

A stochastic generalized Nash equilibrium model for platforms competition in the ride-hail market

Original

A stochastic generalized Nash equilibrium model for platforms competition in the ride-hail market / Fabiani, F., Franci, B.. - (2022), pp. 4455-4460. (61st IEEE Conference on Decision and Control, CDC 2022 Cancun (Messico) 06-09 December 2022) [10.1109/cdc51059.2022.9993244].

Availability:

This version is available at: 11583/3003649 since: 2025-10-10T08:35:59Z

Publisher:

Institute of Electrical and Electronics Engineers - IEEE

Published

DOI:10.1109/cdc51059.2022.9993244

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2022 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

A stochastic generalized Nash equilibrium model for platforms competition in the ride-hail market

Filippo Fabiani and Barbara Franci

Abstract—The presence of uncertainties in the ride-hailing market complicates the pricing strategies of on-demand platforms that compete each other to offer a mobility service while striving to maximize their profit. Looking at this problem as a stochastic generalized Nash equilibrium problem (SGNEP), we design a distributed, stochastic equilibrium seeking algorithm with Tikhonov regularization to find an optimal pricing strategy. Remarkably, the proposed iterative scheme does not require an increasing (possibly infinite) number of samples of the random variable to perform the stochastic approximation, thus making it appealing from a practical perspective. Moreover, we show that the algorithm returns a Nash equilibrium under mere monotonicity assumption and a careful choice of the step size sequence, obtained by exploiting the specific structure of the SGNEP at hand. We finally corroborate our results on a numerical instance of the on-demand ride-hailing market.

I. INTRODUCTION

In the last few years, we have been experiencing a dizzying growth of the ride-hailing market [1], where on-demand ride-hailing platforms, such as Uber, Lyft and Didi Chuxing, have to settle up suitable pricing strategies to be attractive on two nearly complementary fronts: costumers and drivers. Each ride-hailing firm, indeed, not only competes for costumers with the other firms and with traditional transportation systems, but also strives to secure an as wide as possible fleet of “loyal” drivers so that it can meet possibly growing costumers’ demand, which often may not be predicted accurately. In this framework, competition among the platforms can be naturally described through a stochastic generalized Nash equilibrium problem (SGNEP): the firms aim at maximizing their expected valued profit function, trying to satisfy the demand for rides, which is typically uncertain, while sharing the market with the other platforms.

Specifically, SGNEPs amount to a collection of mutually coupled stochastic optimization problems, which are challenging to address especially if one aims at finding a solution, i.e., a stochastic generalized Nash equilibrium (SGNE), in a distributed fashion. The difficulties are mainly due to the constraints coupling the agents’ strategies, and the presence of uncertainty. The first issue is typically accommodate by reformulating the problems as a monotone inclusion [2], [3] obtained by exploiting the Karush-Khun-Tucker conditions

F. Fabiani is with the Department of Engineering Science, University of Oxford, OX1 3PJ, United Kingdom (filippo.fabiani@eng.ox.ac.uk). B. Franci is with the Department of Data Science and Knowledge Engineering, Maastricht University, P.O. Box 616, NL-6200 MD Maastricht, The Netherlands (b.franci@maastrichtuniversity.nl). This work was partially supported through the Government’s modern industrial strategy by Innovate UK, part of UK Research and Innovation, under Project LEO (Ref. 104781).

of the coupled optimization problems. Concerning the uncertainty instead, the usual approach approximates the expected-valued pseudogradient mapping of the game by leveraging available realizations of the uncertainty [4], [5].

As a next step, one should design a sequence of instructions alternating distributed computation and communication steps, i.e., an algorithm, with provable convergence guarantees to an equilibrium solution of the SGNEP at hand. Among the numerous algorithms for classic stochastic optimization [6], [7], only few of them are amenable to solve SGNEPs. For this class of problems, indeed, we include the forward-backward (SFB) algorithm [3] and related variations, such as the relaxed forward-backward (SRFB) [8] or the projected-reflected-gradient (SPRG) algorithms [9]. However, these procedures are affected by common drawbacks: the monotonicity assumption of the operators involved and the number of samples necessary for the approximation. In fact, both SFB and SPRG algorithms converge in case the pseudogradient mapping is strongly monotone or cocoercive, which correspond to rather strong assumptions. Moreover, FB-based methods take as an approximation the average over an increasing, possibly infinite, number of samples of the uncertainty, which is impractical or even unrealistic. These reasons motive us to modify the traditional FB algorithm with a Tikhonov (Tik) regularization method [10, Ch. 12], tailored for SGNEPs. Introducing a regularization sequence is indeed a well-known technique to weaken strong assumptions [5], [11], since it allows one to obtain, e.g., a strongly monotone operator starting from a merely monotone one [10]. We summarize all these considerations in Table I.

