, i.e., minimum information release, and adjacency, i.e., the inability for the aggregator to determine whether the EUs’ readings are biased, and the efficiency loss for the market while taking into account performative effects of the EUs’ hidden ability.

B. Main Contributions

This paper significantly extends [1], with the goal to formally quantify the robustness of the market towards distributional shifts caused by the EUs’ hidden ability. In contrast, [1] focused on a variational analysis of the private network game and its stochastic counter-part, accounting for the EUs’ strategic information. We now describe the most salient contributions of the paper.

i) We consider a peer-to-peer electricity trading problem formulated as a private network game, where strategic EUs want to minimize their cost by computing their demand (flexible around a nominal value), controllable generation, and trades with their peers, while satisfying a set of local constraints including the local balance of supply and demand, and trading reciprocity constraints that couple the EUs’ optimization problems. Their nominal demand constitutes sensitive information that they might want to keep private. We are closely related to [2], [5], [6], which formulate a peer-to-peer electricity market as a generalized Nash game. However, in this paper, we propose a more sophisticated

¹ H. Le Cadre and M. Guckert are with Inria, Univ. Lille, CNRS, Centrale Lille, UMR 9189 CRISTAL, F-59000, Lille, France helene.lecadre@inria.fr, mathis.guckert@centrale.centralelille.fr

² M. Datar is currently with CEA-Leti, Grenoble, France mandar.datar@cea.fr and was associated with Inria Sophia Antipolis and the University of Avignon, France, during the time of this work

³ E. Altman is with Inria Sophia Antipolis, Univ. Avignon, France eitan.altman@inria.fr

*A previous shorter version of this paper [1] appeared at the IEEE European Control Conference (ECC23).

market model taking into account the feedback-loop effect of the EUs' hidden ability. Following [2], [3] and the electricity market literature, the EUs' utility functions are assumed additive and quadratic. We prove that under mild assumptions, the private network game admits a unique Variational Equilibrium (VE), which is efficient and globally stable. In addition, we prove that, at the VE, the EUs' best responses depend on the private information of *all* the EUs.

ii) A Data Aggregator (DA) might then be introduced, which aims to learn the EUs' private information from the readings it collects, by designing a privacy contract incentivizing the EUs to report their private information truthfully. To assess the performative effects of the EUs' hidden abilities on their demands and controllable generations, we formulate a decision-dependent game $\mathcal{G}^{\text{stoch}}$. To interpret $\mathcal{G}^{\text{stoch}}$ outcome, we focus on three solution concepts: (★) Nash Equilibrium (NE), (★) Performatively Stable Equilibrium (PSE), and (★) Stackelberg Equilibrium (SE). PSE arise naturally when the readings are calibrated not against past outcomes, but against the future outcomes that manifest from acting on the readings [7]. When ignored – as is the case for NE – performativity can surface as a form of distribution shift. In case where EUs react to the demand and controllable generation strategies by adapting their hidden ability with the hope to improve their outcomes, the game has a Stackelberg structure. To quantify the market robustness, we evaluate the efficiency loss in each solution concept w.r.t. the social optimum.

iii) We propose distributed and sequential variants of the Repeated Stochastic Gradient Method (RSGM) with provable convergence [8] to compute PSE, that we test against a classical distributed Gradient Descent Method (GDM) for NE computation. For SE learning, we introduce a Two-Timescale Stochastic Method (TTSM) with provable convergence. Closed-form expressions of each algorithm rate of convergence are derived; insights are given regarding the tuning of the stepsize descents and iteration complexity.

iv) Finally, we formulate the DA's optimal contract design as a bilevel optimization problem that we cast as a more tractable nonlinear nonconvex optimization problem which can be solved relying on simulated annealing. The performance of the algorithms is assessed on small and large scale instances, by comparing their convergence rates and the related trade-off between efficiency and information release.

C. Reviewed Literature

The problem of privacy-preserving data collection has been widely considered based on the concept of differential privacy [9]. A DA, with a certain objective function, wants to obtain private information from privacy-aware EUs. However, the model and analysis differ based on the incentive methods adopted. There are mainly three categories: no incentive, i.e., an EU would participate to the data collection if the privacy protection offered by the DA checks its privacy constraints [10]; monetary payments, with the goal to design a mechanism that minimizes the DA's total payment subject to computation accuracy constraints [11]; non-monetary re-

wards which considers that EUs can benefit from the data result computation, and focuses on the EUs' trade-off between privacy and utility [12], [13]. Our work differs significantly from the framework of differential privacy, because in differential privacy, noise is added by a principal to the output of a computation which is subsequently publicly released [9], [14], and the principal performing the computation is trusted (i.e., it does not act strategically). In contrast, we do not assume that the EU is trusted, which motivates input perturbations that can be related to the more recent notion of local differential privacy, though ignoring the strategic aspect of the input perturbations [15].

Over the last few years, there has been a growing body of literature on data markets. On the market design side, Agarwal et al. propose an innovative design involving robust real-time matching, to efficiently buy and sell data [16], while Shilov et al. quantify the impact of privacy constraints on the market equilibrium [3]. Han et al. propose a regression market where EUs can monetize data traded among themselves for wind power forecasting [17]. On the pricing side, Cai et al. in [18] study a data market with many strategic data sources and one DA. The DA aims to build an estimator of a function, thanks to the readings reported by the sources. The quality of the readings depends on the level of effort exerted by the sources, which determines the variance of the sample as a decreasing function of the effort exerted, and the DA wants to design incentives to obtain the best possible estimator. Cai et al. prove that there exists a unique NE, which is also a dominant strategy equilibrium. Westenbroek et al. extend this result in [19] by taking into account the non-rivalry nature of data and considering multiple DAs among which the same data can be freely reproduced and shared. They prove that either no equilibrium exists or an infinite number of Generalized Nash Equilibria (GNEs), depending on the contract payment parameters. They also provide necessary and sufficient conditions for GNEs to be socially inefficient. In the same vein, Gast et al. in [20] propose a non-cooperative game model of strategic data sources, which optimize the variance of the noise level of their readings to minimize their cost, which can be split into a disclosure cost and an estimation cost based on linear regression. The impact of strategic information is quantified through explicit bounds on the price of stability. The above works differ from ours in that we explicitly account for the feedback-loop effect of the EUs' hidden ability.

Within the literature on private network games, [21], [22] demonstrate how private information, including social interaction structure and marginal benefits of EUs, can be inferred or revealed during information aggregation processes, and analyze the trade-off between privacy and utility, measured in terms of solution accuracy. Similarly, [23], [24] study the trade-off between the accuracy of learning equilibrium and privacy or fairness guarantees. An extension to online aggregative games for Nash Equilibrium tracking subject to privacy preservation, where the cost functions are time-varying, is provided in [25]. Finally, although strategic information and privacy often lead to system inefficiencies,

uncertainty can be beneficial in some cases. For example, in transportation networks, the overall congestion can be decreased when a principal carefully calibrates the level of information available to each EU [26], [27].

This paper is organized as follows. In Sec. II, we formulate the peer-to-peer electricity trading problem as a private network game. Relying on performative prediction, we define a decision-dependent game to account for the feedback-loop effect of the EUs' hidden ability, in Sec. III. The DA's contract parameter optimization is formulated as a bilevel optimization problem with induced equilibrium between the EUs, in Sec. IV. We subsequently reformulate it as a more tractable nonconvex program. Algorithms for equilibrium computation are provided in Sec. V, and convergence analysis is performed in Sec. VI. Simulations are run in Sec. VII. We conclude in Sec. VIII.

Notations

$\mathbb{E}_X[f]$ specifies that the expectation of the function $f(\cdot)$ is taken with respect to (w.r.t.) the random variable X (when X is clear from the context, we abbreviate the notation to $\mathbb{E}[f]$). Vectors and matrices are represented by bold letters. The stack of the N vectors $\mathbf{x}_1, \dots, \mathbf{x}_N$ is defined as $\mathbf{x} \stackrel{\text{def}}{=} \text{col}(\mathbf{x}_1, \dots, \mathbf{x}_N)$. We let $\text{diag}(\mathbf{x})$ be the diagonal matrix having on its main diagonal \mathbf{x} . We define \mathbf{x}_{-n} as the stack of the vectors containing the decision variables of all the agents in \mathcal{N} except n . The collection of N variables $\mathbf{x}_1, \dots, \mathbf{x}_N$ is denoted $(\mathbf{x}_n)_{n=1, \dots, N}$. We let \mathbf{x}^\top define the transpose of the vector \mathbf{x} and $\|\mathbf{x}\|$ be its Euclidean norm. The scalar product of vectors \mathbf{x}, \mathbf{y} of same dimension is denoted $\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{x}^\top \mathbf{y}$. $|\mathcal{X}|$ refers to the cardinality of the set \mathcal{X} . For a twice differentiable function $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$, $\nabla_{\mathbf{x}} f(\mathbf{x}, \mathbf{y})$ (resp. $\nabla_{\mathbf{y}} f(\mathbf{x}, \mathbf{y})$) denotes its partial gradient taken w.r.t. \mathbf{x} (resp. \mathbf{y}), and $\nabla_{\mathbf{y}\mathbf{x}}^2 f(\mathbf{x}, \mathbf{y})$ (resp. $\nabla_{\mathbf{x}\mathbf{y}}^2 f(\mathbf{x}, \mathbf{y})$) denotes the Jacobian of $\nabla_{\mathbf{y}} f(\mathbf{x}, \mathbf{y})$ at \mathbf{x} (resp. $\nabla_{\mathbf{x}} f(\mathbf{x}, \mathbf{y})$ at \mathbf{y}). For a given matrix \mathbf{A} , $A_{(i,j)}$ is the entry in position (i, j) and $\mathbf{A}_{(i,:)}$ stands for the i^{th} row. The matrix $\mathbf{A} > \mathbf{0}$ if and only if $A_{(i,j)} > 0, \forall i, j$. The trace and determinant of a square matrix \mathbf{A} are denoted $\text{tr}(\mathbf{A})$ and $\det(\mathbf{A})$ resp. For a closed set $\mathcal{C} \subseteq \mathbb{R}^m$, the mapping $\text{proj}_{\mathcal{C}} : \mathbb{R}^m \rightarrow \mathcal{C}$ denotes the projection onto \mathcal{C} , i.e., $\text{proj}_{\mathcal{C}}(\mathbf{x}) \stackrel{\text{def}}{=} \arg\min_{\mathbf{y} \in \mathcal{C}} \|\mathbf{y} - \mathbf{x}\|$. $\text{SOL}(P)$ is the set of solutions of Problem P . We denote $\mathbb{E}_t \stackrel{\text{def}}{=} \mathbb{E}[\cdot | \mathcal{F}_t]$, where $(\mathcal{F}_t)_t$ is a filtration such that ε^t is \mathcal{F}_{t+1} -measurable. We introduce Lip_1 as the set of 1-Lipschitz continuous functions. Finally, we define the ceiling function $\lceil x \rceil = \min\{n \in \mathbb{N} | n \geq x\}$. Nomenclature, symbols and notation tables are provided in supplementary material.

II. A PRIVATE NETWORK GAME MODEL

We formulate the peer-to-peer electricity trading problem as a private network game, involving strategic EUs with private information. We subsequently analyse the market equilibrium, giving conditions for uniqueness and stability.

A. Description of the Agents

1) *The End Users:* We let \mathcal{N} be a set of N EUs, forming a directed graph $\Gamma \stackrel{\text{def}}{=} (\mathcal{N}, E)$, where $E \subseteq \mathcal{N} \times \mathcal{N}$ is the set

of links between the EUs. Let Γ_n be the set of neighbors (i.e., in and out-neighbors) of EU n .

Each EU n updates the vector of decision variables $\mathbf{x}_n \stackrel{\text{def}}{=} (D_n, G_n^c, (q_{mn})_{m \in \Gamma_n}) \in \mathcal{X}_n$, which contains its demand, controllable generation, and energy trades with its neighbors, where $\mathcal{X}_n \subseteq \mathbb{R}^{m_n}$, $m_n = |\Gamma_n| + 2$, is the strategy set of EU n that captures only local constraints. In addition, we impose the following convention on EU n 's energy trades: $q_{mn} \geq 0$ means that n buys q_{mn} from m ; while $q_{mn} < 0$ means that n sells $-q_{mn}$ to m . Coupling constraints are added to capture bilateral trading reciprocity between EU n and its neighbors such that EU n 's strategy set takes the form: $\mathcal{C}_n(\mathbf{x}_{-n}) \stackrel{\text{def}}{=} \left\{ \mathbf{x}_n = (D_n, G_n^c, (q_{mn})_{m \in \Gamma_n}) \in \mathcal{X}_n \mid q_{mn} = -q_{nm}, \forall m \in \Gamma_n \right\}$, with $\mathcal{X}_n \stackrel{\text{def}}{=} \left\{ \mathbf{x}_n = (D_n, G_n^c, (q_{mn})_{m \in \Gamma_n}) \in \mathbb{R}^{m_n} \mid \underline{D}_n \leq D_n \leq \overline{D}_n, \underline{G}_n^c \leq G_n^c \leq \overline{G}_n^c, D_n \leq G_n^c + G_n^{nc} + Q_n \right\}$ where \underline{D}_n and \overline{D}_n are lower and upper-bounds on n 's demand, G_n^{nc} is the uncontrollable generation of EU n (e.g., solar PV panels generation), and $Q_n \stackrel{\text{def}}{=} \sum_{m \in \Gamma_n} q_{mn}$ is the net import at node n . The constraint $D_n \leq G_n^c + G_n^{nc} + Q_n$ captures the fact that demand at node n should not exceed the sum of local generation and net import. We formulate the balance constraint as an inequality to guarantee that the nodal price λ_n , i.e., the associated dual variable, remains non-negative, $\forall n$.

Typically, the EUs' utility functions, $j_n(\cdot), \forall n \in \mathcal{N}$, depend on some parameters that are sensitive, and that they might not want to disclose [2], [3]. We define $\theta_n \in \Theta_n \subseteq \mathbb{R}$ as the private information of EU n . Let $j_n : \mathbb{R}^{m_n} \times \Theta_n \rightarrow \mathbb{R}$ be the utility function of EU n , $\forall n \in \mathcal{N}$. Following [2], we define EU n 's utility (in the cost minimization sense) as the difference between the sum of the total cost incurred from the trades with its neighbors and the controllable generation cost, and its usage benefit $U_n(\cdot)$ ($-U_n(\cdot)$ is the consumption cost of EU n):

$$j_n(\mathbf{x}_n, \theta_n) = \sum_{m \in \Gamma_n} c_{nm} q_{mn} - U_n(D_n, \theta_n) + C_n^G(G_n^c), \quad (1)$$

where $\mathbf{c}_n > \mathbf{0}$ is a known vector of parameters defining the network charges or taxes, resulting from congestion. Thus, $\sum_{n \in \mathcal{N}} c_{nm} q_{mn}$ represents the amount of money collected by the network operator from the set \mathcal{N} of EUs, for their use of the network. If we let ζ_{nm} be the dual variable related to the bilateral trading reciprocity constraint $q_{mn} = -q_{nm}$ as defined in [1, Prop. 1], the money paid (resp. received) by EU n for buying from (resp. selling to) EU m is given by the perceived price $\zeta_{nm} - c_{nm}$ [28].

The controllable generation cost is assumed to be quadratic in G_n^c , i.e., $C_n^G(G_n^c) = \frac{a_n}{2} (G_n^c)^2 + b_n G_n^c + c_n$ where a_n, b_n, c_n are positive input coefficients. Quadratic generation costs is a classical assumption in the electricity market literature (see, e.g., [29]).

$-U_n(\cdot)$ is EU n 's consumption cost. We assume $U_n(\cdot)$ to be a quadratic function in EU n 's demand, i.e., we can write:

$$U_n(D_n, \theta_n) = -\tilde{a}_n (D_n - \theta_n)^2 + \tilde{b}_n, \quad (2)$$

where \tilde{a}_n, \tilde{b}_n are positive input coefficients, and θ_n is EU n 's nominal demand. EU n 's usage benefit is a strictly concave function in its demand D_n . It leads to the following interpretation: by being flexible around its nominal demand, e.g., agreeing to slightly reduce its demand around its nominal value, EU n will decrease its usage benefit but it might also decrease its bilateral trading or controllable generation costs. Similarly to the consumers' surplus for price, the usage benefit happens when the EU's demand is smaller than its nominal demand. Taking the integral over the demand support, it can therefore be expressed as a quadratic function in the EU's demand.

The impact of $(c_{nm})_{n,m}$ on the solutions of the non-cooperative game among the EUs has been analyzed in [2]. In case $(c_{nm})_{n,m}$ capture network charges or taxes, we assume that they are homogeneous among EUs in \mathcal{N} , since regulators forbid undue discrimination among network EUs.

Assumption 1: We assume that the EUs' preferences for trading are uniform and normalized, i.e., $c_{nm} = 1, \forall n, m$.

Since we do not consider storage and to avoid energy waste, supply and demand need to balance at any node of the network.

Assumption 2: Demand and supply balance each other at each node, i.e., $D_n = G_n^c + G_n^{nc} + Q_n, \forall n \in \mathcal{N}$.

We notice that under Asms. 1 and 2, controlling demand and generation allows to control the net import. Therefore, Q will be treated as a dependent decision variable in the rest of the work.

2) *The Market Operator:* The economic dispatch is classically obtained through a centralized clearing, i.e., the market operator aims to determine the vector \mathbf{x} that minimizes the social cost defined as the sum of the EUs' costs, i.e., $SC(\mathbf{x}, \boldsymbol{\theta}) \stackrel{\text{def}}{=} \sum_n j_n(\mathbf{x}_n, \boldsymbol{\theta}_n)$ under network coupling constraints, i.e., it solves the following optimization problem:

$$\min_{\mathbf{x}} SC(\mathbf{x}, \boldsymbol{\theta}), \quad (3a)$$

$$\text{s.t. } \mathbf{x} \in \mathcal{C}(\mathbf{x}) \stackrel{\text{def}}{=} \prod_n \mathcal{C}_n(\mathbf{x}_{-n}). \quad (3b)$$

We prove in [2] that the solution of (3) can be reached through a distributed clearing, where each EU $n \in \mathcal{N}$ solves the parametrized optimization problem:

$$\min_{\mathbf{x}_n \in \mathcal{C}_n(\mathbf{x}_{-n})} j_n(\mathbf{x}_n, \boldsymbol{\theta}_n). \quad (4)$$

The latter setting gives rise to a non-cooperative game $\mathcal{G} \stackrel{\text{def}}{=} (\mathcal{N}, \mathcal{C}(\cdot), (j_n(\cdot))_{n \in \mathcal{N}})$ between the EUs. We extend the analysis of \mathcal{G} from [2], by proving the global variational stability of its VE.

3) *The Data Aggregator:* Data has seen its role becoming more predominant in the new sharing economy, with a potential for being monetized [30]. As the demand for data increases, we have seen the formation of DAs, which collect data for either use or resale [19]. In our paper, the goal of the DA is to collect readings from the EUs to build an estimate \mathbf{x}^a of the EUs' decisions \mathbf{x} at equilibrium and use it to measure the quality of the readings. To this end, it builds

contracts with the EUs, and pays $p_n^a(\cdot)$ to each EU n , which is an increasing function in the accuracy of its estimate, \mathbf{x}_n^a , w.r.t. \mathbf{x}_n . To allow the DA to interact with the EUs, we need a service for communication network, which can take the form of a service marketplace platform.

B. Existence, Uniqueness and Stability of the Variational Equilibrium

Let $\mathcal{X} \stackrel{\text{def}}{=} \prod_n \mathcal{X}_n$ be the EUs' joint feasibility set. To generalize the notion of NE in the presence of coupling constraints, we use the Generalized Nash Equilibrium (GNE). It is defined as a strategy profile vector \mathbf{x} such that \mathbf{x}_n solves the system KKT_n for each n (see [1, Prop. 1]). The coupling constraint $q_{mn} = -q_{nm}$ written in the problem of n (respectively $m, \forall m \in \Gamma_n$) is associated with the duals ζ_{nm} (respectively ζ_{mn}). At a solution of \mathcal{G} , the bilateral trading reciprocity constraints need to be shared among all the EUs, i.e., $q_{mn} = -q_{nm}, \forall m \in \Gamma_n, \forall n \in \mathcal{N}$. Thus $\mathcal{C}(\mathbf{x})$ is a point-to-set mapping, and \mathcal{G} can be interpreted as a jointly convex generalized Nash game. To compute its set of solutions requires to solve a quasi-variational inequality. To compute a solution of \mathcal{G} , we focus on this special case of a GNE known as a Variational Equilibrium (VE), i.e., a solution $\mathbf{x} = (\mathbf{x}_n)_n$ such that \mathbf{x}_n solves the system KKT_n for each n and, in addition, such that the duals of the coupling constraints are equal, i.e., $\zeta_{nm} = \zeta_{mn}, \forall n \in \mathcal{N}, \forall m \in \Gamma_n$. The term "variational" refers to the variational inequality problem associated with such an equilibrium. To define this inequality, let us introduce the pseudo-gradient mapping $\mathbf{F}^0(\mathbf{x}) \stackrel{\text{def}}{=} \left(\nabla_{\mathbf{x}_n} j_n(\mathbf{x}_n, \boldsymbol{\theta}_n) \right)_n$. Under Asms. 1 and 2, EU n parameterized optimization problem (4) does not depend on $(q_{mn})_m$, but on $Q_n, \forall n$. Thus, we choose to replace the bilateral trading reciprocity constraints by the weaker constraint, $\sum_{n=1}^N Q_n = 0$. To compute a VE solution of \mathcal{G} , we define the joint feasibility set $\mathcal{C} \stackrel{\text{def}}{=} \{\mathbf{x} \in \mathcal{X} \mid \sum_{n=1}^N Q_n = 0\}$. Note that a VE, $\mathbf{x}^{\text{VE}} = (\mathbf{D}^{\text{VE}}, (\mathbf{G}^c)^{\text{VE}}, \mathbf{Q}^{\text{VE}})$, may give rise to multiple optimal $(q_{mn}^{\text{VE}})_m, \forall n$. To compute a solution of \mathcal{G} , we need to check *a posteriori* that $q_{mn}^{\text{VE}} = -q_{nm}^{\text{VE}}, \forall m \in \Gamma_n, \forall n$.

Definition 1 (Variational Equilibrium [31]): The Variational Inequality problem $\text{VI}(\mathbf{F}^0, \mathcal{C})$ consists in finding a vector $\mathbf{x}^{\text{VE}} \in \mathcal{C}$ such that $(\mathbf{y} - \mathbf{x}^{\text{VE}})^\top \mathbf{F}^0(\mathbf{x}^{\text{VE}}) \geq 0, \forall \mathbf{y} \in \mathcal{C}$. A solution of $\text{VI}(\mathbf{F}^0, \mathcal{C})$ is called a VE.

Proposition 1: [1, Prop. 1] Let Asm. 2 hold, then EU n 's best-response strategy \mathbf{x}_n^{VE} at equilibrium depends linearly on the nominal demands of *all* the EUs, i.e.,

$$\mathbf{x}_n^{\text{VE}} = \mathbf{A}_{n,n} \boldsymbol{\theta}_n + \mathbf{A}_{n,-n} \boldsymbol{\theta}_{-n} + \left[\mathbf{B}_{n,n} G_n^{nc} + \sum_{n' \neq n} \mathbf{B}_{n,n'} G_{n'}^{nc} \right] + \mathbf{C}_n,$$

with $\mathbf{A}_{n,n}, \mathbf{B}_{n,n}, \mathbf{B}_{n,-n}, \mathbf{C}_n$ vectors and $\mathbf{A}_{n,-n}$ matrix of appropriate dimensions; otherwise \mathbf{x}_n^{VE} depends only on the EU n 's own nominal demand, $\boldsymbol{\theta}_n$.

Remark 2.1: The coefficients of the best-response strategies are not known by the EUs because they rely on coefficients $(\tilde{a}_n, a_n)_n$ that the other EUs might not want to share.

Assumption 3: We assume that there exists at least one EU $n \in \mathcal{N}$ such that at the VE, $D_n^{\text{VE}} \neq 0$ or $G_n^{c\text{VE}} \neq 0$.

Theorem 1: [1, Thm. 1] Let Asm. 3 hold, then the game \mathcal{G} admits a unique VE, which is in addition efficient.

Definition 2 (Variational Stability [32]): An equilibrium point $\mathbf{x}^{\text{VE}} \in \mathcal{C}$ is said to be variationally stable (or simply stable) if there exists a neighborhood \mathcal{X} of \mathbf{x}^{VE} such that $(\mathbf{y} - \mathbf{x}^{\text{VE}})^\top \mathbf{F}^0(\mathbf{y}) \geq 0$ for all $\mathbf{y} \in \mathcal{X}$. In particular, if this property holds for all $\mathbf{y} \in \mathcal{C}$, we say that \mathbf{x}^{VE} is globally stable.

Proposition 2: Let Asm. 3 hold, then the VE solution of \mathcal{G} is globally stable.

Proof: In Thm. 1 proof, we check that \mathcal{G} belongs to the class of generalized potential games (GPG) [33, Sec. 2], for which the EUs' utility functions are aligned along the game's potential whose minimizers are VEs. From [32, Rk. 1], VEs solutions of convex potential games are automatically stable in the sense of Def. 2. ■

Let Asms. 1–3 hold for the rest of the paper.

III. A DECISION-DEPENDENT GAME WITH FEEDBACK-LOOP EFFECT

In this paper, we focus on the performative prediction setting with decision-dependent distributions, which has been introduced in [8] as a game-theoretic model for decision-dependent learning in the presence of competition. It leads to a weaker solution concept than Nash Equilibrium, but allows to endogenously capture the performative effects caused by the EUs' hidden ability on the market robustness.

A. Balancing Cost and Privacy

We introduce randomization in the reading sent by EU n to the DA, which takes the form $\hat{\theta}_n = \theta_n + \varepsilon_n$. The first part of the reading, $\hat{\theta}_n$, captures the ability of EU n to add a deterministic bias to its reported information. The second part, ε_n , is a noise following a randomized mechanism $M(\cdot)$, added in order to preserve privacy. We assume that the noise in the readings is independent across the EUs [19]. $\hat{\theta}_n$ and ε_n are known by EU n but not by the other participants.

Assumption 4: For each EU $n \in \mathcal{N}$, ε_n follows a centered Gaussian distribution with variance V_n . Furthermore, the random variables $(\varepsilon_n)_{n \in \mathcal{N}}$ are independent.

To incentivize the EUs to report truthful readings that can later be reused or monetized through targeted advertisements, the DA comes into play. The DA defines a payoff contract to each EU; the payment is decreasing in the distance between the EUs' strategy \mathbf{x}_n and the estimate $\mathbf{x}_n^a \stackrel{\text{def}}{=} (D_n^a, G_n^{c\ a})$ built by the DA relying on linear regression, defined as follows:

$$\mathbf{x}_n^a = \mathbf{A}_n^a \tilde{\boldsymbol{\theta}} + \mathbf{B}_n^a \mathbf{G}^{nc} + \mathbf{C}_n^a, \quad (5)$$

where $\mathbf{A}_n^a, \mathbf{B}_n^a, \mathbf{C}_n^a$ are estimates of $\mathbf{A}_n, \mathbf{B}_n, \mathbf{C}_n$ introduced in Prop. 1. The form of the DA's estimator is known by all the EUs. It comes from Prop. 1. The contract between the EU n and the DA takes the form:

$$p_n^a(\mathbf{x}_n, \mathbf{x}_n^a, \tilde{\boldsymbol{\theta}}) = \gamma_n - \beta_n \left((D_n - D_n^a)^2 + (G_n^c - G_n^{c\ a})^2 \right), \quad (6)$$

where γ_n, β_n are non-negative parameters which can be optimized by the DA to check specific market properties [19]. The contract payment targeted at EU n depends on the DA's estimates \mathbf{x}_n^a and the measured demand and controllable generation \mathbf{x}_n that it can obtain through the interface deployed in each node of the network.

Each EU n chooses the deterministic part of its reading $\hat{\theta}_n$ such that $d(\hat{\theta}_n, \theta_n) \stackrel{\text{def}}{=} (\hat{\theta}_n - \theta_n)^2 \leq \alpha_n$ where α_n is EU n 's adjacency level, as defined in differential privacy [34]. Let $L_{M, \theta_n, \hat{\theta}_n}$ be the privacy loss of $M(\cdot)$, which follows the distribution $\mathcal{N}\left(\frac{(\theta_n - \hat{\theta}_n)^2}{2V_n}, \frac{(\theta_n - \hat{\theta}_n)^2}{V_n}\right)$ [35]. How much information can be extracted from the reading $\tilde{\theta}_n$ is measured by the expected privacy loss of the mechanism $M(\cdot)$: $I_n(\theta_n, \tilde{\theta}_n) \stackrel{\text{def}}{=} -\mathbb{E}_{\varepsilon_n} \left[L_{M, \theta_n, \hat{\theta}_n} \right]$, which coincides with the Kullback-Leibler divergence (relative entropy) between $M(\cdot)$'s output distributions on θ_n and $\hat{\theta}_n$. Under Asm. 4, the Kullback-Leibler divergence takes the closed form: $I_n(\theta_n, \hat{\theta}_n, V_n) = \frac{(\hat{\theta}_n - \theta_n)^2}{2V_n}$.

B. Formulation of the Stochastic Game

In the service marketplace platform, EU n 's extended utility is split into four terms. The first term captures the cost it bears in \mathcal{G} ; the second term measures the cost of information release, while the third term is the contract payment it receives from the DA. A last term is added, which measures the penalty associated with the relaxation of the EUs' coupling constraint:

$$J_n(\mathbf{x}_n, \hat{\boldsymbol{\theta}}, \mathbf{V}, \boldsymbol{\varepsilon}) = j_n(\mathbf{x}_n, \theta_n) + c_n^I I_n(\theta_n, \hat{\theta}_n, V_n) - p_n^a(\mathbf{x}_n, \mathbf{x}_n^a, \tilde{\boldsymbol{\theta}}) + \xi \left(\sum_{m=1}^N Q_m \right)^2, \quad (7)$$

with c_n^I, ξ non-negative parameters. From (5), since $D_n^a, G_n^{c\ a}$ depend on $\hat{\boldsymbol{\theta}}, \boldsymbol{\varepsilon}$, and $I_n(\cdot)$ on $\hat{\boldsymbol{\theta}}, \mathbf{V}$, $J_n(\cdot)$ also depends on $\hat{\boldsymbol{\theta}}, \mathbf{V}, \boldsymbol{\varepsilon}$. To avoid coupling between the feasibility sets of the EUs, we put the global coupling constraint as a penalty in their utility function, with a weight ξ that is the same for all the EUs. The penalty ξ calibration is discussed in Lem. 1.

We introduce $\mathcal{V} \stackrel{\text{def}}{=} \text{diag}(\mathbf{V})$. Each EU n solves:

$$\min_{\mathbf{x}_n, \hat{\theta}_n, V_n} \mathbb{E}_{\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \mathcal{V})} \left[J_n(\mathbf{x}_n, \hat{\boldsymbol{\theta}}, \mathbf{V}, \boldsymbol{\varepsilon}) \right], \quad (8a)$$

$$\text{s.t. } \underline{D}_n \leq D_n \leq \overline{D}_n, \quad (8b)$$

$$\underline{G}_n^c \leq G_n^c \leq \overline{G}_n^c, \quad (8c)$$

$$D_n = G_n^c + G_n^{nc} + Q_n, \quad (8d)$$

$$d(\hat{\theta}_n, \theta_n) \leq \alpha_n, \quad (8e)$$

$$\underline{V}_n \leq V_n \leq \overline{V}_n. \quad (8f)$$

Let \mathcal{Y}_n be EU n 's feasibility set. Problem (8) taken for each EU $n \in \mathcal{N}$ gives rise to the stochastic game $\mathcal{G}^{\text{stoch}} \stackrel{\text{def}}{=} (\mathcal{N}, (\mathcal{Y}_n)_{n \in \mathcal{N}}, (\mathbb{E}_{\boldsymbol{\varepsilon}} [J_n(\cdot)])_{n \in \mathcal{N}})$. Each EU n defines a minimum hidden ability, so it imposes a lower bound \underline{V}_n on V_n . Furthermore, to limit the instability of the market, the DA imposes an upper bound \overline{V}_n on V_n . This upper bound does not need to be shared by the EUs, instead it might be encoded

by the DA in the communication interface deployed in the EUs' nodes when communicating with the DA. Note that constraint (8e) also imposes an upper bound on the amount of information released by EU n , i.e., $0 \leq I_n(\cdot) \leq \frac{\alpha_n}{2\bar{V}_n}, \forall n$.

The game between the EUs, which involves the DA as a principal, takes place as follows:

- a) Taking $\hat{\theta}, \mathbf{V}$ as parameters, each EU n computes its $\mathbf{x}_n(\hat{\theta}, \mathbf{V})$ by solving (8) in \mathbf{x}_n .
- b) Each EU n computes its hidden ability parameters by solving (8) w.r.t. $\hat{\theta}_n, V_n$. Then, it samples a reading $\hat{\theta}_n = \hat{\theta}_n + \varepsilon_n$, with $\varepsilon_n \sim \mathcal{N}(0, V_n)$ that it reports to the DA.

The readings $\tilde{\theta}$ are then broadcast to all the EUs by the DA. Finally, each EU n receives a contract payment $p_n^a(\mathbf{x}_n, \mathbf{x}_n^a, \tilde{\theta}), \forall n \in \mathcal{N}$.

Note that the timing of steps a) and b) determines the structure of the non-cooperative game $\mathcal{G}^{\text{stoch}}$. If both steps occur simultaneously, then $\mathcal{G}^{\text{stoch}}$ can be interpreted as a Nash game. If $\mathcal{G}^{\text{stoch}}$ leads to EUs' strategies which are optimal for the distribution they induce, the resulting equilibria coincide with the stable points of retraining, i.e., are Performatively Stable. If step a) is performed independently and selfishly by each EU while anticipating the hidden ability (defined through bias and variance) of the EUs, and the EUs' reading distributions are best-response strategies to step a) decision variables, then $\mathcal{G}^{\text{stoch}}$ is said to have a Stackelberg structure.

Let $\mathbf{v}_n \stackrel{\text{def}}{=} (\hat{\theta}_n, V_n)$ and $\mathbf{y}_n \stackrel{\text{def}}{=} (\mathbf{x}_n, \hat{\theta}_n, V_n)$ be the decision variables of EU n in $\mathcal{G}^{\text{stoch}}$. We denote $\mathcal{U}_n \stackrel{\text{def}}{=} \{\mathbf{v}_n \in \mathbb{R}_+^2 \mid (8e), (8f) \text{ hold}\}$ and $\mathcal{Y}_n = \{\mathbf{y}_n \in \mathbb{R}^{m_n} \times \mathbb{R}_+^2 \mid (8b)-(8f) \text{ hold}\}$. Due to the penalty in the objective function, it is required that all the EUs have access to the sum of the net imports. It is a reasonable hypothesis because the market operator can broadcast this information without revealing much information about the EUs.

In the non-penalized case, we let $\check{J}_n(\cdot)$ be the function $J_n(\cdot)$ without the penalty for EU $n, \forall n$. This gives rise to the stochastic generalized Nash game $\check{\mathcal{G}}^{\text{stoch}} \stackrel{\text{def}}{=} (\mathcal{N}, \mathcal{C}, (\mathbb{E}_\varepsilon[\check{J}_n(\cdot)]))$ where each EU n solves:

$$\begin{aligned} & \min_{\mathbf{x}_n, \hat{\theta}_n, V_n} \mathbb{E}_{\varepsilon \sim \mathcal{N}(0, \mathbf{V})} \left[\check{J}_n(\mathbf{x}_n, \hat{\theta}, \mathbf{V}, \varepsilon) \right], \\ & \text{s.t. } (8b)-(8f), \\ & \sum_m Q_m = 0. \end{aligned}$$

We clarify the relationship between $\check{\mathcal{G}}^{\text{stoch}}$ and $\mathcal{G}^{\text{stoch}}$ in the lemma below.