In [5], a distributed version of this Tikhonov regularization-based algorithm was introduced for SNEPs, i.e., without coupling constraints. We show here that the generalization to SGNEPs is possible, albeit non-trivial, as the time-varying nature of both the step sizes and regularization step poses technical challenges to be treated

	Tik	SpFB [3]	SPRG [9]	SRFB [8]
MONOTONICITY	✓	✗	✗	✓
# SAMPLE(S)	1	N_k	N_k	N_k
STEP SIZE	α_k	α	α	α

TABLE I: Existing FB-based algorithms for SGNEPs which converge with monotonicity (✓) or stronger assumptions (✗). The quantity N_k indicates an increasing (possibly infinite) number of samples, while α_k indicates a the time-varying step size sequence, as opposed to α .

carefully when using operator splitting techniques (§IV-B). We can hence summarize our contributions as follows:

- Inspired by the cognate literature, we propose a noncooperative model for on-demand competing ride-hailing platforms under a regulated pricing scenario and uncertainties, and we recast it as a SGNP (§II, III);
- We propose a distributed Tikhonov regularization-based algorithm that leverages a finite number of samples of the uncertainty to perform the stochastic approximation, thus circumventing a crucial issue in equilibrium seeking algorithm design for SGNPs (§IV);
- We show that the algorithm converges to a SGNP i) under mere monotonicity of the pseudogradient mapping, one of the weakest assumptions to establish convergence [12], and ii) with a careful choice of the step size sequence that exploits the structure of the problem.

In conclusion, the performance of the designed Tikhonov-like algorithm is tested on a numerical instance of the proposed ride-hailing competition market model (§V).

A. Notation and Preliminaries

a) *Notation*: \mathbb{N} indicates the set of natural numbers and \mathbb{R} ($\bar{\mathbb{R}} = \mathbb{R} \cup \{\infty\}$) is the set of (extended) real numbers. $\langle \cdot, \cdot \rangle : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ denotes the standard inner product and $\|\cdot\|$ is the associated Euclidean norm. Given a vector $x \in \mathbb{R}^n$, $x_{\min} = \min_{i=1, \dots, n} x_i$. We indicate that a matrix A is positive definite, i.e., $x^\top A x > 0$, with $A \succ 0$. Given a symmetric $W \succ 0$, the W -induced inner product is $\langle x, y \rangle_W = \langle Wx, y \rangle$ and the associated norm is defined as $\|x\|_W = \sqrt{\langle Wx, x \rangle}$. Id is the identity operator. $\iota_{\mathcal{X}}$ is the indicator function of the set \mathcal{X} , that is, $\iota_{\mathcal{X}}(x) = 0$ if $x \in \mathcal{X}$ and $\iota_{\mathcal{X}}(x) = \infty$ otherwise. The set-valued mapping $N_{\mathcal{X}} : \mathbb{R}^n \rightarrow \bar{\mathbb{R}}^n$ denotes the normal cone operator of the set \mathcal{X} , i.e., $N_{\mathcal{X}}(x) = \emptyset$ if $x \notin \mathcal{X}$, $N_{\mathcal{X}}(x) = \{v \in \mathbb{R}^n \mid \sup_{z \in \mathcal{X}} \langle v, z - x \rangle \leq 0\}$ otherwise.

b) *Operator theory*: Let $\text{gra}(F) = \{(x, u) : u \in F(x)\}$ be the graph of $F : \mathcal{X} \subseteq \mathbb{R}^n \rightarrow \mathbb{R}^n$. Then, F is said to be: monotone on \mathcal{X} if $\langle F(x) - F(y), x - y \rangle \geq 0$, for all $x, y \in \mathcal{X}$; μ -strongly monotone on \mathcal{X} if there exists a constant $\mu > 0$ such that $\langle F(x) - F(y), x - y \rangle \geq \mu \|x - y\|^2$, for all $x, y \in \mathcal{X}$; maximally monotone if there exists no monotone operator $G : \mathcal{X} \rightarrow \mathbb{R}^n$ such that $\text{gra } G$ properly contains $\text{gra } F$: ℓ -Lipschitz continuous with constant $\ell > 0$ if for all $x, y \in \mathcal{X}$ $\|F(x) - F(y)\| \leq \ell \|x - y\|$. The projection operator onto \mathcal{C} is the operator defined as $\text{proj}_{\mathcal{C}}(x) = \text{argmin}_{z \in \mathcal{C}} \|z - x\|^2$.

c) *Graph theory*: The weighted adjacency matrix associated to \mathcal{G}^λ is denoted with $W = [w_{ij}]_{i, j \in \mathcal{I}} \in \mathbb{R}^{N \times N}$, where $w_{ij} > 0$ if agents i and j can communicate with each other and $w_{ij} = 0$ otherwise. Then, letting $D = \text{diag}\{d_1, \dots, d_N\}$ where $d_{i,h} = \sum_{j=1}^N w_{ij}$ is the degree of agent i , the associated Laplacian is given by $L = D - W \in \mathbb{R}^{N \times N}$. It follows from Assumption 5 that $L = L^\top$.

d) *Stochastic setting*: \mathbb{E}_ξ represent the mathematical expectation with respect to the distribution of the random variable $\xi(\omega)$, in the probability space $(\Xi, \mathcal{F}, \mathbb{P})$. We avoid expressing the dependency on ω when clear from the context and often simply write \mathbb{E} . We use i.i.d. to indicate independent identically distributed random variables.