Lemma 1: If the penalty coefficient ξ in $J_n(\cdot)$ is bigger than the Lagrangian multiplier ϕ of the constraint $(\sum_{m=1}^N Q_m)^2 = 0$, then the penalized game $\mathcal{G}^{\text{stoch}}$ has the same set of solutions as $\check{\mathcal{G}}^{\text{stoch}}$.

We aim to use $\mathcal{G}^{\text{stoch}}$, which is easier to solve because it does not involve any coupling constraint, to compute solutions of the generalized Nash game $\check{\mathcal{G}}^{\text{stoch}}$. This requires an appropriate calibration of the penalty coefficient ξ detailed in Appendix A.

C. Equilibrium Analysis

We recall three equilibrium solution concepts, which can be used to interpret the outcome of $\mathcal{G}^{\text{stoch}}$. While NE does not enable to capture the distribution shifts caused by the EUs' hidden ability on the market, we introduce PSE and SE as alternative solution concepts.

Definition 3 (Nash Equilibrium): A Nash Equilibrium (NE) is a strategy profile $\mathbf{y}^{\text{NE}} = (\mathbf{y}_n^{\text{NE}})_n \in \mathcal{Y} = \prod_n \mathcal{Y}_n$ such that for all $n \in \mathcal{N}$

$$\mathbf{y}_n^{\text{NE}} \in \operatorname{argmin}_{\mathbf{y}_n \in \mathcal{Y}_n} \mathbb{E}_{\varepsilon \sim \mathcal{N}(0, \operatorname{diag}(\mathbf{V}_n, \mathbf{V}_n^{\text{NE}}))} \left[J_n(\mathbf{y}_n, \mathbf{y}_{-n}^{\text{NE}}, \varepsilon) \right].$$

There is another solution concept which is more amenable to numerical methods [8], and reduces to Performatively Stable points (PSEs). The idea of the PSE is to decouple the effect of a decision \mathbf{y} on the integrand $J_n(\mathbf{y}_n, \mathbf{y}_{-n}, \varepsilon)$ and on the distribution $\mathcal{N}(0, \mathbf{V})$. Any vector \mathbf{V}^{PSE} induces a static game wherein the distribution for EU n is fixed at $\mathcal{N}(0, \mathbf{V}^{\text{PSE}})$, with $\mathbf{V}^{\text{PSE}} \stackrel{\text{def}}{=} \operatorname{diag}(\mathbf{V}^{\text{PSE}})$, that is:

Definition 4 (Performatively Stable Equilibrium [8]):

A Performatively Stable Equilibrium (PSE) is a strategy profile $\mathbf{y}^{\text{PSE}} = (\mathbf{y}_n^{\text{PSE}})_n \in \mathcal{Y}$ such that for all $n \in \mathcal{N}$

$$\mathbf{y}_n^{\text{PSE}} \in \operatorname{argmin}_{\mathbf{y}_n \in \mathcal{Y}_n} \mathbb{E}_{\varepsilon \sim \mathcal{N}(0, \mathbf{V}^{\text{PSE}})} \left[J_n(\mathbf{y}_n, \mathbf{y}_{-n}^{\text{PSE}}, \varepsilon) \right]. \quad (10)$$

PSE has a clear intuitive meaning, which makes it more relevant in our setting than NE: each EU has no incentive to deviate from \mathbf{y}^{PSE} having access only to readings drawn from $\mathcal{N}(0, \mathbf{V}^{\text{PSE}})$.

Performative Optimality is another solution concept resulting from the casting of the problem as a Stackelberg game. EUs react to the economic dispatch by adapting their hidden ability with the hope to improve their outcomes. This game is said to have a Stackelberg structure since EUs adapt their bias and variance only after the market has cleared. In general, Performative Optimality and Performative Stability are two distinct solution concepts [7].

Definition 5 (Stackelberg Equilibrium [36], [37]): A Stackelberg Equilibrium (SE) is a strategy profile $\mathbf{y}^{\text{SE}} \in \mathcal{Y}$ obtained when each EU $n \in \mathcal{N}$ aims to solve the following optimization problem:

$$\begin{aligned} & \min_{\mathbf{x}_n \in \mathcal{X}_n} \left\{ \mathbb{E}_{\varepsilon \sim \mathcal{N}(0, \mathbf{V})} \left[J_n(\mathbf{y}_n, \mathbf{y}_{-n}, \varepsilon) \right] \mid \mathbf{v}_m \in \bar{\mathcal{U}}_m, \forall m \right\}, \\ & \text{where } \forall m \in \mathcal{N}, \\ & \bar{\mathcal{U}}_m \stackrel{\text{def}}{=} \operatorname{argmin}_{\mathbf{u}_m \in \mathcal{U}_m} \mathbb{E}_{\varepsilon \sim \mathcal{N}(0, \mathbf{V})} \left[J_m(\mathbf{x}_m, \mathbf{u}_m, \mathbf{x}_{-m}, \mathbf{u}_{-m}, \varepsilon) \right], \\ & \min_{\mathbf{v}_m \in \bar{\mathcal{U}}_m} \mathbb{E}_{\varepsilon \sim \mathcal{N}(0, \mathbf{V})} \left[J_m(\mathbf{y}_m, \mathbf{y}_{-m}, \varepsilon) \right], \forall m \in \mathcal{N}. \end{aligned} \quad (11)$$

The following assumption states that the cost of information should be larger than EU n 's maximum noise variance.

Assumption 5: For any EUs $n, m \in \mathcal{N}, n \neq m$, $c_n^I \geq \beta_m (A_m^a(1, n) + A_m^a(2, n))^2 \bar{V}_n$.

Let $D_n^a(\hat{\theta}) \stackrel{\text{def}}{=} \mathbb{E}_\varepsilon[D_n^a]$ and $G_n^a(\hat{\theta}) \stackrel{\text{def}}{=} \mathbb{E}_\varepsilon[G_n^a]$ where $\mathbf{x}_n^a = (D_n^a, G_n^a)$ has been introduced in Equation (5). In addition, we define $\mathbf{F}_{\mathbf{y}_n}$ and $\mathbf{H}_{\mathbf{y}_n}$ as the partial derivative and partial Hessian of $\mathbb{E}_\varepsilon[J_n(\cdot)]$ w.r.t. $\mathbf{y}_n, \forall n$, and $\widehat{\mathbf{H}}_{\mathbf{y}}$ as the pseudo Hessian of the utilities of the EUs. The following proposition will be used to prove the uniqueness of the

NE and PSE in Prop. 4 and 5 resp., and later to derive convergence rates in Sec. VI.

Proposition 3: Let Asm. 5 hold. Then, the pseudo Hessian matrix $\tilde{\mathbf{H}}_{\mathbf{y}} = \mathbf{H}_{\mathbf{y}} + \mathbf{H}_{\mathbf{y}}^T$ of utilities of the EUs is positive definite.

We set $K_n \stackrel{\text{def}}{=} (A_n^a)_{(1,n)}^2 + (A_n^a)_{(2,n)}^2, \forall n$.

Theorem 2: There exists a unique NE solution of $\mathcal{G}^{\text{stoch}}$.

Proof: From Prop. 3, the pseudo-gradient map $\mathbf{F}_{\mathbf{y}}(\mathbf{y}) = (F_{\mathbf{y}_n}(\mathbf{y}))_n$ is strongly monotone on $\mathcal{Y} = \prod_n \mathcal{Y}_n$. Furthermore, for all n , the feasibility set \mathcal{Y}_n is convex and compact. We immediately have that $\mathcal{G}^{\text{stoch}}$ has a unique solution [31]. ■

The following proposition gives a necessary condition for the NE solution of $\mathcal{G}^{\text{stoch}}$ to be efficient. This result has an intuitive interpretation: by introducing bias and noise in its reading, the EU contributes to decrease the efficiency of the overall system.

Proposition 4: Let the DA's estimate \mathbf{x}^a be perfect for the game \mathcal{G} , i.e., $\mathbf{x}^a \in \text{SOL}(\mathcal{G})$. Then, the NE solution of $\mathcal{G}^{\text{stoch}}$ is efficient.

Proof: We denote $\mathbf{x}^{\mathcal{G}}$ the VE solution of the game \mathcal{G} . We have $\forall \hat{\boldsymbol{\theta}}, \mathbf{V}, \mathbf{x}_n, j_n(\mathbf{x}^{\mathcal{G}}) \leq j_n(\mathbf{x}_n) + c_n^I \frac{(\hat{\theta}_n - \theta_n)^2}{2V_n} + \beta_n \left((D_n - D_n^a(\hat{\boldsymbol{\theta}}))^2 + (G_n - G_n^a(\hat{\boldsymbol{\theta}}))^2 \right) + \xi \left(\sum_{m=1}^N Q_m \right)^2 \implies j_n(\mathbf{x}^{\mathcal{G}}) + \sum_{m=1}^N K_m V_m \leq j_n(\mathbf{x}_n) + c_n^I \frac{(\hat{\theta}_n - \theta_n)^2}{2V_n} + \beta_n \left((D_n - D_n^a(\hat{\boldsymbol{\theta}}))^2 + (G_n - G_n^a(\hat{\boldsymbol{\theta}}))^2 \right) + \xi \left(\sum_{m=1}^N Q_m \right)^2 + \sum_{m=1}^N K_m V_m$. Note that the term $-\gamma_n$ is constant and does not change the equilibrium, so we omit it. This lower bound is reached in $\hat{\boldsymbol{\theta}} = \boldsymbol{\theta}, V_n = \underline{V}_n, \forall n$ and $\mathbf{x}_n = \mathbf{x}_n^{\mathcal{G}}$. This means that we have proved that assuming \mathbf{x}^a is perfect for \mathcal{G} , $\mathbf{x}^{\mathcal{G}}, \boldsymbol{\theta}, \underline{\mathbf{V}}$ is the NE solution of $\mathcal{G}^{\text{stoch}}$. If we sum the former inequalities over n , we observe that $\hat{\boldsymbol{\theta}} = \boldsymbol{\theta}, V_n = \underline{V}_n, \forall n$ and $\mathbf{x} = \mathbf{x}^{\mathcal{G}}$ also minimizes the EUs' expected social cost. Thus, assuming \mathbf{x}^a is perfect for \mathcal{G} , the NE solution of $\mathcal{G}^{\text{stoch}}$ is efficient. ■

Next, we look at the properties of the PSE.

Theorem 3: There exists a unique PSE solution of $\mathcal{G}^{\text{stoch}}$.

Proof: To prove this result, we rewrite $\mathbb{E}_{\varepsilon}[J_n(\cdot)]$ as follows, to use [8, Thm. 1]:

$$\mathbb{E}_{\varepsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{V})} [J_n(\mathbf{x}_n, \hat{\boldsymbol{\theta}}, \mathbf{V}, \varepsilon)] = \mathbb{E}_{\varepsilon_n \sim \mathcal{N}(0, V_n)} [\bar{J}_n(\mathbf{x}_n, \hat{\boldsymbol{\theta}}, \mathbf{V}, \varepsilon_n)],$$

with $\bar{J}_n(\cdot)$ an approximating function of $J_n(\cdot)$ such that

$$\bar{J}_n(\mathbf{x}_n, \hat{\boldsymbol{\theta}}, \mathbf{V}, \varepsilon_n) \stackrel{\text{def}}{=} J_n(\mathbf{x}_n, \hat{\boldsymbol{\theta}}, \mathbf{V}, \mathbf{0}) + \sum_{m=1, m \neq n}^N K_m V_m + K_n \varepsilon_n^2, \quad (12)$$

where $J_n(\mathbf{x}_n, \hat{\boldsymbol{\theta}}, \mathbf{V}, \mathbf{0})$ is $J_n(\cdot)$ taken in $\varepsilon = \mathbf{0}$. To prove the strong convexity of $\bar{J}_n(\cdot)$, we proceed the same way as to prove the strong convexity of $J_n(\cdot)$, because the partial Hessians of both objective functions are the same. We also check the \mathcal{C}^1 continuity of $\bar{J}_n(\cdot)$, the Lipschitz continuity of the application $\varepsilon \mapsto \bar{J}_n(\mathbf{y}, \varepsilon)$, and that the distribution of the ε varies in a Lipschitz way according to [8, Asm. 2]. From [8, Thm. 1], we obtain the existence and uniqueness of the PSE. ■

IV. FORMULATION OF THE DATA AGGREGATOR PROBLEM

We introduce some notation for brevity here: we let $p_n^a(\cdot)$ denote the expected value of the payment to EU n , $p_n^a(\cdot)$, i.e., $p_n^a(\gamma_n, \beta_n, \mathbf{y}) \stackrel{\text{def}}{=} \mathbb{E}_{\varepsilon} [p_n^a(\mathbf{x}_n, \mathbf{x}_n^a, \tilde{\boldsymbol{\theta}})]$. The DA defines payment so as to minimize:

$$L^a(p^a, \mathbf{y}) = \pi \sum_n p_n^a(\gamma_n, \beta_n, \mathbf{y}) + \mathbb{E}_{\varepsilon} [l(\mathbf{x}, \mathbf{x}^a)], \quad (13)$$

where the parameter $\pi > 0$ denotes a conversion between monetary amounts allocated by the payment functions and the expectation of the loss $l(\mathbf{x}, \mathbf{x}^a) \stackrel{\text{def}}{=} \sum_n [(D_n - D_n^a)^2 + (G_n^c - G_n^{c,a})^2]$ generated by the estimation gap, and $D_n^a = x_{n(1,1)}^a, G_n^{c,a} = x_{n(2,1)}^a$ are obtained from (5).

We describe the DA's optimization problem, and specify constraints that the parameter choice $\boldsymbol{\gamma}, \boldsymbol{\beta}$ must satisfy.

The first constraint is Individual Rationality (IR). IR requires that each EU prefers contracting with the DA, than with the conventional retailer. Based on the definition of $J_n(\cdot)$ and $p_n^a(\cdot)$, we observe that $\mathbb{E}_{\varepsilon} [J_n(\mathbf{x}_n, \hat{\boldsymbol{\theta}}, \mathbf{V}, \varepsilon)] = \mathbb{E}_{\varepsilon} [J_n(p_n^a(\mathbf{x}_n, \mathbf{x}_n^a, \tilde{\boldsymbol{\theta}}), \mathbf{y}_n)] = \mathbb{E}_{\varepsilon} [J_n(p_n^a(\gamma_n, \beta_n, \mathbf{y}), \mathbf{y}_n)]$. Using these relationships, IR means that *ex-ante*:

$$\mathbb{E}_{\varepsilon} [J_n(p_n^a(\gamma_n, \beta_n, \mathbf{y}), \mathbf{y}_n)] \leq p^r D_n(\boldsymbol{\beta}),$$

where p^r is the unit price fixed by the conventional retailer.

The second constraint is non-negativity of the payments from the DA, i.e., the payment that the DA offers to each EU n is positive *ex-ante*.

The third constraint is Incentive Compatibility (IC). It is enforced by the DA finding $\mathbf{y}^{\text{NE}}(\cdot)$ that minimizes its cost (13), and then designing $p_n^a(\cdot)$ such that:

$$\mathbf{y}_n^{\text{NE}}(\boldsymbol{\beta}) = \underset{\mathbf{y}_n \in \mathcal{Y}_n}{\text{argmin}} \mathbb{E}_{\varepsilon} [J_n(p_n^a(\gamma_n, \beta_n, \mathbf{y}), \mathbf{y}_n)], \forall \mathbf{y}_{-n} \in \mathcal{Y}_{-n}.$$

With IC, IR, non-negativity constraints, we formulate a bilevel optimization problem for the DA:

$$\min_{\boldsymbol{\gamma}, \boldsymbol{\beta}} L^a(p^a, \mathbf{y}^{\text{NE}}(\boldsymbol{\beta})), \quad (14a)$$

$$\text{s.t. } \mathbf{y}_n^{\text{NE}}(\boldsymbol{\beta}) = \underset{\mathbf{y}_n \in \mathcal{Y}_n}{\text{argmin}} \mathbb{E}_{\varepsilon} [J_n(p_n^a(\gamma_n, \beta_n, \mathbf{y}_n, \mathbf{y}_{-n}), \mathbf{y}_n)],$$

$$\forall \mathbf{y}_{-n} \in \mathcal{Y}_{-n}, \forall n, \quad (14b)$$

$$\sum_{m=1}^N (D_m^{\text{NE}}(\boldsymbol{\beta}) - G_m^{c,\text{NE}}(\boldsymbol{\beta}) - G_m^{nc}) = 0, \quad (14c)$$

$$p^r D_n^{\text{NE}}(\boldsymbol{\beta}) - \mathbb{E}_{\varepsilon} [J_n(p_n^a(\gamma_n, \beta_n, \mathbf{y}^{\text{NE}}(\boldsymbol{\beta})), \mathbf{y}_n^{\text{NE}}(\boldsymbol{\beta}))] \geq 0,$$

$$\forall n \in \mathcal{N}, \quad (14d)$$

$$p_n^a(\gamma_n, \beta_n, \mathbf{y}^{\text{NE}}(\boldsymbol{\beta})) \geq 0, \forall n \in \mathcal{N}, \quad (14e)$$

$$\gamma_n \geq 0, \beta_n \geq 0, \forall n \in \mathcal{N}. \quad (14f)$$

This problem actually has N coupled optimization problems as constraints, making it a difficult bilevel program. However, we will reformulate the DA's problem to a more manageable nonlinear program in the sequel.

A. Induced Equilibrium between the End Users

To ensure the notion of IC in equilibrium, we show that there is a well-defined mapping $(\gamma, \beta) \mapsto \mathbf{y}^{\text{NE}}$.