II. ON-DEMAND COMPETING RIDE-HAILING FIRMS

We take inspiration from [13]–[15] to examine how $N \in \mathbb{N}$ on-demand competing ride-hailing platforms (e.g., Uber, Didi Chuxing, Lyft, Juno and Via) design their pricing strategies under a regulated pricing scenario and the presence of uncertainties. Specifically, given a continuum of potential riders of mass $C_h > 0$ for each area of interest $h \in \mathcal{H} \subset \mathbb{N}$ (e.g., suburbs, city centres, airports), each firm $i \in \mathcal{I} := \{1, \dots, N\}$ aims at maximizing its profit by setting i) a price $p_{i,h} \geq 0$ for the on-demand ride-hailing service to attract as many costumers as possible in the h -th area, with

$$\frac{1}{N} \sum_{i \in \mathcal{I}} p_{i,h} \leq \bar{p}_h, \quad \forall h \in \mathcal{H}, \quad (1)$$

capping the averaged maximum price allowed $\bar{p}_h > 0$, usually imposed by consumers' associations; ii) a wage $w_{i,h} \geq \underline{w} > 0$ for the registered drivers on the i -th ride-hailing platform, which is also typically regulated by institutions [13], [16], to meet the resulting costumers' demand. As one may expect, since profit maximization is a consideration, $p_{i,h}$ and $w_{i,h}$ shall be necessarily interdependent.

Similar to [14], [17], we assume that the fraction of costumers who choose the i -th platform's service in the h -th area, i.e., the demand for the i -th firm, is characterized as:

$$d_{i,h} = \frac{C_h K_{i,h}}{\bar{p} \sum_{j \in \mathcal{I}} K_{j,h}} \left(\bar{p} - p_{i,h} + \frac{\theta_i}{N-1} \sum_{j \in \mathcal{I} \setminus \{i\}} p_{j,h} \right), \quad (2)$$

where $K_{i,h} > 0$ denotes the number of registered drivers on the i -th ride-hailing platform who prefer to work in the h -th area, $\bar{p} > \max_{h \in \mathcal{H}} \bar{p}_h$ is a maximum service price, and $\theta_i \in [0, 1]$ models the substitutability of the service provided by each firm. When θ_i is close to 0, the service of the i -th platform is almost independent from the others, while θ_i close to 1 means that it is fully substitutable, thus obtaining a perfect competition market. We discuss in details the role played by this parameter in §V with a numerical example.

Note that, however, the demand request in (2) does not account for the willingness of the drivers to actually provide a service, which is key to meet the costumers' demand and hence maximize the profit. In fact, any driver provides service in a prescribed area $h \in \mathcal{H}$ only if its earn is greater than their *opportunity cost*, here denoted by $\delta_{i,h}$, which we let coincide with a random variable. For each firm and area, we assume the wage be given by $w_{i,h} = \beta p_{i,h}$, where the parameter $\beta > 0$ denotes the commission ratio that the platform should pay to its driver, typically regulated by a third party (e.g., governments [13]). Thus, for all $i \in \mathcal{I}$ and $h \in \mathcal{H}$, we introduce the *effective demand* as

$$d_{i,h}^e = d_{i,h} \mathbb{P}[w_{i,h} \geq \delta_{i,h}], \quad (3)$$

which coincides with the portion of costumers' demand for which a ride-hailing company can actually claim a payment. Roughly speaking, a driver is willing to provide a service only if the wage she gets matches (at least) her expectations.

This directly leads us to define the fraction of drivers that give services on the i -th platform as:

$$k_{i,h} = K_{i,h} \mathbb{P}[w_{i,h} \geq \delta_{i,h}]. \quad (4)$$

As a consequence, the effective demand in (3) can be equivalently obtained from (2) by replacing $K_{i,h}$ with $k_{i,h}$. However, the costumers' service request in (2) may be affected by an additional source of uncertainty. In fact, it is unlikely that the i -th firm is aware of the total number of potential drivers, $\sum_{j \in \mathcal{I}} K_{j,h}$, both for privacy reasons and possible multiple registrations. Thus, we define $C_h(\xi) = \frac{C_h(\xi)}{\sum_{j \in \mathcal{I}} K_{j,h}}$ as the unknown fraction of passengers in area $h \in \mathcal{H}$, which allows us to explicitly account for the uncertain parameter ξ . In accordance, the demands in (2) and (3) turn into random variables $d_{i,h}(\xi)$ and $d_{i,h}^c(\xi)$.

As commonly adopted in the literature [13], [15], we restrict attention to the case where the drivers' willingness is uniformly distributed in $[\underline{w}, \bar{w}_{i,h}]$, so that $\mathbb{P}[w_{i,h} \geq \delta_{i,h}] = (w_{i,h} - \underline{w}) / (\bar{w}_{i,h} - \underline{w})$. This not only allows us to consider one source of uncertainty, but also to make the constraint in (4) convex. Note that a similar argument can be adopted for any distribution concave in the induced demand $d_{i,h}$ (e.g., exponential, Pareto). Then, the stochastic optimization problem associated to each platform amounts to:

$$\forall i \in \mathcal{I} : \begin{cases} \max_{(p_{i,h}, w_{i,h})_{h \in \mathcal{H}}} & \mathbb{E}_\xi \left[\sum_{h \in \mathcal{H}} (p_{i,h} d_{i,h}^c(\xi) - w_{i,h} k_{i,h}) \right] \\ \text{s.t.} & (1), (3), (4), p_{i,h} \geq 0, \forall h \in \mathcal{H}, \\ & w_{i,h} \in [\underline{w}, \bar{w}_{i,h}], \forall h \in \mathcal{H}. \end{cases} \quad (5)$$

The first part of the cost function amounts to the profit of the i -th firm to provide a service to the costumers, while the second one considers the costs for providing a service to the drivers. Unlike [13], [14], [17], the cost functions in (5) accounts for the number of actual drivers who decide to provide a service rather than the whole fleet of registered ones, $K_{i,h}$. After replacing the equality constraints in (3) and (4), we obtain a collection of mutually coupled stochastic optimization problems, where the cost functions (cubic in $p_{i,h}$, due to the distribution of $\delta_{i,h}$) are affected by the uncertain fraction of potential riders, i.e., $C_h(\xi)$. The proposed model relies on a common assumption in the literature: those riders who do not get assigned to a driver when they seek service from the ride-hailing platforms, e.g., because of excess demand for rides, use a different mean of transportation. The need for a common platform handling competition in ride-hailing mobility to satisfy the costumers' demand has been indeed recently explored in, e.g., [18], [19].

Throughout the paper we treat this model as a SGNEP, and we propose an algorithm to compute an equilibrium solution, according to the discussion presented in the next section.

III. STOCHASTIC GENERALIZED NASH EQUILIBRIUM PROBLEM

To compact the notation, we rewrite the problem in (5) as

$$\forall i \in \mathcal{I} : \begin{cases} \min_{x_i \in \Omega_i} & \mathbb{J}_i(x_i, \mathbf{x}_{-i}) \\ \text{s.t.} & g(x_i, \mathbf{x}_{-i}) \leq 0. \end{cases} \quad (6)$$

where $x_i = \text{col}((p_{i,h})_{h \in \mathcal{H}}) \in \mathbb{R}^n$, $n := |\mathcal{H}|$, $\mathbf{x} = \text{col}((x_i)_{i \in \mathcal{I}}) \in \mathbb{R}^{nN}$, and $\mathbf{x}_{-i} = \text{col}((x_j)_{j \neq i})$. Moreover, we indicate the set of local constraints of firm i as $\Omega_i \in \mathbb{R}^n$ and let $\Omega := \prod_{i \in \mathcal{I}} \Omega_i$. On the other hand, the set of coupling constraints arising from (1) in a general form read as

$$\mathcal{X} := \Omega \cap \{\mathbf{y} \in \mathbb{R}^{nN} \mid g(\mathbf{y}) \leq \mathbf{0}_m\}, \quad (7)$$

where $g : \mathbb{R}^{nN} \rightarrow \mathbb{R}^m$. We indicate with $\mathcal{X}_i(\mathbf{x}_{-i})$ the piece of coupling constraints corresponding to agent i , which is affected by the decision variables of the other agents \mathbf{x}_{-i} . We stress that the formulation of the SGNEP in (6) is standard [2], [8], [20], and hence the theory we develop applies to all the SGNEPs satisfying the assumptions introduced next.

Assumption 1 (Constraint qualification): For each $i \in \mathcal{I}$, the set Ω_i is nonempty, closed and convex. The set \mathcal{X} satisfies Slater's constraint qualification. \square

Assumption 2 (Separable convex coupling constraints): The mapping g in (7) has a separable form, i.e., $g(\mathbf{x}) := \sum_{i=1}^N g_i(x_i)$, for some convex differentiable functions $g_i : \mathbb{R}^n \rightarrow \mathbb{R}^m$, $i \in \mathcal{I}$ and it is ℓ_g -Lipschitz continuous. Its gradient ∇g is bounded, i.e., $\sup_{\mathbf{x} \in \mathcal{X}} \|\nabla g(\mathbf{x})\| \leq B_{\nabla g}$. \square

Given the stochastic nature of our collection of problems, we indicate the cost function of each agent $i \in \mathcal{I}$ as

$$\mathbb{J}_i(x_i, \mathbf{x}_{-i}) := \mathbb{E}[J_i(x_i, \mathbf{x}_{-i}, \xi(\omega))], \quad (8)$$

for some measurable function $J_i : \mathbb{R}^n \times \mathbb{R}^d \rightarrow \mathbb{R}$. We assume that $\mathbb{E}[J_i(\mathbf{x}, \xi)]$ is well defined for all feasible $\mathbf{x} \in \mathcal{X}$ [21].