Definition 6: For fixed contract payments $(\mathbf{p}_n^a)_n$, we say that \mathbf{y}^{NE} is an induced NE if for each EU $n \in \mathcal{N}$:

$$\mathbf{y}_n^{\text{NE}} = \underset{\mathbf{y}_n \in \mathcal{Y}_n}{\operatorname{argmin}} \mathbb{E}_\varepsilon \left[\check{J}_n(\mathbf{p}_n^a(\gamma_n, \beta_n, \mathbf{y}_n, \mathbf{y}_{-n}^{\text{NE}}), \mathbf{y}_n) \right]. \quad (15)$$

By definition $x_n^a(\hat{\theta}) = \mathbb{E}_\varepsilon[x_n^a] = \mathbf{A}_n^a \hat{\theta} + \mathbf{B}_n^a \mathbf{G}^{mc} + \mathbf{C}_n^a$, and $D_n^a(\hat{\theta}) = x_n^{a(1,1)}(\hat{\theta})$, $G_n^{ca}(\hat{\theta}) = x_n^{a(2,1)}(\hat{\theta})$. Suppose that we have a set of payments (γ, β) , EU n chooses \mathbf{y}_n^{NE} (which is a function of β) according to:

$$\begin{aligned} \mathbf{y}_n^{\text{NE}}(\beta) = \underset{\mathbf{y}_n \in \mathcal{Y}_n}{\operatorname{argmin}} & \left[j_n(\mathbf{x}_n, \theta_n) + c_n^I \frac{(\hat{\theta}_n - \theta_n)^2}{2V_n} - \gamma_n \right. \\ & \left. + \beta_n \left((D_n - D_n^a(\hat{\theta}))^2 + (G_n - G_n^{ca}(\hat{\theta}))^2 \right) + \mathbf{K}^\top \mathbf{V} \right], \quad (16) \end{aligned}$$

for each choice of $\mathbf{y}_{-n} \in \mathcal{Y}_{-n}$ made by the other EUs. According to Prop. 2, (16) admits a unique globally optimal solution for EU n . Moreover, the choice of the optimal \mathbf{y}_n^{NE} is affected by the choice of \mathbf{y}_{-n} . Thus, each choice of contract parameters β selected by the DA leads to an induced strategy NE for the EUs. Notice that \mathbf{y}_n^{NE} are coupled in $\beta, \forall n$.

Next, we define $\Psi_n : \mathbb{R}_+^N \rightarrow \mathcal{Y}_n$ to be the implicitly defined map such that $\Psi_n : \beta \mapsto \mathbf{y}_n^{\text{NE}}$ returns the solution of (16) for each EU $n \in \mathcal{N}$, for a given choice of $\beta \in \mathbb{R}_+^N$. By definition, for each EU n , \mathcal{Y}_n is compact as the product of compact sets, therefore there exists $\underline{\mathbf{y}}_n, \bar{\mathbf{y}}_n$ such that $\underline{\mathbf{y}}_n \leq \mathbf{y}_n \leq \bar{\mathbf{y}}_n, \forall n \in \mathcal{N}$. Let $\underline{\beta} \stackrel{\text{def}}{=} \min\{n \in \mathcal{N}, \beta \in \mathbb{R}_+^N \mid \Psi_n(\beta) = \underline{\mathbf{y}}_n\}$. We define $B \stackrel{\text{def}}{=} [\underline{\beta}, \bar{\beta}]$ where $\bar{\beta} \stackrel{\text{def}}{=} \min\{n \in \mathcal{N}, \beta \in \mathbb{R}_+^N \mid \Psi_n(\beta) = \bar{\mathbf{y}}_n\}$. Thus, when reformulating each EU's optimization problem in the following section, we will additionally constraint $\beta \in B$.

B. Reformulation of the End Users' Optimization Problems

To solve (16), we compute the gradient of the right-hand side expression. We obtain:

$$\frac{\partial \mathbb{E}_\varepsilon[\check{J}_n(\cdot)]}{\partial D_n} = 2\tilde{a}_n(D_n - \theta_n) + 2\beta_n \left(D_n - D_n^a(\hat{\theta}) \right) + 1, \quad (17a)$$

$$\frac{\partial \mathbb{E}_\varepsilon[\check{J}_n(\cdot)]}{\partial G_n^c} = a_n G_n^c + b_n + 2\beta_n \left(G_n^c - G_n^{ca}(\hat{\theta}) \right) - 1, \quad (17b)$$

$$\begin{aligned} \frac{\partial \mathbb{E}_\varepsilon[\check{J}_n(\cdot)]}{\partial \hat{\theta}_n} &= c_n^I \frac{(\hat{\theta}_n - \theta_n)}{V_n} + 2\beta_n \left((A_n^a)_{(1,n)} (D_n^a(\hat{\theta}) - D_n) \right. \\ & \left. + (A_n^a)_{(2,n)} (G_n^{ca}(\hat{\theta}) - G_n^c) \right), \quad (17c) \end{aligned}$$

$$\frac{\partial \mathbb{E}_\varepsilon[\check{J}_n(\cdot)]}{\partial V_n} = -c_n^I \frac{(\hat{\theta}_n - \theta_n)^2}{2V_n^2} + \beta_n K_n. \quad (17d)$$

Set $E_n^a \stackrel{\text{def}}{=} C_n^a - \mathbf{B}_n^a \mathbf{G}^{mc}$ and $F_n^a \stackrel{\text{def}}{=} C_n^a - \mathbf{B}_n^a \mathbf{G}^{mc}$. In a stationary point, (17a)-(17d) vanish, and by combining the above expressions, we find the following expression for $\mathbf{y}_n(\beta)$:

$$\begin{aligned} & \sum_n \left(\frac{\beta_n}{\tilde{a}_n + \beta_n} \mathbf{A}_n^a - \frac{2\beta_n}{a_n + \beta_n} \mathbf{A}_n^a \right) \hat{\theta}(\beta) \\ &= \sum_n \left[\frac{1}{\tilde{a}_n + \beta_n} (\tilde{a}_n \theta_n + \beta_n E_n^a + 1) + \frac{1}{a_n + 2\beta_n} (b_n - 2\beta_n F_n^a + 1) \right. \\ & \quad \left. - G_n^{mc} \right], \quad (18) \end{aligned}$$

and

$$G_n^c(\beta) = \frac{1}{a_n + 2\beta_n} \left[2\beta_n \mathbf{A}_n^a \hat{\theta}(\beta) - b_n + 2\beta_n F_n^a + 1 \right], \quad (19a)$$

$$\begin{aligned} D_n(\beta) &= \frac{1}{2\tilde{a}_n A_n^a} \left[-\sqrt{2\beta_n K_n} + A_n^a \right. \\ & \quad \left. + A_n^a (2,n) (1 - b_n) - a_n A_n^a \right] \frac{1}{a_n + 2\beta_n} \left[2\beta_n \mathbf{A}_n^a \hat{\theta}(\beta) \right. \\ & \quad \left. - b_n + 2\beta_n F_n^a + 1 \right], \quad (19b) \end{aligned}$$

$$V_n(\beta) = \frac{1}{\sqrt{2\beta_n K_n}} \left(\hat{\theta}_n(\beta) - \theta_n \right). \quad (19c)$$

We introduce now the following result which is classical in the optimal contract literature.

Proposition 5: Let $\mathbf{z} = (\gamma_n, \beta_n)_n$ be a solution of the DA's bilevel optimization problem. Then, for each EU $n \in \mathcal{N}$

$$\mathbb{E}_\varepsilon \left[\check{J}_n(\mathbf{z}) \right] = p^r D_n(\beta). \quad (20)$$

In other words, the IR constraint is always binding in equilibrium.

Proof: We reason by contradiction. Suppose that there exists an equilibrium \mathbf{z} in which the IR constraint is non binding for some EU n . From (15d), this implies $p^r D_n^{\text{NE}}(\beta) > \mathbb{E}_\varepsilon \left[\check{J}_n \left(\mathbf{p}_n^a(\gamma_n, \beta_n, \mathbf{y}_n^{\text{NE}}(\beta)), \mathbf{y}_n^{\text{NE}}(\beta) \right) \right]$. Then, the DA can unilaterally decrease its expected cost $L^a(\mathbf{p}^a, \mathbf{y}^{\text{NE}}(\beta)) = \pi \sum_n \gamma_n + (1 - \pi \sum_n \beta_n) \mathbb{E}_\varepsilon \left[(D_n^{\text{NE}}(\beta) - D_n^a)^2 + (G_n^{\text{NE}}(\beta) - G_n^{ca})^2 \right]$ in (15a) by decreasing γ_n in its contract payment $\mathbf{p}^a(\cdot)$ without causing any of the constraints to be violated. Indeed, constraints (15b) and (15c) depend only on β relying on the form of the induced equilibrium in (19)-(20); while (15d)-(15f) are checked by definition. This contradicts the assumption that \mathbf{z} is a solution of the DA's bilevel optimization problem. ■

According to Prop. 5, (20) can be interpreted as an implicit constraint that will appear both in the DA's objective function and non-negativity constraints. Using (20), the DA's objective function takes the following closed form: $L^a(\mathbf{p}^a, \mathbf{y}(\beta)) = \sum_n \frac{\gamma_n}{\beta_n} + \sum_n \left(\pi - \frac{1}{\beta_n} \right) \left[j_n(\mathbf{x}_n(\beta), \theta_n) + c_n^I \frac{(\hat{\theta}_n(\beta) - \theta_n)^2}{2V_n(\beta)} - p^r D_n(\beta) \right]$. Using (20), the non-negativity constraints become:

$$j_n(\mathbf{x}_n(\beta), \theta_n) + c_n^I \frac{(\hat{\theta}_n(\beta) - \theta_n)^2}{2V_n(\beta)} - p^r D_n(\beta) \geq 0, \forall n \in \mathcal{N}.$$

Substituting the previous expressions, the DA's bilevel optimization can then be reformulated as follows:

$$\min_{\gamma, \beta} \left\{ \sum_n \frac{\gamma_n}{\beta_n} + \sum_n \left(\pi - \frac{1}{\beta_n} \right) \left[j_n(\mathbf{x}_n^{\text{NE}}(\beta), \theta_n) + c_n^I \frac{(\hat{\theta}_n^{\text{NE}}(\beta) - \theta_n)^2}{2V_n^{\text{NE}}(\beta)} \right. \right. \quad (21a)$$

$$\left. \left. - p^r D_n^{\text{NE}}(\beta) \right] \right\}, \quad (21b)$$

$$\text{s.t. } \mathbf{y}_n^{\text{NE}}(\beta) \text{ checks (18), (19a)-(19c), } \forall n \in \mathcal{N}, \quad (21c)$$

$$j_n(\mathbf{x}_n^{\text{NE}}(\beta), \theta_n) + c_n^I \frac{(\hat{\theta}_n^{\text{NE}}(\beta) - \theta_n)^2}{2V_n^{\text{NE}}(\beta)} - p^r D_n^{\text{NE}}(\beta) \geq 0, \quad (21d)$$

$$\forall n \in \mathcal{N}, \quad (21d)$$

$$\sum_n (D_n^{\text{NE}}(\beta) - G_n^{\text{NE}}(\beta) - G_n^{mc}) = 0, \quad (21e)$$

$$\Psi_n(\beta) \in \mathcal{Y}_n, \gamma_n \geq 0, \forall n \in \mathcal{N}. \quad (21f)$$

The DA's problem (21) is nonlinear nonconvex in β , and linear in γ . From (18), the bias of one EU can be compensated by the bias of another EU because it couples the EUs' biases. This implies that the map $\Psi_n(\cdot)$ is not necessarily bijective. Therefore, we cannot guarantee that the DA's bilevel optimization problem admits a unique solution when a NE is considered at the lower level in (21c).

Proposition 6: If the EUs reach a PSE in (21c) instead of a NE, then the DA's bilevel optimization problem admits a unique equilibrium $z^{\text{PSE}} = (\beta^{\text{PSE}}, \gamma^{\text{PSE}})$ if and only if $A_n^a \theta + B_n^a G^{c\text{PSE}}(\beta) > C_n^a, \forall n$.

The equivalence condition in Prop. 6 only concerns the controllable generation. To guarantee that it holds, we need to find a solution β of the bilevel problem such that $B_n^a G^{c\text{PSE}}(\beta) > C_n^a - A_n^a \theta$. This inequality can be checked *a posteriori* after the bilevel optimization problem has been solved in β, γ to determine if the solution is indeed unique.

The DA's bilevel optimization problem being linear increasing in each component of γ and unconstrained on \mathbb{R}_+^N , the DA's objective is minimized in $\gamma = \mathbf{0}$. In β , the problem is more difficult as it is nonlinear nonconvex. To compute a solution, we rely on simulated annealing. Numerical experiments illustrate the results in Sec. VII-E.

V. ALGORITHMS

Going back to the definition of $\mathcal{G}^{\text{stoch}}$ involving steps a) and b), we notice that the stochastic game can be split into two distinct parametrized optimization problems for each EU $n \in \mathcal{N}$. Taking $\hat{\theta}, \mathbf{V}$ as parameters, we introduce the game $\check{\mathcal{G}}^{\text{stoch}}(\hat{\theta}, \mathbf{V})$, where each EU $n \in \mathcal{N}$ solves:

$$\min_{\mathbf{x}_n} \mathbb{E}_{\varepsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{V})} \left[\check{J}_n(\mathbf{x}_n, \hat{\theta}, \mathbf{V}, \varepsilon) \right], \quad (22a)$$

$$\text{s.t. } \underline{D}_n \leq D_n \leq \overline{D}_n, \quad (22b)$$

$$\underline{G}_n^c \leq G_n^c \leq \overline{G}_n^c, \quad (22c)$$

$$D_n = G_n^c + G_n^{mc} + Q_n, \quad (22d)$$

$$\sum_m Q_m = 0. \quad (22e)$$

Thus, \mathbf{x}_n and \mathbf{v}_n can be updated simultaneously by each EU n , or sequentially, leading to two algorithms: **Dist-RSGM** and **Seq-RSGM**, that we compare against benchmark algorithms for NE and social optimum computation.

A. Nash Equilibrium and Social Optimum Computation

To compute the NE, we formulate (8) as a deterministic optimization problem, by making the expectation in the objective function explicit. This is possible because the $(\varepsilon_n)_n$ are independent and $\varepsilon_n \sim \mathcal{N}(0, V_n), \forall n$ under Asm. 4. For the social optimum (SO), we sum these deterministic objective functions over all the EUs. To compute the NE and the SO, we apply a classic Gradient Descent Method (GDM) to compute the SO and a distributed GDM to compute the NE. The drawback of this approach is that EU n 's expected utility function depends on $\mathbf{v}_{-n} = (\hat{\theta}_{-n}, \mathbf{V}_{-n}), \forall n$. Thus, to implement a distributed GDM, the EUs need to reveal their bias and variance, i.e., their hidden ability.

B. Repeated Stochastic Gradient Methods for Performatively Stable Equilibrium Computation

To avoid having to compute the expectation of the EUs' utility functions in each iteration, we use the Repeated Stochastic Gradient Method (RSGM) that approximates the expectation in each step of gradient with an unbiased estimator [8]. In RSGM, having initialized \mathbf{y} to \mathbf{y}^0 , at each step t , each EU n samples its ε_n^t , and updates its variable \mathbf{y}_n^t using the following formula:

$$\mathbf{y}_n^{t+1} = \text{proj}_{\mathcal{Y}_n} \left(\mathbf{y}_n^t - \eta \nabla_{\mathbf{y}_n} \bar{J}_n(\mathbf{y}_n^t, \varepsilon_n^t) \right), \quad (23)$$

where $\eta > 0$ is the RSGM constant stepsize, $\bar{J}_n(\cdot)$ is defined in Equation (12). However, RSGM requires that each EU n shares its decision variables \mathbf{y}_n with the other EUs at each iteration. But, in $\mathcal{G}^{\text{stoch}}$, at iteration t , the EUs are only aware of their own decision variables \mathbf{y}_n^t , and of $\tilde{\theta}^t$. To overcome this issue, we propose a distributed variant of the algorithm, **Dist-RSGM**, in which at each time step t , each EU n samples its $\hat{\theta}_n^t$ and shares the sampled reading with the DA, which then broadcasts it to the other EUs. In fact, in our game $J_n(\mathbf{y}_n^t, \varepsilon^t)$ coincides with $J_n(\mathbf{x}_n^t, \hat{\theta}_n^t, V_n^t, \tilde{\theta}^t)$. As a result, each EU n can update its decision variables following the rule:

$$\mathbf{y}_n^{t+1} = \text{proj}_{\mathcal{Y}_n} \left(\mathbf{y}_n^t - \eta \nabla_{\mathbf{y}_n} J_n(\mathbf{x}_n^t, \hat{\theta}_n^t, V_n^t, \tilde{\theta}^t) \right). \quad (24)$$

The **Dist-RSGM** is described in Alg. 1.

Algorithm 1 Distributed Repeated Stochastic Gradient Method (Dist-RSGM)

- 1: Initialize $\mathbf{y}^0 \in \mathcal{Y}$.
 - 2: For iteration $t = 0, \dots, T$
 - a: Sample $\hat{\theta}_n^t \sim \mathcal{N}(\hat{\theta}_n^t, V_n^t), \forall n$.
 - b: Find $\mathbf{y}_n^{t+1} = \text{proj}_{\mathcal{Y}_n} \left(\mathbf{y}_n^t - \eta \nabla_{\mathbf{y}_n} J_n(\mathbf{x}_n^t, \hat{\theta}_n^t, V_n^t, \tilde{\theta}^t) \right), \forall n$.
-

In **Dist-RSGM**, the updates of the bias, variance ($\hat{\theta}_n^t, V_n^t, \forall n$), demand and controllable generation ($\mathbf{x}_n^t, \forall n$) is done simultaneously by all the EUs. However, it might be desirable to split $\mathcal{G}^{\text{stoch}}$ into two levels, thus allowing a sequential formulation. Thus, the stochastic gradients of both problems may be required to be updated at the same time because of the sequential structure of the game. To address these issues, we introduce another variant of RSGM: **Seq-RSGM**, which is described in Alg. 2.