Assumption 3 (Cost function convexity): For each $i \in \mathcal{I}$ and $\mathbf{x}_{-i} \in \mathcal{X}_{-i}$ the function $\mathbb{J}_i(\cdot, \mathbf{x}_{-i})$ is convex and continuously differentiable. \square

The goal of the firms is hence to solve (6) to find a SGNE, i.e., a strategy profile where no agent can decrease its cost function by unilaterally deviating from its decision. Formally, a SGNE is a collective vector $\mathbf{x}^* \in \mathcal{X}$ such that for all $i \in \mathcal{I}$

$$\mathbb{J}_i(x_i^*, \mathbf{x}_{-i}^*) \leq \inf \{ \mathbb{J}_i(y, \mathbf{x}_{-i}^*) \mid y \in \mathcal{X}_i(\mathbf{x}_{-i}^*) \}.$$

Existence of a SGNE for the game in (6) is guaranteed under suitable assumptions [21, §3.1], though uniqueness does not hold in general [21, §3.2]. Within all possible Nash equilibria, we focus on those coinciding with the solutions of a suitable stochastic variational inequality (SVI) [5]. Thus, we introduce the pseudogradient mapping of the game as

$$\mathbb{F}(\mathbf{x}) := \text{col}((\mathbb{E}[\nabla_{x_i} J_i(x_i, \mathbf{x}_{-i})])_{i \in \mathcal{I}}). \quad (9)$$

Flipping the expected value and the gradient follows from differentiability of $J_i(\cdot, \mathbf{x}_{-i})$ (Assumption 3) [22, Th. 7.44]. Then, the SVI associated to the SGNEP in (6) reads as

$$\langle \mathbb{F}(\mathbf{x}^*), \mathbf{x} - \mathbf{x}^* \rangle \geq 0, \text{ for all } \mathbf{x} \in \mathcal{X}. \quad (10)$$

If Assumptions 1–3 hold, any solution to SVI(\mathcal{X}, \mathbb{F}) in (10) is a SGNE of the game in (6), while the converse does not necessarily hold. A game may have a Nash equilibrium while the associated (S)VI may have no solution [23, Prop. 12.7].

Assumption 4 (Existence of a variational equilibrium): The SVI in (10) has at least one solution. \square

Algorithm 1 Distributed stochastic Tikhonov relaxation

Initialization: $x_i^0 \in \Omega_i, \lambda_i^0 \in \mathbb{R}_{>0}^m$, and $z_i^0 \in \mathbb{R}^m$.

Iteration k : Agent i receives x_j^k for all $j \in \mathcal{N}_i^J$ and z_j^k, λ_j^k for $j \in \mathcal{N}_i^\lambda$, then updates:

$$\begin{aligned} x_i^{k+1} &= \text{proj}_{\Omega_i} \{x_i^k - \alpha_i^k \gamma_i (\hat{F}_i(x_i^k, \mathbf{x}_{-i}^k, \xi_i^k) + \nabla g_i(x_i)^\top \lambda_i^k \\ &\quad + \epsilon_i^k x_i^k)\} \\ z_i^{k+1} &= z_i^k - \alpha_i^k \nu_i (\sum_{j \in \mathcal{N}_i^\lambda} w_{ij} (\lambda_i^k - \lambda_j^k) + \epsilon_i^k z_i^k) \\ \lambda_i^{k+1} &= \text{proj}_{\mathbb{R}_{>0}^m} \{\lambda_i^k + \alpha_i^k \tau_i (g_i(x_i^k) - \epsilon_i^k \lambda_i^k) \\ &\quad + \alpha_i^k \tau_i \sum_{j \in \mathcal{N}_i^\lambda} w_{ij} [(z_i^k - z_j^k) - (\lambda_i^k - \lambda_j^k)]\} \end{aligned}$$

We call variational equilibria (v-SGNE) the SGNE that are also solution to SVI(\mathcal{X}, \mathbb{F}) in (10) with \mathbb{F} in (9) and \mathcal{X} in (7). These equilibria can be characterized in terms of the Karush–Kuhn–Tucker (KKT) conditions of the coupled optimization problems in (6), i.e., a \mathbf{x}^* is a v-SGNE if and only if the following inclusion is satisfied for $\lambda \in \mathbb{R}_{>0}^m$ [12, Th. 4.6]:

$$0 \in \mathcal{T}(\mathbf{x}, \lambda) := \begin{bmatrix} \mathbb{F}(\mathbf{x}) + \mathbf{N}_{\Omega}(\mathbf{x}) + \nabla g(\mathbf{x})^\top \lambda \\ \mathbf{N}_{\mathbb{R}_{>0}^m}(\lambda) - g(\mathbf{x}) \end{bmatrix}, \quad (11)$$

where $\mathcal{T} : \mathcal{X} \times \mathbb{R}_{>0}^m \rightrightarrows \mathbb{R}^{nN} \times \mathbb{R}^m$ is a set-valued mapping. According to [24, Th. 3.1], [25, Th 3.1], the v-SGNE are those equilibria such that the shared constraints have the same dual variable for all the agents, i.e., $\lambda_i = \lambda$ for all $i \in \mathcal{I}$, and solve the SVI(\mathcal{X}, \mathbb{F}) in (10). Then, the v-SGNE of the game in (6) correspond to the zeros of \mathcal{T} , which can be split as the sum of two operators, $\mathcal{T} = \mathcal{A} + \mathcal{B}$, where

$$\begin{aligned} \mathcal{A} : \begin{bmatrix} \mathbf{x} \\ \lambda \end{bmatrix} &\mapsto \begin{bmatrix} \mathbb{F}(\mathbf{x}) \\ 0 \end{bmatrix} + \begin{bmatrix} \nabla g(\mathbf{x})^\top \lambda \\ -g(\mathbf{x}) \end{bmatrix}, \\ \mathcal{B} : \begin{bmatrix} \mathbf{x} \\ \lambda \end{bmatrix} &\mapsto \begin{bmatrix} \mathbf{N}_{\Omega}(\mathbf{x}) \\ \mathbf{N}_{\mathbb{R}_{>0}^m}(\lambda) \end{bmatrix}. \end{aligned} \quad (12)$$

IV. DISTRIBUTED STOCHASTIC TIKHONOV RELAXATION

We now discuss in details the sequence of instructions summarized in Algorithm 1. For the local decision variable x_i the projection onto Ω_i guarantees that the local constraints are always satisfied, while the coupling constraints are enforced asymptotically through the (nonnegative, due to the projection onto $\mathbb{R}_{>0}^m$) dual variable λ_i . The auxiliary variable z_i , instead, forces consensus on the dual variables [2], [24].

We assume that the decision-maker i knows its feasible set Ω_i , and its part of the coupling constraints $\mathcal{X}_i(x_{-i})$. The set of agents j whose decision variables affect the cost function of agent i , are denoted by \mathcal{N}_i^J . Specifically, some $j \in \mathcal{I}$ belongs to \mathcal{N}_i^J if $J_i(x_i, x_{-i})$ explicitly depends on x_j .

Let us then introduce the graph $\mathcal{G}^\lambda = (\mathcal{I}, \mathcal{E}^\lambda)$ through which a local copy of the dual variable is shared, along with of the auxiliary one, $z_i \in \mathbb{R}^m$. The set of edges \mathcal{E}^λ of the multiplier graph \mathcal{G}^λ , is given by: $(i, j) \in \mathcal{E}^\lambda$ if player j share its $\{\lambda_j, z_j\}$ with player i . For all $i \in \mathcal{I}$, the neighboring agents in \mathcal{G}^λ form the set $\mathcal{N}_i^\lambda = \{j \in \mathcal{I} : (i, j) \in \mathcal{E}^\lambda\}$. Under these premises, Algorithm 1 is distributed in the sense that

each agent knows its own problem data and communicates with the other agents through \mathcal{G}^λ . To guarantee that consensus can be reached, we make the following assumption.

Assumption 5 (Graph connectivity): The multiplier graph \mathcal{G}^λ is undirected and connected. \square

By making use of \mathcal{G}^λ , we hence note that consensus on the dual variable can be enforced via equality constraint $\mathbf{L}\lambda = 0$, where $\mathbf{L} = L \otimes \text{Id}_m \in \mathbb{R}^{Nm \times Nm}$, L being the Laplacian of the graph, and $\lambda = \text{col}(\lambda_1, \dots, \lambda_N) \in \mathbb{R}^{Nm}$. Following [2], the operators \mathcal{A} and \mathcal{B} in (12) can thus be extended to

$$\begin{aligned} \bar{\mathcal{A}} : \begin{bmatrix} \mathbf{x} \\ \mathbf{z} \\ \lambda \end{bmatrix} &\mapsto \begin{bmatrix} \mathbb{F}(\mathbf{x}) \\ 0 \\ \mathbf{L}\lambda \end{bmatrix} + \begin{bmatrix} \nabla G(\mathbf{x})^\top \lambda \\ \mathbf{L}\lambda \\ -G(\mathbf{x}) - \mathbf{L}\mathbf{z} \end{bmatrix}, \\ \bar{\mathcal{B}} : \begin{bmatrix} \mathbf{x} \\ \mathbf{z} \\ \lambda \end{bmatrix} &\mapsto \begin{bmatrix} \mathbf{N}_{\Omega}(\mathbf{x}) \\ \mathbf{0} \\ \mathbf{N}_{\mathbb{R}_{>0}^m}(\lambda) \end{bmatrix}, \end{aligned} \quad (13)$$

where $\mathbf{z} = \text{col}(z_1, \dots, z_N) \in \mathbb{R}^{Nm}$, $G(\mathbf{x}) = \text{diag}((g_i(x_i))_{i \in \mathcal{I}})$ and $\nabla G(\mathbf{x}) = \text{diag}((\nabla_{x_i} g_i(x_i))_{i \in \mathcal{I}})$. From now on, we indicate the local decision variable of each agent taking part to the SGNEP as $\mathbf{u} = \text{col}(\mathbf{x}, \mathbf{z}, \lambda)$.