Algorithm 2 Sequential Repeated Stochastic Gradient Method (Seq-RSGM)

- 1: Find $\mathbf{x}^{\mathcal{G}} \in \text{SOL}(\mathcal{G})$ and set $\mathbf{x}^0 = \mathbf{x}^{\mathcal{G}}$. Initialize $\mathbf{v}_n^0 \in \mathcal{U}_n, \forall n$.
 - 2: At each time step $t = 0, \dots, T$
 - a: Sample $\tilde{\theta}_n^t \sim \mathcal{N}(\hat{\theta}_n^t, V_n^t), \forall n$, broadcast $\tilde{\theta}^t$ to all the EUs.
 - b: Find $\mathbf{x}_n^{t+1} = \text{argmin}_{\mathbf{x}_n \in \mathcal{X}_n} \mathbb{E}_{\varepsilon} [J_n(\mathbf{x}_n, \mathbf{v}_n^t, \tilde{\theta}^t)], \forall n$, using an exact method.
 - c: Find $\mathbf{v}_n^{t+1} = \text{proj}_{\mathcal{U}_n} (\mathbf{v}_n^t - \bar{\eta} \nabla_{\mathbf{v}_n} J_n(\mathbf{x}_n^{t+1}, \mathbf{v}_n^t, \tilde{\theta}^t)), \forall n$.
-

Because of the GPG structure of $\check{\mathcal{G}}(\hat{\theta}, \mathbf{V})$ proved in Prop. 6, step 2.b can be solved by an exact algorithm from a nonlinear optimization solver.

C. Two-Timescale Stochastic Method for Stackelberg Equilibrium Computation

$\mathcal{G}^{\text{stoch}}$ might be interpreted as a Stackelberg game [37] with step b) giving rise to the inner problem, while step a) is referred as the outer problem. We let $\mathbf{v}_n^{\text{SE}}(\mathbf{x})$, $\forall n$ denote a solution of the inner loop, i.e., EU n 's reaction function. Many algorithms involving a double-loop structure are proposed in the literature to compute stationary solutions. Since we are in a setting where the inner and outer problems are required to be updated at the same time and $\mathbf{v}_n^{\text{SE}}(\mathbf{x})$, $\forall n$ can be characterized analytically, we use a single-loop stochastic algorithm which updates both outer and inner solutions simultaneously, by using estimates of the stochastic gradient of the outer utility function [38].

We proved the strong convexity of $J_n(\cdot)$ w.r.t. \mathbf{y}_n , $\forall n$. Therefore $J_n(\cdot)$ is also strongly convex w.r.t. \mathbf{v}_n , $\forall n$. As a result, for fixed $(\mathbf{x}_n)_n$, the inner problem admits a single interior point solution defined as the reaction function $\mathbf{v}_n^{\text{SE}}(\mathbf{x})$, $\forall n$. Furthermore, assuming $\underline{V}_n < \sqrt{\frac{c_n^2 \alpha_n}{2\beta_n K_n}} \leq \bar{V}_n$, constraints (8e), (8f) are non binding, $\forall n$ (see Supplementary material). Hence, $\det(\nabla_{\mathbf{v}_n, \mathbf{v}_n}^2 J_n(\mathbf{x}_n, \mathbf{v}_n, \varepsilon)) \neq 0$, $\forall \mathbf{v}_n$. Thus, the inverse function theorem can be used and we can rely on a surrogate for $\nabla_{\mathbf{x}_n} J_n(\mathbf{x}_n, \mathbf{v}_n^{\text{SE}}(\mathbf{x}), \varepsilon)$, e.g., $\forall n$, $\bar{\nabla}_{\mathbf{x}_n} J_n(\mathbf{x}_n, \mathbf{v}_n, \varepsilon) = \nabla_{\mathbf{x}_n} J_n(\mathbf{x}_n, \mathbf{v}_n, \varepsilon) - \nabla_{\mathbf{x}_n, \mathbf{v}_n}^2 J_n(\mathbf{x}_n, \mathbf{v}_n, \varepsilon) \left[\nabla_{\mathbf{v}_n, \mathbf{v}_n}^2 J_n(\mathbf{x}_n, \mathbf{v}_n, \varepsilon) \right]^{-1} \nabla_{\mathbf{v}_n} J_n(\mathbf{x}_n, \mathbf{v}_n, \varepsilon)$. Details about the analytical computation of the surrogate can be found in Appendix B.

Algorithm 3 Two-Timescale Stochastic Method (TTSM)

- 1: Find $\mathbf{x}^{\mathcal{G}} \in \text{SOL}(\mathcal{G})$ and set $\mathbf{x}^0 = \mathbf{x}^{\mathcal{G}}$. Initialize $\mathbf{v}_n^0 \in \mathcal{U}_n$, $\forall n$.
- 2: At each time step $t = 0, \dots, T$
 - a: Find $\mathbf{v}_n^{t+1} = \text{proj}_{\mathcal{U}_n}(\mathbf{v}_n^t - \eta_1^t \nabla_{\mathbf{v}_n} J_n(\mathbf{x}_n^t, \mathbf{v}_n^t, \tilde{\boldsymbol{\theta}}^t))$, $\forall n$,
 - b: Sample $\tilde{\boldsymbol{\theta}}^{t+1} \sim \mathcal{N}(\hat{\boldsymbol{\theta}}^{t+1}, \mathbf{V}_n^{t+1})$, $\forall n$, broadcast $\tilde{\boldsymbol{\theta}}^{t+1}$ to all the EUs.
 - c: Find

$$\mathbf{x}_n^{t+1} = \text{proj}_{\mathcal{X}_n}(\mathbf{x}_n^t - \eta_2^t \bar{\nabla}_{\mathbf{x}_n} J_n(\mathbf{x}_n^t, \mathbf{v}_n^{t+1}, \tilde{\boldsymbol{\theta}}^{t+1})), \forall n.$$

VI. CONVERGENCE ANALYSIS

We derive closed-form expressions of each algorithm rate of convergence; insights are given regarding the tuning of the stepsize descents and iteration complexity.

A. Preliminary Result

Let us first introduce an intermediate result that will be used later, to prove the convergence of Seq-RSGM.

Proposition 7: $\check{\mathcal{G}}^{\text{stoch}}(\hat{\boldsymbol{\theta}}, \mathbf{V})$ is a generalized potential game (GPG), $\forall \hat{\boldsymbol{\theta}}, \mathbf{V}$.

Proof: Relying on the GPG definition [33, Sec. 2], this result follows from the coincidence of the KKTs of $\check{\mathcal{G}}^{\text{stoch}}(\hat{\boldsymbol{\theta}}, \mathbf{V})$ with the KKTs of the minimization problem with $\sum_n \mathbb{E}_{\varepsilon} [\check{J}_n(\mathbf{x}_n, \hat{\boldsymbol{\theta}}, \mathbf{V}, \varepsilon)]$ in the objective, under the concatenation of the constraints (22b)-(22d), $\forall n$, and (22e). ■

From Prop. 7, $\check{\mathcal{G}}^{\text{stoch}}(\hat{\boldsymbol{\theta}}, \mathbf{V})$ can be equivalently formulated as an optimization problem, and its VE computed through exact mathematical optimization algorithms.

B. Convergence Proofs for Dist-RSGM and Seq-RSGM

We introduce generic notation for any $\mathbf{z} \in \{\mathbf{y}, \mathbf{x}, \mathbf{v}\}$ and $\mathbf{z}_n \in \{\mathbf{y}_n, \mathbf{x}_n, \mathbf{v}_n\}$, $g_{\mathbf{z}}(\mathbf{y}, \varepsilon) \stackrel{\text{def}}{=} (\nabla_{\mathbf{z}_n} J_n(\mathbf{y}, \varepsilon))_{n \in \mathcal{N}}$, $g_{\mathbf{z}_n}(\mathbf{y}, \varepsilon) = \nabla_{\mathbf{z}_n} J_n(\mathbf{y}, \varepsilon)$, $G_{\mathbf{z}}(\mathbf{y}) \stackrel{\text{def}}{=} \mathbb{E}_{\varepsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{V})} [g_{\mathbf{z}}(\mathbf{y}, \varepsilon)]$ and $G_{\mathbf{z}_n}(\mathbf{y}) \stackrel{\text{def}}{=} \mathbb{E}_{\varepsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{V})} [g_{\mathbf{z}_n}(\mathbf{y}, \varepsilon)]$, $\forall n$.

As a benchmark, GDM converges linearly to the NE of $\mathcal{G}^{\text{stoch}}$, requiring $O(\frac{1}{\delta})$ iterations to achieve $\|G_{\mathbf{y}^t}(\mathbf{y}^t)\|^2 \leq \delta$. We compare GDM with Dist and Seq-RSGM on their convergence rate and iteration complexity.

Theorem 4: Let $\rho \in]0, 1[$. Then, for any stepsize $\eta < \frac{\omega(1-\rho)}{8L^2}$, Dist-RSGM generates a sequence $(\mathbf{y}^t)_t$ satisfying

$$\mathbb{E}_t \|\mathbf{y}^{t+1} - \mathbf{y}^{\text{PSE}}\|^2 \leq \frac{1 + 2\eta\omega\rho + 2\eta^2\omega^2\rho^2}{1 + 2\eta\omega(\frac{1+\rho}{2})} \mathbb{E}_t \|\mathbf{y}^t - \mathbf{y}^{\text{PSE}}\|^2 + \frac{4\eta^2\sigma^2}{1 + 2\eta\omega(\frac{1+\rho}{2})},$$

where \mathbf{y}^{PSE} is the PSE of $\mathcal{G}^{\text{stoch}}$.

Proof: We observe that $g_{\mathbf{y}_n}(\cdot)$ is linear in ε , hence it is κ_n -Lipschitz in ε . Now, we want to obtain the L -Lipschitz continuity of $G_{\mathbf{y}}(\cdot)$ in \mathbf{y} . To do so, we observe that $G_{\mathbf{y}_n}(\mathbf{y}) - G_{\mathbf{y}'_n}(\mathbf{y}') = F_{\mathbf{y}_n}(\mathbf{y}) - F_{\mathbf{y}'_n}(\mathbf{y}')$. We proceed the same way as we did to prove the strong monotonicity of the objective function, and sum over all the EUs. We obtain $L = N \max_{\mathbf{y}_n \in \mathcal{Y}_n} \max_{i \in [1, m_n + 2]} \lambda_i(\mathbf{H}_{\mathbf{y}_n})$.

We check a property on the regularity of the distribution used. Using [39, Lem. 2.4], we obtain an upper bound on the Wasserstein-1 distance between two Gaussian distributions:

$$\mathbb{W}_1(\mathcal{N}(\mathbf{0}, \mathbf{V}), \mathcal{N}(\mathbf{0}, \mathbf{V}')) \leq \sqrt{\sum_{i=1}^N (\sqrt{V_i} - \sqrt{V'_i})^2}, \quad (25a)$$

$$\leq \sqrt{\sum_{i=1}^N \tau_i^2 (V_i - V'_i)^2} \leq \tau \|\mathbf{V} - \mathbf{V}'\|, \quad (25b)$$

with τ_i obtained by the Mean Value Theorem and $\tau = \max_i \tau_i$. Let $\rho \stackrel{\text{def}}{=} \sqrt{\sum_{i=1}^N (\frac{\kappa_i \tau}{\omega})^2}$. We aim to prove the convergence of Dist-RSGM, a distributed variant of RSGM [8]. First, we need to obtain the following property on the variance of the gradient of $J_n(\cdot)$:

$$\mathbb{E}_t \|g_{\mathbf{y}}(\mathbf{y}, \varepsilon) - \mathbb{E}_t [g_{\mathbf{y}}(\mathbf{y}, \varepsilon)]\|^2 \leq \sigma^2. \quad (26)$$

To do so, we use the explicit expression of the gradient $g_{\mathbf{y}}(\cdot)$ and its linearity in ε . This leads us to the following bound:

$$\sigma^2 = 4 \sum_{n=1}^N \beta_n^2 S_n \bar{V}_n, \quad (27)$$

with S_n a scalar being quadratic in the elements of \mathbf{A}_n^a defined as:

$$S_n = \left(\sum_{m=1}^N (A_{n(1,m)}^a)^2 \right) \left((A_{n(1,n)}^a)^2 + 1 \right) + \left(\sum_{m=1}^N (A_{n(2,m)}^a)^2 \right) \left((A_{n(2,n)}^a)^2 + 1 \right) + A_{n(1,n)}^a A_{n(2,n)}^a \times \left(\sum_{m=1}^N A_{n(1,m)}^a A_{n(2,m)}^a \right).$$

Let \mathbf{y}^{PSE} be the PSE solution of $\mathcal{G}^{\text{stoch}}$. Then, we need to check the next condition on the bias of our estimator of the gradient:

$$\|\mathbb{E}_t g_{\mathbf{y}}(\mathbf{y}, \varepsilon) - G_{\mathbf{y}^{\text{PSE}}}(\mathbf{y})\| \leq P \|\mathbf{y} - \mathbf{y}^{\text{PSE}}\|. \quad (28)$$

To do so, we decompose the norm in (28) over each EU n . We use the κ_n -Lipschitz continuity of $g_{\mathbf{y}_n}(\cdot)$ in ε and the bound on the Wasserstein-1 distance between two Gaussian density functions. This gives us:

$$\|\mathbb{E}_t g_{\mathbf{y}_n}(\mathbf{y}, \varepsilon) - G_{\mathbf{y}^{\text{PSE}}, n}(\mathbf{y})\| \quad (29a)$$

$$\leq \kappa_n \sup_{h \in \text{Lip}_1} \left[\mathbb{E}_{\varepsilon \sim \mathcal{N}(\mathbf{0}, \mathcal{V})} [h(\mathbf{y}, \varepsilon)] - \mathbb{E}_{\varepsilon' \sim \mathcal{N}(\mathbf{0}, \mathcal{V}^{\text{PSE}})} [h(\mathbf{y}, \varepsilon')] \right], \quad (29b)$$

$$\leq \kappa_n \tau \|\mathbf{V} - \mathbf{V}^{\text{PSE}}\|. \quad (29c)$$

To go from equation (29b) to (29c), we use the conditions obtained in (25a)-(25b), because this sup is the definition of the Wasserstein-1 distance.

Now that we have this inequality, we square it, sum it over all the EUs, major $\|\mathbf{V} - \mathbf{V}^{\text{PSE}}\|$ by $\|\mathbf{y} - \mathbf{y}^{\text{PSE}}\|$, and apply a square root on the inequality. We obtain:

$$\|\mathbb{E}_t g_{\mathbf{y}}(\mathbf{y}, \varepsilon) - G_{\mathbf{y}^{\text{PSE}}}(\mathbf{y})\| \leq \sqrt{\sum_{n=1}^N (\kappa_n \tau)^2} \|\mathbf{y} - \mathbf{y}^{\text{PSE}}\|. \quad (30)$$

We obtain $P = \omega \rho$. Now that we have proved those two inequalities, we can apply [8, Thm. 8], and derive the statement of Thm. 4. \blacksquare

Note that we can approach even closer \mathbf{y}^{PSE} by reducing the stepsize η , but Dist-RSGM will converge slowly. Thus, there is a trade-off to find between the accuracy of the PSE computation and the convergence speed. We focus on a step-decay schedule that drops the stepsize by a factor every epoch. We set $\eta_0 \stackrel{\text{def}}{=} \frac{\omega(1-\rho)}{4} \min\{1, \frac{1}{2L^2}\}$, and run Dist-RSGM in $k = 0, \dots, k_{\max}$ epochs for T_k iterations each, with constant stepsize $\eta_k = 2^{-k} \eta_0$; the last iterate of epoch k being used as the first iterate in epoch $k+1$. For target accuracy $\delta > 0$ defined such that $\mathbb{E}[\|\mathbf{y}^{T_{k_{\max}}} - \mathbf{y}^{\text{PSE}}\|^2] \leq \delta$, $\rho < \frac{1}{2}$, and a constant $R \geq \|\mathbf{y}^0 - \mathbf{y}^{\text{PSE}}\|^2$, we set $T_0 = \lceil \frac{10}{(1-\rho)\omega\eta_0} \log(\frac{2R}{\delta}) \rceil$, $T_l = \lceil \frac{10 \log(4)}{(1-\rho)\omega\eta_l} \rceil$, $\forall l \geq 1$, $k_{\max} = \lceil 1 + \log_2(\frac{40\eta_0\sigma^2}{(1-\rho)\omega\delta}) \rceil$, Dist-RSGM converges in at most $O\left(\frac{L^2}{(1-\rho)\omega^2} \log(\frac{2R}{\delta}) + \frac{\sigma^2}{(1-\rho)^2\omega^2\delta}\right)$ iterations [8, Cor. 1].

Remark 6.1: Since \mathbf{y}^{PSE} is initially unknown, T_0 can be determined based on the desired accuracy δ and the diameter of the set \mathcal{Y} .

Proposition 8: Seq-RSGM converges with the same convergence rate and iteration complexity as Dist-RSGM.

Proof: From Prop. 7, $\check{\mathcal{G}}(\hat{\theta}, \mathbf{V})$ is a GPG, and, thus, can be solved as an optimization problem. As a result, we can find \mathbf{x}^{t+1} based on an exact method. Then, Thm. 4 convergence rate also applied to Dist-RSGM, assuming $\mathbf{y}^{t+1} = \text{col}(\mathbf{0}, \mathbf{v}^{t+1})$, $\forall t$. \blacksquare

C. Convergence Proof for TTSM

Let $\mathbf{x}^{\text{SE}}, \mathbf{v}^{\text{SE}}(\mathbf{x}^{\text{SE}})$ be a solution of $\mathcal{G}^{\text{stoch}}$ interpreted as a Stackelberg game. The optimality gap is defined as $\Delta_x^t \stackrel{\text{def}}{=} \mathbb{E}_t[\|\mathbf{x}^t - \mathbf{x}^{\text{SE}}\|^2]$ and the tracking error, as $\Delta_v^t \stackrel{\text{def}}{=} \mathbb{E}_t[\|\mathbf{v}^t - \mathbf{v}^{\text{SE}}(\mathbf{x}^{t-1})\|^2]$.

Proposition 9: $\mathbb{E}_{\varepsilon} [J_n(\mathbf{x}_n, \mathbf{v}_n^{\text{SE}}(\mathbf{x}), \varepsilon)]$ is strongly convex with a modulus $\mu_l > 0$, $\forall n$.