Remark 1: It can be proven that if Assumptions 1–5 are satisfied, then the following hold [2], [3], [8]:

- (i) Given any $\mathbf{u}^* \in \text{zer}(\bar{\mathcal{A}} + \bar{\mathcal{B}})$, \mathbf{x}^* is a v-SGNE of game in (6), i.e., \mathbf{x}^* solves the SVI(\mathcal{X}, \mathbb{F}) in (10), $\lambda^* = \mathbf{1}_N \otimes \lambda^*$, and $(\mathbf{x}^*, \lambda^*)$ satisfy the KKT condition in (11), i.e., $\text{col}(\mathbf{x}^*, \lambda^*) \in \text{zer}(\mathcal{A} + \mathcal{B})$.
- (ii) $\text{zer}(\mathcal{A} + \mathcal{B}) \neq \emptyset$ and $\text{zer}(\bar{\mathcal{A}} + \bar{\mathcal{B}}) \neq \emptyset$. \square

Note that the update of the primal variable x_i in Algorithm 1 makes use of an approximation \hat{F} of the pseudogradient mapping \mathbb{F} , since the distribution of the random variable is unknown and hence the expected value mapping can be hard to compute. Thus, at each iteration k we let

$$\begin{aligned} \hat{F}(\mathbf{x}^k, \xi^k) &= \text{col}((\hat{F}_i(\mathbf{x}^k, \xi^k))_{i \in \mathcal{I}}) \\ &= \text{col}(\nabla_{x_1} J_1(\mathbf{x}^k, \xi_1^k), \dots, \nabla_{x_N} J_N(\mathbf{x}^k, \xi_N^k)), \end{aligned} \quad (14)$$

where $\xi^k = \text{col}(\xi_1^k, \dots, \xi_N^k) \in \mathbb{R}^N$ is a collection of i.i.d. random variables drawn from \mathbb{P} . Essentially, we replace the expected-valued pseudogradient mapping with the realization of *one sample* of the random variable. It follows that

$$\hat{F}(\mathbf{x}^k, \xi^k) = \mathbb{F}(\mathbf{x}^k) - \delta^k(\mathbf{x}^k, \xi^k),$$

where δ^k is typically called stochastic or approximation error. By exploiting the approximation \hat{F} in (14) of the expected value mapping \mathbb{F} , we replace the operator $\bar{\mathcal{A}}$ in (13) with

$$\hat{\mathcal{A}} : \begin{bmatrix} (\mathbf{x}, \xi) \\ \mathbf{z} \\ \lambda \end{bmatrix} \mapsto \begin{bmatrix} \hat{F}(\mathbf{x}, \xi) \\ 0 \\ \mathbf{L}\lambda \end{bmatrix} + \begin{bmatrix} \nabla G(\mathbf{x})^\top \lambda \\ \mathbf{L}\lambda \\ -G(\mathbf{x}) - \mathbf{L}\mathbf{z} \end{bmatrix}. \quad (15)$$

Algorithm 1 can now be rewritten in compact form as

$$\mathbf{u}^{k+1} = (\text{Id} + \Phi_k^{-1} \bar{\mathcal{B}})^{-1} (\mathbf{u}^k - \Phi_k^{-1} (\hat{\mathcal{A}} \mathbf{u}^k + \varepsilon^k \mathbf{u}^k)), \quad (16)$$

where $\Phi_k \succ 0$ contains the inverse of step size sequences

$$\Phi_k = \alpha_k^{-1} \Phi = \alpha_k^{-1} \text{diag}(\gamma^{-1}, \nu^{-1}, \tau^{-1}), \quad (17)$$

with γ^{-1} , ν^{-1} , τ^{-1} being diagonal matrices, and $\varepsilon^k = \text{diag}(\varepsilon_j^k) \in \mathbb{R}^{T \times T}$, $T = nN + 2Nm$, contains the regularization steps. This last term is what typically characterize the Tikhonov regularization scheme [5], [10].

A. Convergence analysis

We now study the convergence properties of Algorithm 1. First, to ensure that $\bar{\mathcal{A}}$ and $\bar{\mathcal{B}}$ have the properties that we use for the analysis, we make the following assumption, which allows us to state the result immediately below.

Assumption 6 (Monotonicity): \mathbb{F} in (9) is monotone and $\ell_{\mathbb{F}}$ -Lipschitz continuous for some $\ell_{\mathbb{F}} > 0$. \square

Lemma 1: [8, Lemma 2 and 4] Let Assumptions 5 and 6 hold true, and let $\Phi \succ 0$. Then, the operators $\bar{\mathcal{A}}$ and $\bar{\mathcal{B}}$ in (13) have the following properties.