Proof: Using (17c) and (17d), we infer that there exists $\mathbf{U} > \mathbf{0}$ such that $\hat{\theta} = \mathbf{U}^{-1} \text{col}\left(\left(A_n^a(1,n)D_n + A_n^a(2,n)G_n^c - \frac{K_n}{2\beta_n c_n^I}\right)_n\right)$ where $\mathbf{U}_{(n,:)} = A_n^a(1,n)\mathbf{A}_n^a(1,:) + A_n^a(2,n)\mathbf{A}_n^a(2,:)$. This expression can be rewritten in a more compact way as $\hat{\theta} = \underbrace{\begin{bmatrix} \mathbf{U}^{-1}\mathbf{X}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{U}^{-1}\mathbf{X}_2 \end{bmatrix}}_{\mathbf{M}} \begin{bmatrix} \mathbf{D} \\ \mathbf{G}^c \end{bmatrix} + \mathbf{X}_3$ with

$$\mathbf{X}_n \stackrel{\text{def}}{=} \begin{bmatrix} A_1^a(1,n) & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & A_N^a(n,N) \end{bmatrix}, \quad \forall n \in \{1, 2\} \text{ and } \mathbf{X}_3 \text{ a}$$

vector of constants. Applying twice the chain rule, we infer that $\mathbf{H}_{\mathbf{x}_n} = \nabla_{\mathbf{x}_n}^2 \mathbb{E}_{\varepsilon} [J_n(\mathbf{x}_n, \mathbf{v}_n^{\text{SE}}(\mathbf{x}), \varepsilon)] = (\mathbf{M}^{\top})^{-1} \nabla_{\mathbf{x}_n}^2 \mathbb{E}_{\varepsilon} [J_n(\mathbf{x}_n, \mathbf{x}, \varepsilon)] \mathbf{M}^{-1} \geq \tilde{\omega} (\mathbf{M}^{\top})^{-1} \mathbf{M}^{-1}$ with $\tilde{\omega} > 0$ since, from the proof of Prop. 3, $\mathbb{E}_{\varepsilon} [J_n(\cdot)]$ is strongly monotone in \mathbf{y}_n and therefore in \mathbf{x}_n . Since $\mathbf{M} > \mathbf{0}$, $\mathbf{M}^{\top} > \mathbf{0}$, $(\mathbf{M}^{\top})^{-1} > \mathbf{0}$ and $\mathbf{M}^{-1} > \mathbf{0}$. Thus, there exists $\mu_l > 0$ such that $\nabla_{\mathbf{x}_n}^2 \mathbb{E}_{\varepsilon} [J_n(\mathbf{x}_n, \mathbf{v}_n^{\text{SE}}(\mathbf{x}), \varepsilon)] \geq \mu_l \mathbf{I}$. \blacksquare

Proposition 10: Suppose $\bar{V}_n \leq \frac{c_n^I}{2\beta_n K_n}$ hold. Then, $\mathbb{E}_{\varepsilon} [J_n(\mathbf{x}_n, \cdot, \varepsilon)]$ is strongly convex w.r.t. \mathbf{v}_n with a modulus μ_g , $\forall n$.

Proof: We compute the partial Hessian $\mathbf{H}_{\mathbf{v}_n}$ of $\mathbb{E}_{\varepsilon} [J_n(\cdot)]$ w.r.t. \mathbf{v}_n , $\mathbf{H}_{\mathbf{v}_n} = \begin{bmatrix} \frac{c_n^I}{\bar{V}_n} + 2\beta_n K_n - \frac{c_n^I(\hat{\theta}_n - \theta_n)}{V_n^2} \\ \frac{c_n^I(\hat{\theta}_n - \theta_n)}{V_n^2} & \frac{c_n^I(\hat{\theta}_n - \theta_n)^2}{2V_n^3} \end{bmatrix}$. From this 2×2 matrix, we can extract the trace as a closed-form expression $\text{tr}(\mathbf{H}_{\mathbf{v}_n}) = \frac{c_n^I}{\bar{V}_n} \left(1 + \frac{(\hat{\theta}_n - \theta_n)^2}{2V_n^2}\right) + 2\beta_n K_n > 0$, and the determinant $\det(\mathbf{H}_{\mathbf{v}_n}) = \frac{(\hat{\theta}_n - \theta_n)^2}{V_n^3} c_n^I [\beta_n K_n - \frac{c_n^I}{2\bar{V}_n}] > 0$ since $\bar{V}_n \leq \frac{c_n^I}{2\beta_n K_n}$ by assumption. Using the classical relationships between $\text{tr}(\cdot)$, $\det(\cdot)$, and the matrix eigenvalues, we infer that $\mu_g = \min_n \min_{\mathbf{v}_n \in \mathcal{U}_n} \frac{1}{2} \left(\text{tr}(\mathbf{H}_{\mathbf{v}_n}) - \sqrt{\text{tr}(\mathbf{H}_{\mathbf{v}_n})^2 - 4 \det(\mathbf{H}_{\mathbf{v}_n})} \right)$. \blacksquare

The following assumption guarantees Prop. 10 holds and the lower-level constraints (8e), (8f) are non binding.

Assumption 6: Suppose that constants are chosen such that $\bar{V}_n < \sqrt{\frac{c_n^I \alpha_n}{2\beta_n K_n}}$ and $\sqrt{\frac{c_n^I \alpha_n}{2\beta_n K_n}} \leq \bar{V}_n \leq \frac{c_n^I}{2\beta_n K_n}$.

We now introduce conditions on $g_{\mathbf{v}_n}$ which are standard when the latter is taken as a stochastic gradient of J_n w.r.t. \mathbf{v}_n , while $g_{\mathbf{x}_n}$ is a potentially biased estimate of $\mathbb{E}_{\varepsilon} [\nabla_{\mathbf{x}_n} J_n]$.

Assumption 7: There exist constants σ_g , σ_f , and nonincreasing sequence $\{b_n^t\}_t$ such that $\forall n$, $\mathbb{E}_t \left[\|g_{\mathbf{v}_n}(\mathbf{x}_n^t, \mathbf{v}_n^t, \varepsilon^t) - G_{\mathbf{v}_n}(\mathbf{x}_n^t, \mathbf{v}_n^t, \varepsilon^t)\|^2 \right] \leq \sigma_g^2 \left\{ 1 + \|G_{\mathbf{v}_n}(\mathbf{x}_n^t, \mathbf{v}_n^t, \varepsilon^t)\|^2 \right\}$, and $\mathbb{E}_{t+1} \left[\|g_{\mathbf{x}_n}(\mathbf{x}_n^t, \mathbf{v}_n^{t+1}, \varepsilon^{t+1}) - b_n^t - \mathbb{E}_{\varepsilon} [\nabla_{\mathbf{x}_n} J_n(\mathbf{x}_n^t, \mathbf{v}_n^{t+1}, \varepsilon^{t+1})]\|^2 \right] \leq \sigma_f^2$.

Set $\tilde{\sigma}_f^2 \stackrel{\text{def}}{=} \sigma_f^2 + 3 \sup_{n, \mathbf{x}_n \in \mathcal{X}_n} \|\nabla_{\mathbf{x}_n} \mathbb{E}_{\varepsilon} [J_n(\mathbf{x}_n, \mathbf{v}_n^{\text{SE}}(\mathbf{x}), \varepsilon)]\|^2$.

Theorem 5: Let **Asms. 6 and 7** hold. Then, for any stepsizes satisfying $\eta_2^t \leq c_0(\eta_1^t)^{3/2}$, $\eta_1^t \leq c_1(\eta_2^t)^{2/3}$, $\frac{\eta_1^{t-1}}{\eta_1^t} \leq 1 + \eta_1^t \frac{\mu_g}{8}$, $\frac{\eta_2^{t-1}}{\eta_2^t} \leq 1 + 3\eta_2^t \frac{\mu_l}{4}$, $\eta_2^t \leq \frac{1}{\mu_l}$, $\eta_1^t \leq \min \left\{ \frac{1}{\mu_g}, \frac{\mu_g}{L_y^2(1+\sigma_f^2)}, \frac{\mu_g^2}{48c_0^2L^2L_y^2} \right\}$, $8\mu_l\eta_2^t \leq \mu_g\eta_1^t$, $\forall t$, where c_0, c_1 are free parameters, TTSM generates a sequence $(\mathbf{x}^t, \mathbf{v}^{\text{SE}}(\mathbf{x}^t))$ satisfying

$$\Delta_x^t \leq \prod_{k=0}^{t-1} (1 - \eta_2^k \mu_l) (\Delta_x^0 + \frac{L^2}{\mu_l^2} \Delta_y^0) + \frac{c_1 L^2}{\mu_l^2} (\frac{\sigma_g^2}{\mu_g} + \frac{c_0^2 L_y^2}{\mu_g^2} \tilde{\sigma}_f^2) (\eta_2^{t-1})^{2/3},$$

$$\Delta_v^t \leq \prod_{k=0}^{t-1} (1 - \eta_1^k \mu_g / 4) \Delta_y^0 + (\frac{\sigma_g^2}{\mu_g} + \frac{c_0^2 L_y^2}{\mu_g^2} \tilde{\sigma}_f^2) \eta_1^{t-1}.$$

Proof: We use the same reasoning as in the proof of Thm. 3 to check the L_g -Lipschitz continuity of $G_{v,n}(\mathbf{x}_n, \mathbf{v}_n, \varepsilon)$. Furthermore, by definition, the expected utility functions check Asms. 1–2 from [38]. As a result, from [38, Lem. 2], there exist constants L and L_y such that:

$$\begin{aligned} & \|\bar{\nabla}_{\mathbf{x}_n} \mathbb{E}_\varepsilon [J_n(\mathbf{x}_n, \mathbf{v}_n, \varepsilon)] - \nabla_{\mathbf{x}_n} \mathbb{E}_\varepsilon [J_n(\mathbf{x}_n, \mathbf{v}_n^{\text{SE}}(\mathbf{x}), \varepsilon)]\| \\ & \leq L \|\mathbf{v}_n^{\text{SE}}(\mathbf{x}) - \mathbf{v}_n\|, \forall n, \\ & \|\mathbf{v}_n^{\text{SE}}(\mathbf{x}) - \mathbf{v}_n^{\text{SE}}(\mathbf{x}')\| \leq L_y \|\mathbf{x} - \mathbf{x}'\|, \forall n. \end{aligned}$$

In step 2.a of TTSM, under Assumption 6, the EUs' decision variables $(\hat{\theta}_n, V_n)_n$ always remain within the feasible set, thereby eliminating the need for projection¹. Further, steps 2.a and 2.b can be solved independently by each EU, and the Stackelberg game can be framed as a single leader-single follower game. Thus, [38, Thm. 1] can be applied. ■

To compute a δ -optimal solution with $\Delta_x^t \leq \delta$, TTSM requires a total of $O(\log(1/\delta)/\delta^{3/2})$ calls of stochastic oracles of both outer and inner functions. The polynomial bias decay $\{b_n^t\}_t$ requires using $O(1 + \log(t))$ samples per iteration, justifying the log factor in the bound [38].

VII. SIMULATIONS

A. A Private Network Game Abstraction

We consider a network of EUs, with one EU located in each node of the network, involved in peer-to-peer energy trading. Each EU holds both controllable and non-controllable (e.g., solar PV power) generations, thus producing energy it can use for its own means and/or sell to its neighbors. While renewable generation at each node is measured through a communication interface and subsequently broadcast via a service platform to all the nodes; the generation and consumption costs parameters, as well as nominal demands, are private to the nodes. In this asymmetric information setting, each EU computes its demand and controllable generation. To perform this computation, the EUs contract with a DA which shares readings through a service platform, i.e., randomized nominal demands reported by the EUs. The randomization process is performed for each EU by the node interface, which samples the noise to add to the EU's nominal demand while checking constraints

¹However, without this assumption, step 2.a would require a projection operation. Our numerical simulations demonstrate that TTSM converges even when projection is applied. Nonetheless, a rigorous theoretical proof of convergence in the presence of projection would likely require more in-depth analysis, which we leave for future work.

on the adjacency level and bounds on variance. Finally, the DA contract payment remunerates *a posteriori* the EUs depending on the accuracy of their readings.

Numerical results are presented on randomly generated problem instances with 5, 10 and 15 EUs. For these instances, we choose $(\tilde{a}_n)_n$ as random integer variables on [8, 15], $\tilde{b}_n = 0, \forall n$, $(a_n)_n, (b_n)_n$ as random Uniform variables on [1, 3], $c_n = 1, \forall n$, $(G^{nc})_n$ as random Uniform variables on [0, 1], $\bar{D}_n = 10, \forall n$, $(\underline{D}_n)_n$ as random Uniform variables on [0, 0.5], $\bar{G}_n^c = 3, \forall n$, $\underline{G}_n^c = 0, \forall n$, $\underline{V}_n = 0, \forall n$, $\bar{V}_n = 2, \forall n$, $(\theta_n)_n$ as a random Uniform variables on [0.1, 1], $c_n^l = 3, \forall n$, $\pi = 1$, $p^r = 1.0$, and $\xi = 10^3$.

B. Numerical Comparison of Iteration Complexity

In Fig. 1, we plot the error $\|\hat{\theta} - \theta^*\|$ where θ^* coincides with $\mathcal{G}^{\text{stoch}}$ PSE for RSGM, Dist-RSGM, and Seq-RSGM, and with $\mathcal{G}^{\text{stoch}}$ SE for TTSM resp. For each algorithm, we consider 50 random initializations of \mathbf{y}^0 . We proved in Prop. 5 that the PSE solution of $\mathcal{G}^{\text{stoch}}$ is unique. However, RSGM (for more than 5 EUs) and variants of RSGM only enable to approximate the PSE. This is due to the noise sampling in the EUs' readings and to the fact that we impose an upper bound on the number of iterations.

Comparing the algorithms performance, we observe that:

- For a fixed number of iterations and small number of EUs (less than 5), RSGM learns more accurately the PSE than Dist-RSGM and Seq-RSGM;
- When the number of EUs increases, Seq-RSGM performs better than RSGM, but at the cost of more volatility in the learning process that might be caused by oversampling;
- TTSM is scalable to a large number of EUs with small volatility in the learning process, but slower than RSGM to reach the same accuracy level.

C. Balancing Efficiency and Privacy

Let \mathbf{y}^{so} be the social optimum solution of the optimization problem with objective the sum over n of the expected utilities $SC(\mathbf{y}) = \sum_n \mathbb{E}_{\varepsilon \sim \mathcal{N}(0, \mathcal{V})} [J_n(\mathbf{x}_n, \hat{\theta}, \mathbf{V}, \varepsilon)]$, under the concatenation of the constraints (8b)-(8f) for all the EUs. We evaluate the efficiency gap as the difference between the social cost in the equilibrium \mathbf{y}^* and the optimal social cost: $SC(\mathbf{y}^*) - SC(\mathbf{y}^{\text{so}})$ where \mathbf{y}^* is resp. the NE, PSE and SE of $\mathcal{G}^{\text{stoch}}$. We plot the distribution of the efficiency gap in each solution concept in Fig. 2. We observe that the NE is economically efficient, minimizing the social cost, while PSE and SE generate large efficiency gaps. However, the efficiency of an equilibrium solution concept should be balanced by the amount of information released for the EUs to reach it. The total information released by the EUs on the market, illustrated in Fig. 3, is defined as $\sum_n I_n(\theta_n, \hat{\theta}_n, V_n)$. We observe that increasing the number of EUs, also increases the total information released for the NE computation and, in a lesser extent, for the PSE and SE. Furthermore, while the NE achieves economic efficiency, it is at the cost of large information released. PSE and SE, while being inefficient, requires far less information released.

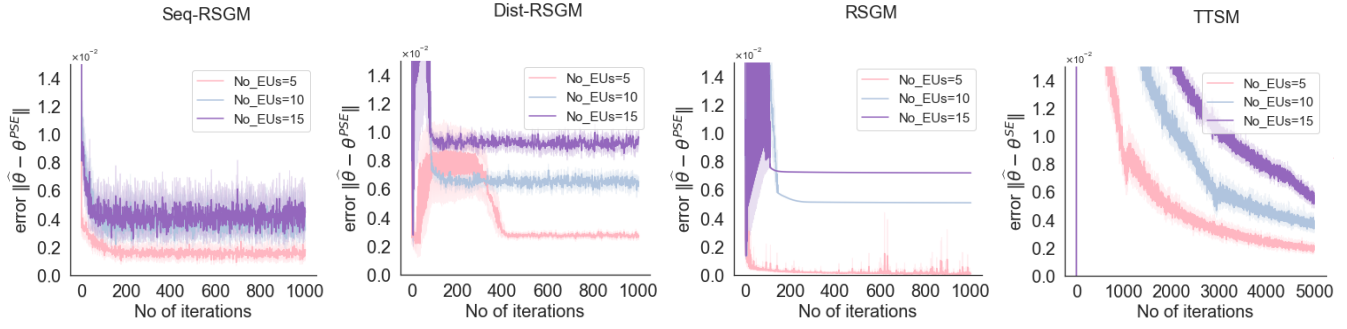


Fig. 1: Iteration complexity of the norm-squared error to the PSE bias for RSGM, Dist, Seq-RSGM, and to the SE bias for TTSM for randomly generated problem instances with 5, 10 and 15 EUs. The empirical rates match the theory in Sec. VI.

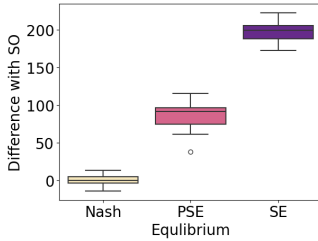


Fig. 2: Box plots of the efficiency gaps in NE, PSE, SE.




Fig. 3: Total information released in NE, PSE, SE.




Fig. 4: Mean and standard deviation of the cost $(j_n)_n$ per EU.

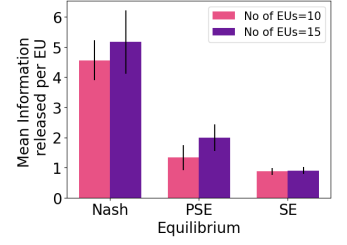


Fig. 5: Mean and standard deviation of the information released $(I_n)_n$ per EU.

D. Benefits of Accounting for Performativity

In Fig. 4-6, we have represented the mean and standard deviation of the cost $(j_n)_n$, information released $(I_n)_n$, and contract payment $(p_n^a)_n$ per EU. We observe that:

- PSE and SE require far less information released from the EUs than NE;
- Dispersion around mean information released among the EUs is tiny under SE while it is big, thus favoring free-rider behaviors, under NE;
- for a large number of EUs (more than 15), costs and contract payments are aligned among the three solution concepts, with large dispersion among the EUs;
- for a small number of EUs (less than 10), costs are reduced under SE, by comparison with NE and PSE. However, this is compensated by low contract payments.

These plots demonstrate that performative effects modeling is beneficial to the EUs by requiring less information released to reach the related equilibrium (PSE, SE) and potentially acting as a barrier to free-rider behaviors.

E. Solving the DA's Bilevel Optimization Problem

In Fig. 7, we plot the DA's Bi-Level (BL) objective in the last sample solution that was selected (last BL objective) and in the optimal sample solution (optimal BL objective). We observe that the simulated annealing algorithm converges in less than $2 \cdot 10^3$ iterations to $\beta^{\text{NE}} = 0.8$, with an initial temperature of 10^3 and an empirically-chosen geometric decrease with factor 0.98.




Fig. 6: Mean and standard deviation of the contract payment $(p_n^a)_n$ per EU.

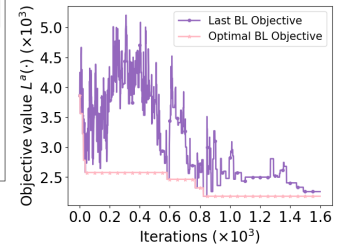


Fig. 7: Simulated annealing algorithm convergence.

VIII. CONCLUSION

We formulate a peer-to-peer electricity market as a private network game, leading to a unique and globally stable VE. To protect their privacy, the EUs might provide biased and noisy readings of their nominal demand. In this asymmetric information setting, a DA is introduced, which aims to learn the EUs' private information, while remunerating them depending on the quality of their readings. To assess the robustness of the market towards distributional shifts caused by the EUs' hidden ability, we formulate a decision-dependent game $\mathcal{G}^{\text{stoch}}$. We rely on three solution concepts (NE, PSE, SE) to interpret $\mathcal{G}^{\text{stoch}}$ outcome, each allowing to capture different degrees of performativity. To compute these equilibria, we propose two variants of RSGM and a two-time-scale stochastic approximation algorithm. We prove the convergence of each algorithm; insights are given regarding the tuning of the stepsizes descents and iteration complexity. The algorithm performance is subsequently compared on small and large scale problem instances, by measuring the

iteration complexity. Solution concepts are compared on the basis of their efficiency and information release, thus highlighting the benefits of performative effects modeling to incentivize the EUs to report truthful readings and penalize free-rider behaviors.

APPENDIX

A. Proofs

Proof of Lemma 1: Let $\nabla_{\mathbf{y}_n} h(\mathbf{y})$ be the partial derivative of $h(\cdot)$ w.r.t. \mathbf{y}_n , $\forall n$. We adapt the proof of [40, Thm. 2.2]. We set $h(\mathbf{y}) \stackrel{\text{def}}{=} (\sum_{m=1}^N Q_m)^2$ and ϕ is the Lagrangian multiplier related to the constraint $h(\mathbf{y}) = 0$. We denote $\hat{\mathbf{y}}$ the point that minimizes $\check{J}_n(\cdot)$ over \mathcal{Y} and satisfies the constraint $h(\hat{\mathbf{y}}) = 0$. For any $\mathbf{y} \in \mathcal{Y}$, we have:

$$\check{J}_n(\mathbf{y}) + \xi h(\mathbf{y}) \geq \check{J}_n(\hat{\mathbf{y}}) + \phi h(\hat{\mathbf{y}}) \quad (32a)$$

$$\geq \check{J}_n(\hat{\mathbf{y}}) + \phi (h(\hat{\mathbf{y}}) + \nabla_{\mathbf{y}_n} h(\hat{\mathbf{y}})^\top (\mathbf{y} - \hat{\mathbf{y}})), \quad (32b)$$

$$= \check{J}_n(\hat{\mathbf{y}}) + \phi \nabla_{\mathbf{y}_n} h(\hat{\mathbf{y}})^\top (\mathbf{y} - \hat{\mathbf{y}}), \quad (32c)$$

$$= \check{J}_n(\mathbf{y}) - \nabla_{\mathbf{y}_n} \check{J}_n(\hat{\mathbf{y}})^\top (\mathbf{y} - \hat{\mathbf{y}}), \quad (32d)$$

$$\geq \check{J}_n(\hat{\mathbf{y}}) = \check{J}_n(\hat{\mathbf{y}}) + \xi h(\hat{\mathbf{y}}). \quad (32e)$$

To go from (32a) to (32b), we use the convexity of $h(\cdot)$. To go from (32c) to (32d), we use the stationarity condition. To go from (32d) to (32e), we use the convexity of $\check{J}_n(\cdot)$, which can be proved by invoking similar arguments as in the proof of the strict convexity of $\mathbb{E}_\varepsilon[J_n(\cdot)]$ in Prop. 3. Hence, $\hat{\mathbf{y}}$ also minimizes the penalized problem.

In addition, it is possible to derive a closed-form expression for ϕ . Under Asms. 1 and 2, EU n optimization problem (4) depends only on D_n, G_n^c since $Q_n = D_n - G_n^c - G_n^{nc}$. Problem (4) KKT conditions give the following expressions at the VE: $D_n^{\text{VE}} = \theta_n - \frac{1}{2\bar{a}_n} (1 + (\bar{\mu}_n - \underline{\mu}_n) + \phi)$,

$$\widehat{\mathbf{H}}_{\mathbf{y}} = \begin{bmatrix} \left[\begin{array}{ccc} \frac{2c_1^I}{V_1^2} + 4\beta_1 \left((A_1^a)_{(1,1)}^2 + (A_1^a)_{(2,1)}^2 \right) \frac{-2c_1^I(\bar{\theta}_n - \theta_1)}{V_1^2} & -4\beta_1 A_1^a(1,1) & -4\beta_1 A_1^a(2,1) \\ -\frac{c_1^I(\bar{\theta}_n - \theta_1)}{V_1^2} & \frac{2c_1^I(\bar{\theta}_n - \theta_1)^2}{V_1^3} & 0 \\ -4\beta_1 A_1^a(1,1) & 0 & 4\bar{a}_1 + 4\beta_1 + 4\xi \\ -4\beta_1 A_1^a(2,1) & 0 & -4\xi \end{array} \right] \left[\begin{array}{ccc} 2 \sum_{n=1,2} \sum_{m,m \neq n} \beta_n \sum_{l=1,2} A_n^a(l,n) A_n^a(l,m) & 0 & -2\beta_2 A_2^a(1,1) & -2\beta_2 A_2^a(2,1) \\ 0 & 0 & 0 & 0 \\ -2\beta_2 A_2^a(1,1) & 0 & 4\xi & -4\xi \\ -2\beta_2 A_2^a(2,1) & 0 & -4\xi & 4\xi \end{array} \right] \\ \left[\begin{array}{ccc} 2 \sum_{n=1,2} \sum_{m,m \neq n} \beta_n \sum_{l=1,2} A_n^a(l,n) A_n^a(l,m) & 0 & -2\beta_1 A_1^a(1,2) & -2\beta_1 A_1^a(2,2) \\ 0 & 0 & 0 & 0 \\ -2\beta_2 A_2^a(1,1) & 0 & 4\xi & -4\xi \\ -2\beta_2 A_2^a(2,1) & 0 & -4\xi & 4\xi \end{array} \right] \left[\begin{array}{ccc} \frac{2c_2^I}{V_2^2} + 4\beta_2 \left((A_2^a)_{(1,2)}^2 + (A_2^a)_{(2,2)}^2 \right) \frac{-2c_2^I(\bar{\theta}_n - \theta_2)}{V_2^2} & -4\beta_2 A_2^a(1,2) & -4\beta_2 A_2^a(2,2) \\ -\frac{c_2^I(\bar{\theta}_n - \theta_2)}{V_2^2} & \frac{2c_2^I(\bar{\theta}_n - \theta_2)^2}{V_2^3} & 0 \\ -4\beta_2 A_2^a(1,2) & 0 & 4\bar{a}_2 + 4\beta_2 + 4\xi \\ -4\beta_2 A_2^a(2,2) & 0 & -4\xi \end{array} \right] \end{bmatrix}$$

By definition, $\widehat{\mathbf{H}}_{\mathbf{y}}$ is positive definite if and only if $\mathbf{y}^\top \widehat{\mathbf{H}}_{\mathbf{y}} \mathbf{y} > 0, \mathbf{y} \neq 0$. Dealing with ξ terms from diagonal matrices and off-diagonal matrices, we derive the compact expression $4\xi ((y_3 - y_4) + (y_7 - y_8))^2 \geq 0$. Now, dealing with the β terms, we obtain: $\beta_1 [2(A_1^a(1,1)y_1 - y_3) + A_1^a(1,2)y_5]^2 - \beta_1 (A_1^a(1,2)y_5)^2 + \beta_1 [2(A_1^a(2,1)y_1 - y_4) + A_1^a(2,2)y_5]^2 - \beta_1 (A_1^a(2,2)y_5)^2 + \beta_2 [2(A_2^a(1,2)y_5 - y_7) + A_2^a(1,1)y_1]^2 - \beta_2 (A_2^a(1,1)y_1)^2 + \beta_2 [2(A_2^a(2,2)y_5 - y_8) + A_2^a(2,1)y_1]^2 - \beta_2 (A_2^a(2,1)y_1)^2$. Adding to the previous expression the information release cost at the top corner left of the partial Hessians, i.e., $\frac{2c_1^I}{V_1} y_1^2$ and $\frac{2c_2^I}{V_2} y_5^2$, Asm. 5 guarantees that the sum of the above expression

$(G_n^c)^{\text{VE}} = \frac{1}{a_n} (1 - b_n - (\bar{\nu}_n - \underline{\nu}_n) + \phi)$ where $\underline{\mu}_n, \bar{\mu}_n, \underline{\nu}_n, \bar{\nu}_n$ are Lagrangian multipliers related to lower and upper bounds on the demand and controllable generation resp., for EU n . From the primal condition $\sum_m Q_m^{\text{VE}} = 0$, thus we infer that: $\phi = \frac{\sum_m (\theta_m - G_m^{nc}) - \sum_m (\frac{1}{2\bar{a}_m} + \frac{1}{a_m} (1 - b_m))}{\sum_m (\frac{1}{2\bar{a}_m} + \frac{1}{a_m})}$.

Proof of Proposition 3: We explicit the expectation of EU n 's extended utility function $\mathbb{E}_{\varepsilon \sim \mathcal{N}(0, \nu)} [J_n(\mathbf{x}_n, \hat{\boldsymbol{\theta}}, \mathbf{V}, \varepsilon)] = j_n(\mathbf{x}_n, \theta_n) + c_n^I \frac{(\hat{\theta}_n - \theta_n)^2}{2V_n} - \gamma_n + \beta_n \left((D_n - D_n^a(\hat{\boldsymbol{\theta}}))^2 + (G_n^c - G_n^{ca}(\hat{\boldsymbol{\theta}}))^2 + \sum_{m=1}^N (A_n^a)_{(1,m)}^2 V_m + \sum_{m=1}^N (A_n^a)_{(2,m)}^2 V_m \right) + \xi \left(\sum_{m=1}^N (D_m - G_m^c - G_m^{nc}) \right)^2$. Using the former expression, we first compute the partial derivative $\mathbf{F}_{\mathbf{y}_n}$ and the partial Hessian $\mathbf{H}_{\mathbf{y}_n}$. We get

$$\mathbf{H}_{\mathbf{y}_n} = \begin{bmatrix} \frac{c_n^I}{V_n} + 2\beta_n \left((A_n^a)_{(1,n)}^2 + (A_n^a)_{(2,n)}^2 \right) \frac{-c_n^I(\bar{\theta}_n - \theta_n)}{V_n^2} & -2\beta_n A_n^a(1,n) & -2\beta_n A_n^a(2,n) \\ -\frac{c_n^I(\bar{\theta}_n - \theta_n)}{V_n^2} & 0 & 0 \\ -2\beta_n A_n^a(1,n) & 0 & 2\bar{a}_n + 2\beta_n + 2\xi \\ -2\beta_n A_n^a(2,n) & 0 & -2\xi \end{bmatrix} \begin{matrix} a_n + 2\beta_n + 2\xi \\ a_n + 2\beta_n + 2\xi \\ a_n + 2\beta_n + 2\xi \end{matrix}$$

Consider the matrix $\mathbf{H}_{\mathbf{y}}$ of utilities of the EUs where each n^{th} row and m^{th} column block matrix is given as

$$\mathbf{H}_{\mathbf{y}(n,m)} = \begin{bmatrix} \frac{\partial \mathbf{F}_{\mathbf{y}_n}}{\partial \mathbf{y}_m} \\ 2\beta_n \sum_{l=1,2} A_n^a(l,n) A_n^a(l,m) & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ -2\beta_n A_n^a(1,m) & 0 & 2\xi & -2\xi \\ -2\beta_n A_n^a(2,m) & 0 & 2\xi & -2\xi \end{bmatrix}$$

We explicit the expression of the pseudo Hessian for two EUs (1, 2). The results can be extended easily to N EUs:

$$\mathbf{H}_{\mathbf{y}} = \begin{bmatrix} [\mathbf{H}_{\mathbf{y}_1}] & [\mathbf{H}_{\mathbf{y}(1,2)}] \\ [\mathbf{H}_{\mathbf{y}(2,1)}] & [\mathbf{H}_{\mathbf{y}_2}] \end{bmatrix}. \quad (33)$$

Taking transpose and adding, we have (see supplementary material for the detailed calculation):

remains positive. For the remaining terms we check that $(4\bar{a}_l + 4\beta_l)y_{3l+(l-1)}^2 > 0$ and $(2a_l + 4\beta_l)y_{4l}^2 > 0, \forall l = 1, 2$. We conclude that $\widehat{\mathbf{H}}_{\mathbf{y}}$ is positive definite.

Second, (we prove the strong monotonicity of $\mathbf{F}_{\mathbf{y}}$, i.e., that $\langle \mathbf{F}_{\mathbf{y}}(\mathbf{y}) - \mathbf{F}_{\mathbf{y}}(\mathbf{y}'), \mathbf{y} - \mathbf{y}' \rangle \geq \omega \|\mathbf{y} - \mathbf{y}'\|^2, \forall \mathbf{y}, \mathbf{y}' \in \mathcal{Y}$. Using the multi-dimensional Mean Value Theorem on the convex set \mathcal{Y}_n , we obtain the existence of a $\mathbf{z} \in [\mathbf{y}, \mathbf{y}']$, i.e., component-wise $\mathbf{z}_i \in [(y)_i, (y')_i] \forall i$, such that $\langle \mathbf{F}_{\mathbf{y}}(\mathbf{y}) - \mathbf{F}_{\mathbf{y}}(\mathbf{y}'), \mathbf{y} - \mathbf{y}' \rangle = (\mathbf{y} - \mathbf{y}')^\top \nabla_{\mathbf{y}} \mathbf{F}_{\mathbf{y}}(\mathbf{z}) (\mathbf{y} - \mathbf{y}')$. We then use the Spectral Theorem on $\nabla_{\mathbf{y}} \mathbf{F}_{\mathbf{y}}(\cdot) = \mathbf{H}_{\mathbf{y}}(\cdot)$ to ob-

tain the existence of D_n , a diagonal matrix such that $\widehat{\mathbf{H}}_{\mathbf{y}_n}(z) = \mathbf{PDP}^{-1}$, and \mathbf{P} the orthogonal transition matrix. By replacing $\widehat{\mathbf{H}}_{\mathbf{y}}$ by its expression, identifying the variable in the new base, and using the property that the product of a vector by an orthogonal matrix preserves the Euclidean norm of the vector, we obtain: $(\mathbf{y}-\mathbf{y}')^\top \mathbf{PDP}^{-1} (\mathbf{y}-\mathbf{y}') \geq \min_i \left\{ \lambda_i \left(\widehat{\mathbf{H}}_{\mathbf{y}} \right) \right\} \|(\mathbf{y}-\mathbf{y}')\|^2$, where the $\left(\lambda_i \left(\widehat{\mathbf{H}}_{\mathbf{y}} \right) \right)_{i \in \llbracket 1, m_n+2 \rrbracket}$ are the eigenvalues of $\widehat{\mathbf{H}}_{\mathbf{y}}$. We obtain the ω -strong monotonicity by taking the minimum over all the feasible set of the eigenvalues of $\widehat{\mathbf{H}}_{\mathbf{y}}$ (which are positive due to the strict convexity), with $\omega = \min_{\mathbf{y} \in \mathcal{Y}} \min_{i \in \llbracket 1, \sum_n m_n+2 \rrbracket} \lambda_i \left(\widehat{\mathbf{H}}_{\mathbf{y}} \right)$. Since the conditions listed in [31] are met, we conclude that $\mathbb{E}_\varepsilon [J_n(\cdot)]$ is strongly convex.

Proof of Proposition 6: The IC definition assumes that the induced equilibrium between the EUs is a NE. It can be extended considering that the EUs reach a PSE. In that case, IC is enforced by the DA designing $p_n^a(\cdot)$ such that:

$$\mathbf{y}_n^{\text{PSE}}(\beta_n) = \underset{\mathbf{y}_n \in \mathcal{Y}_n}{\operatorname{argmin}} \mathbb{E}_{\varepsilon \sim \mathcal{N}(0, \mathbf{V}_n^{\text{PSE}})} \left[J_n \left(p_n^a \left(\gamma_n, \beta_n, \mathbf{y}_n, \mathbf{y}_{-n}^{\text{PSE}} \right), \mathbf{y}_n \right) \right]. \quad (34)$$

Define $\Psi_n : \mathbb{R}_+ \rightarrow \mathcal{Y}_n$ to be the implicitly defined map such that $\Psi_n : \beta_n \mapsto \mathbf{y}_n^{\text{PSE}}$ returns the solution of (34). Furthermore, from Prop. 3, we know that the lower level of the DA's optimization problem admits a unique PSE, which is characterized by the set of conditions:

$$\begin{aligned} \widehat{\theta}_n^{\text{PSE}} &= \theta_n, \\ D_n^{\text{PSE}}(\beta_n) &= \frac{\tilde{a}_n \theta_n + \beta_n (\mathbf{A}_n^a(1, \cdot) \boldsymbol{\theta} - E_n^a)}{\tilde{a}_n + \beta_n} - \frac{1}{2(\tilde{a}_n + \beta_n)}, \\ G_n^{\text{cPSE}}(\beta_n) &= \frac{2\beta_n (\mathbf{A}_n^a(2, \cdot) \boldsymbol{\theta} - F_n^a)}{a_n + 2\beta_n} - \frac{1}{a_n + 2\beta_n}, \\ V_n^{\text{PSE}} &= \bar{V}_n. \end{aligned}$$

We elucidate the transformations that occur in the EUs' extended utility functions giving rise to the following bilevel optimization problem:

$$\begin{aligned} \min_{\gamma, \beta} \sum_n \left(\frac{\gamma_n}{\beta_n} + \left(\pi - \frac{1}{\beta_n} \right) \left[j_n \left(\mathbf{x}_n^{\text{PSE}}(\beta_n), \theta_n \right) - p^r D_n^{\text{PSE}}(\beta_n) \right] \right), \\ \text{s.t. } \widehat{\theta}_n^{\text{PSE}} &= \theta_n, \forall n, \\ D_n^{\text{PSE}}(\beta_n) &= \frac{\tilde{a}_n \theta_n + \beta_n (\mathbf{A}_n^a(1, \cdot) \boldsymbol{\theta} - E_n^a)}{\tilde{a}_n + \beta_n} - \frac{1}{2(\tilde{a}_n + \beta_n)}, \forall n, \\ G_n^{\text{cPSE}}(\beta_n) &= \frac{2\beta_n (\mathbf{A}_n^a(2, \cdot) \boldsymbol{\theta} - F_n^a)}{a_n + 2\beta_n} - \frac{1}{a_n + 2\beta_n}, \forall n, \\ V_n^{\text{PSE}} &= \bar{V}_n, \forall n, \\ j_n \left(\mathbf{x}_n^{\text{PSE}}(\beta_n), \theta_n \right) - p^r D_n^{\text{PSE}}(\beta_n) &\geq 0, \forall n \in \mathcal{N}, \\ \sum_n \left(D_n^{\text{PSE}}(\beta_n) - G_n^{\text{cPSE}}(\beta_n) - G_n^{\text{nc}} \right) &= 0, \\ \Psi_n(\beta_n) &\in \mathcal{Y}_n, \gamma_n \geq 0, \forall n \in \mathcal{N}. \end{aligned}$$

By differentiating $D_n^{\text{PSE}}(\beta_n)$ and $G_n^{\text{cPSE}}(\beta_n)$ w.r.t. β_n , we obtain that $\frac{\partial D_n^{\text{PSE}}(\beta_n)}{\partial \beta_n} > 0$ and $\frac{\partial^2 D_n^{\text{PSE}}(\beta_n)}{\partial \beta_n^2} > 0$ if and only if $\mathbf{A}_n^a(1, \cdot) \boldsymbol{\theta} + \mathbf{B}_n^a(1, \cdot) \mathbf{G}^{\text{cPSE}}(\boldsymbol{\beta}) > C_n^a(1, 1)$ and $\frac{\partial G_n^{\text{cPSE}}}{\partial \beta_n} > 0$ and $\frac{\partial^2 G_n^{\text{cPSE}}}{\partial \beta_n^2} > 0$ if and only if $\mathbf{A}_n^a(2, \cdot) \boldsymbol{\theta} + \mathbf{B}_n^a(2, \cdot) \mathbf{G}^{\text{cPSE}}(\boldsymbol{\beta}) >$

$C_n^a(2, 1)$. Under these assumptions, $D_n^{\text{PSE}}(\cdot)$ and $G_n^{\text{cPSE}}(\cdot)$ are strictly increasing concave in β_n , therefore $\Psi_n(\cdot)$ is bijective over $B_n \stackrel{\text{def}}{=} [\underline{\beta}_n, \bar{\beta}_n]$, where $\Psi_n(\underline{\beta}_n) = \underline{\mathbf{y}}_n$ and $\Psi_n(\bar{\beta}_n) = \bar{\mathbf{y}}_n$. Finally, we check that the bilevel optimization problem's objective function is strongly convex in \mathbf{x} , thus it admits a unique solution.

B. Computation of the TTSM Surrogate $\bar{\nabla}_{\mathbf{x}_n} J_n(\cdot)$

To compute the TTSM surrogate $\bar{\nabla}_{\mathbf{x}_n} J_n(\mathbf{x}_n, \mathbf{v}_n, \varepsilon)$, $\forall n$, we rely on the partial Hessian $\mathbf{H}_{\mathbf{y}_n}$ and Jacobians

$$\begin{aligned} \nabla_{\mathbf{v}_n \mathbf{v}_n}^2 J_n(\mathbf{x}_n, \mathbf{v}_n \varepsilon) \\ = \begin{bmatrix} \frac{c_n^I}{V_n} + 2\beta_n \left((A_n^a)_{(1,n)}^2 + (A_n^a)_{(2,n)}^2 \right) & -\frac{c_n^I(\widehat{\theta}_n - \theta_n)}{V_n^2} \\ -\frac{c_n^I(\widehat{\theta}_n - \theta_n)}{V_n^2} & \frac{c_n^I(\widehat{\theta}_n - \theta_n)^2}{V_n^3} \end{bmatrix}, \end{aligned} \quad (36)$$

as well as

$$\begin{aligned} \left[\nabla_{\mathbf{v}_n \mathbf{v}_n}^2 J_n(\mathbf{x}_n, \mathbf{v}_n \varepsilon) \right]^{-1} \\ = \frac{V_n^3}{c_n^I(\widehat{\theta}_n - \theta_n)^2 2\beta_n \left((A_n^a)_{(1,n)}^2 + (A_n^a)_{(2,n)}^2 \right)} \\ \times \begin{bmatrix} \frac{c_n^I(\widehat{\theta}_n - \theta_n)^2}{V_n^3} & \frac{c_n^I(\widehat{\theta}_n - \theta_n)}{V_n^2} \\ \frac{c_n^I(\widehat{\theta}_n - \theta_n)}{V_n^2} & \frac{c_n^I}{V_n} + 2\beta_n \left((A_n^a)_{(1,n)}^2 + (A_n^a)_{(2,n)}^2 \right) \end{bmatrix}, \\ \nabla_{\mathbf{x}_n \mathbf{v}_n}^2 J_n(\mathbf{x}_n, \mathbf{v}_n \varepsilon) = \begin{bmatrix} -2\beta_n (A_n^a)_{1,n} & 0 \\ -2\beta_n (A_n^a)_{2,n} & 0 \end{bmatrix}, \end{aligned} \quad (38)$$

$$\begin{aligned} \nabla_{\mathbf{x}_n \mathbf{v}_n}^2 J_n(\mathbf{x}_n, \mathbf{v}_n \varepsilon) \left[\nabla_{\mathbf{v}_n \mathbf{v}_n}^2 J_n(\mathbf{x}_n, \mathbf{v}_n, \varepsilon) \right]^{-1} \\ = \frac{V_n^3}{c_n^I(\widehat{\theta}_n - \theta_n)^2 2\beta_n \left((A_n^a)_{(1,n)}^2 + (A_n^a)_{(2,n)}^2 \right)} \\ \times \begin{bmatrix} -\frac{2\beta_n (A_n^a)_{1,n} c_n^I(\widehat{\theta}_n - \theta_n)}{V_n^3} & -\frac{2\beta_n (A_n^a)_{1,n} c_n^I(\widehat{\theta}_n - \theta_n)}{V_n^2} \\ -\frac{2\beta_n (A_n^a)_{2,n} c_n^I(\widehat{\theta}_n - \theta_n)}{V_n^3} & -\frac{2\beta_n (A_n^a)_{2,n} c_n^I(\widehat{\theta}_n - \theta_n)}{V_n^2} \end{bmatrix}. \end{aligned} \quad (39)$$

REFERENCES

- [1] R. Taisant, M. Datar, H. Le Cadre, and E. Altman, "Learning market equilibria using performative prediction: Balancing efficiency and privacy," in *European Control Conference*. IEEE, 2023.
- [2] H. Le Cadre, P. Jacquot, C. Wan, and C. Alasseur, "Peer-to-peer electricity market analysis: From variational to generalized nash equilibrium," *European Journal of Operational Research*, vol. 282, no. 2, pp. 753–771, 2020.
- [3] I. Shilov, H. Le Cadre, and A. Basic, "Privacy impact on generalized nash equilibrium in peer-to-peer electricity market," *Operations Research Letters*, vol. 49, no. 5, pp. 759–766, 2021.
- [4] C. Campaigne, M. Balandat, and L. Ratliff, "Welfare effects of dynamic electricity pricing," *Working Paper*, 2016.
- [5] G. Belgioioso, W. Ananduta, S. Grammatico, and C. Ocampo-Martinez, "Operational-safe peer-to-peer energy trading in distribution grids: A game-theoretic market-clearing mechanism," *IEEE Transactions on Smart Grid*, vol. 13, no. 4, pp. 2897–2907, 2022.
- [6] —, "Energy management and peer-to-peer trading in future smart grids: A distributed game-theoretic approach," in *European Control Conference*. IEEE, 2020, pp. 1324–1329.
- [7] J. C. Perdomo, T. Zrnica, C. Mendler-Dünner, and M. Hardt, "Performative prediction," in *Machine Learning Research*. PMLR, 2020, pp. 7599–7609.
- [8] A. Narang, E. Faulkner, D. Drusvyatskiy, M. Fazel, and L. J. Ratliff, "Multiplayer performative prediction: Learning in decision-dependent games," *Journal of Machine Learning Research*, vol. 24, pp. 1–56, 2023.
- [9] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Theory of cryptography conference*. Springer, 2006, pp. 265–284.

[10] A. Gosh and K. Ligett, "Privacy and coordination: Computing on databases with endogenous participation," in *ACM Conference on Electronic Commerce*. ACM, 2013.

[11] A. Ghosh and A. Roth, "Selling privacy at auction," *Games and Economic Behavior*, vol. 91, pp. 334–346, 2015.

[12] R. Dong, L. J. Ratliff, A. A. Cárdenas, H. Ohlsson, and S. S. Sastry, "Quantifying the utility–privacy tradeoff in the internet of things," *ACM Trans. Cyber-Phys. Syst.*, vol. 2, no. 2, 2018.

[13] G. Liao, X. Chen, and J. Huang, "Prospect theoretic analysis of privacy-preserving mechanism," *IEEE/ACM Transactions on Networking*, vol. 28, no. 1, pp. 71–83, 2020.

[14] C. Dwork, A. Roth, *et al.*, "The algorithmic foundations of differential privacy," *Foundations and Trends® in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.

[15] P. Kairouz, S. Oh, and P. Viswanath, "Extremal mechanisms for local differential privacy," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 492–542, 2016.

[16] A. Agarwal, M. Dahleh, and T. Sarkar, "A marketplace for data: An algorithmic solution," in *Proceedings of the 2019 ACM Conference on Economics and Computation*, 2019, pp. 701–726.

[17] L. Han, P. Pinson, and J. Kazempour, "Trading data for wind power forecasting: A regression market with lasso regularization," *Electric Power Systems research*, vol. 212, no. 108442, 2022.

[18] Y. Cai, C. Daskalakis, and C. Papadimitriou, "Optimum statistical estimation with strategic data sources," in *Conference on Learning Theory*. PMLR, 2015, pp. 280–296.

[19] T. Westenbroek, R. Dong, L. J. Ratliff, and S. S. Sastry, "Competitive statistical estimation with strategic data sources," *IEEE Transactions on Automatic Control*, vol. 65, no. 4, pp. 1537–1551, 2019.

[20] N. Gast, S. Ioannidis, P. Loiseau, and B. Roussillon, "Linear regression from strategic data sources," *ACM Transactions on Economics and Computation (TEAC)*, vol. 8, no. 2, pp. 1–24, 2020.

[21] Y. Chen, K. Ding, and G. Shi, "Network learning from best-response dynamics in lq games," in *American Control Conference*. IEEE, 2023, pp. 1680–1685.

[22] Y. Chen and G. Shi, "Differentially private games via payoff perturbations," in *American Control Conference*. IEEE, 2023, pp. 429–434.

[23] Y. Wang, J.-F. Zhang, and A. Nedich, "Differentially private distributed algorithms for stochastic aggregative games," *Automatica*, vol. 142, p. 110440, 2022.

[24] F. Fioretto, C. Tran, P. Van Hentenryck, and K. Zhu, "Differential privacy and fairness in decisions and learning tasks: A survey," in *In Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2022, pp. 5470–5477.

[25] Y. Lin, K. Liu, D. Han, and Y. Xia, "Statistical privacy-preserving online distributed nash equilibrium tracking in aggregative games," *IEEE Transactions on Automatic Control*, 2023, DOI 10.1109/TAC.2023.3264164.

[26] L. J. Ratliff, R. Dong, S. Sekar, and T. Fiez, "A perspective on incentive design: Challenges and opportunities," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 2, pp. 305–338, 2019.

[27] D. Acemoglu, A. Makhdomi, A. Malekian, and A. Ozdaglar, "Informational braess' paradox: The effect of information on traffic congestion," *Operations Research*, vol. 66, no. 4, pp. 893–917, 2018.

[28] T. Baroche, P. Pinson, R. L. G. Latimier, and H. B. Ahmed, "Exogenous cost allocation in peer-to-peer electricity markets," *IEEE Transactions on Power Systems*, vol. 34, no. 4, pp. 2553–2564, 2019.

[29] J. Yao, I. Adler, and S. S. Oren, "Modeling and computing two-settlement oligopolistic equilibrium in a congested electricity network," *Operations Research*, vol. 56, no. 1, pp. 34–47, 2008.

[30] M. Rasouli and M. I. Jordan, "Data sharing markets," 2021, <https://arxiv.org/abs/2107.08630>.

[31] G. Scutari, D. P. Palomar, F. Facchinei, and J.-S. Pang, "Convex optimization, game theory, and variational inequality theory," *IEEE Signal Processing Magazine*, vol. 27, no. 3, pp. 35–49, 2010.

[32] P. Mertikopoulos, E. V. Belmega, R. Negrel, and L. Sanguinetti, "Distributed stochastic optimization via matrix exponential learning," *IEEE Transactions on Signal Processing*, vol. 65, no. 9, pp. 2277–2290, 2017.

[33] F. Facchinei, V. Piccialli, and M. Sciandrone, "Decomposition algorithms for generalized potential games," *Computational Optimization and Applications*, vol. 50, no. 2, pp. 237–262, 2011.

[34] K. Chatzikokolakis, M. E. Andrés, N. E. Bordenabe, and C. Palamidessi, "Broadening the scope of differential privacy using

metrics," in *International Symposium on Privacy Enhancing Technologies Symposium*. Springer, 2013, pp. 82–102.

[35] B. Balle and Y.-X. Wang, "Improving the gaussian mechanism for differential privacy: Analytical calibration and optimal denoising," in *International Conference on Machine Learning*. PMLR, 2018, pp. 394–403.

[36] T. Basar and G. J. Olsder, *Dynamic Noncooperative Game Theory, 2nd Edition*. SIAM, 1998.

[37] H. Von Stackelberg, *Market Structure and Equilibrium*. Springer Science & Business Media, 2010.

[38] M. Hong, H.-T. Wai, Z. Wang, and Z. Yang, "A two-timescale stochastic algorithm framework for bilevel optimization: Complexity analysis and application to actor-critic," *SIAM Journal on Optimization*, vol. 33, no. 1, 2023.

[39] D. Chafai and F. Malrieu, "On fine properties of mixtures with respect to concentration of measure and sobolev type inequalities," in *Annales de l'IHP Probabilités et statistiques*, vol. 46, no. 1, 2010, pp. 72–96.

[40] R. M. Freund, "Penalty and barrier methods for constrained optimization," *Lecture Notes, Massachusetts Institute of Technology*, 2004.



Hélène Le Cadre received the M.Sc.A. degree in Applied Mathematics from the ENS Rennes and University of Rennes, Dipl. Ing. with specialization in Information Science and Control from Institut Mines-Télécom Atlantique, and Ph.D. degree in Applied Mathematics from the University of Rennes. She is currently a tenured Inria Researcher, France. Her current research interests revolve around game theory, optimization, learning in games, with applications in markets, security, and societal-scale system problems.



Mandar Datar received the M.Tech. degree in Electrical Engineering with a specialization in Control Systems from the University of Mumbai, India, and a PhD in Computer Science from the University of Avignon, France. Previously, he worked as an assistant professor at VJTI, as a Research Assistant at the IIT Bombay, and as a postdoc at Inria. He is currently working as a researcher at CEA-Leti. His interest includes game theory, optimization, and learning in games, MDP with applications in the wireless communications.



Mathis Guckert is a final-year student at the École Centrale, in Lille, France, and is pursuing a M.Sc. degree in Applied Mathematics with specialization in Data Science. During his studies, he has developed a keen interest in algorithmic game theory and its applications to market-related problems.



Eitan Altman received a B.Sc. degree in Electrical Engineering, B.A. in Physics, B. Mus. in Music Composition, and PhD in Electrical Engineering from Technion Institute of Technology, Haifa, Israel. He is currently a Senior Researcher with Inria, France. His areas of interests include network games, social networks, and game theoretical models of network neutrality issues. He has received the Grand Prix France Télécom, the ISAACS Award, and the Distinguished Technical Achievement Recognition Award.