- 1) $\bar{\mathcal{A}}$ is monotone and $\ell_{\bar{\mathcal{A}}}$ -Lipschitz continuous.
- 2) $\bar{\mathcal{B}}$ is maximally monotone.
- 3) $\Phi^{-1}\bar{\mathcal{A}}$ is monotone and ℓ_{Φ} -Lipschitz continuous.
- 4) $\Phi^{-1}\bar{\mathcal{B}}$ is maximally monotone. \square

The Lipschitz constants in Lemma 1 depends on those of g_i , $i \in \mathcal{I}$, and \mathbb{F} (Assumptions 1 and 6, respectively). Their specific expressions, however, are not relevant for our analysis, and therefore we point to [8] for additional details.

Remark 2: Introducing the regularization term makes $\bar{\mathcal{A}} + \varepsilon^k$ strongly monotone [10, Th. 12.2.3]. Thus, mere monotonicity of the pseudogradient mapping is enough to show convergence. In addition, note that the regularization term is added to the operator $\bar{\mathcal{A}}$ and not to force strong monotonicity of the SVI in (10). In fact, starting from a strongly monotone mapping, the resulting operator $\bar{\mathcal{A}}$ can be at most cocoercive [2, Lemma 5 and 7], [26, Lemma 2 and 4]. \square

Taking few samples as in (14) is realistic and computationally tractable, at the price of requiring an additional assumption on the step sizes. Specifically, we indicate next how to choose the parameters in Algorithm 1 to ensure that they are vanishing to control the approximation error [5].

Assumption 7: The step size sequence $(\alpha^k)_{k \in \mathbb{N}}$ and the regularizing sequences $(\varepsilon_j^k)_{k \in \mathbb{N}}$, $j = 1, \dots, T$, are such that $\alpha^k = (k + \eta)^{-a}$ and $\varepsilon_j^k = (k + \zeta_j)^{-b}$ for $k \geq 0$, where each η and ζ_j are selected from a uniform distribution on the intervals $[\eta, \bar{\eta}]$ and $[\zeta, \bar{\zeta}]$, respectively, for some $0 < \eta < \bar{\eta}$ and $0 < \zeta < \bar{\zeta}$ and $a, b \in (0, 1)$, $a + b < 1$, and $a > b$. \square

Let us define the filtration $\mathcal{F} = \{\mathcal{F}_k\}_{k \in \mathbb{N}}$, that is, a family of σ -algebras such that $\mathcal{F}_0 = \sigma(X_0)$ and $\mathcal{F}_k = \sigma(X_0, \xi_1, \xi_2, \dots, \xi_k)$ for all $k \geq 1$, such that $\mathcal{F}_k \subseteq \mathcal{F}_{k+1}$ for all $k \in \mathbb{N}$. In words, \mathcal{F}_k contains the information up to iteration k . Since we consider the approximation in (14), we let $\Delta^k = \text{col}(\delta^k, 0, 0)$ be the stochastic error, i.e., $\Delta^k = \bar{\mathcal{A}}(\mathbf{u}^k) - \hat{\mathcal{A}}(\mathbf{u}^k, \boldsymbol{\xi}^k)$. Then, an additional assumption is finally needed to regulate its asymptotic behaviour.

Assumption 8: The step size sequence $(\alpha^k)_{k \in \mathbb{N}}$, regularization sequence $(\varepsilon^k)_{k \in \mathbb{N}}$ and the stochastic error Δ^k satisfy $\lim_{k \rightarrow \infty} (\alpha^k / \varepsilon_{k, \min}) \mathbb{E}[\|\Delta^k\|_{\Phi}^2 \mid \mathcal{F}_k] = 0$ and $\sum_{k=0}^{\infty} \alpha_k^2 \mathbb{E}[\|\Delta^k\|_{\Phi}^2 \mid \mathcal{F}_k] < \infty$ a.s. \square

We are now ready to state our main convergence result.

Theorem 1: Let Assumptions 1 - 8 hold true. Then, the sequence $(\mathbf{x}^k)_{k \in \mathbb{N}}$ generated by Algorithm 1 with \bar{F} as in (14) converges a.s. to a v-SGNE of the game in (6). \square

Proof: Convergence to the primal-dual solution follows similarly to [5] by using the Φ -induced metric. Specifically, with similar steps as [5, Prop. 1 and 2] and by letting \mathbf{y}^k be the sequence generated by the centralized Tikhonov method [5, Lemma 3], we obtain

$$\mathbb{E}[\|\mathbf{u}^{k+1} - \mathbf{y}^k\|_{\Phi}^2 \mid \mathcal{F}_k] \leq (1 - c(\alpha^k, \varepsilon^k)) \|\mathbf{u}^{k+1} - \mathbf{y}^k\|_{\Phi}^2 + d(\alpha^k, \varepsilon^k) + \alpha_k^2 \mathbb{E}[\|\Delta^k\|_{\Phi}^2 \mid \mathcal{F}_k]$$

where $d(\alpha^k, \varepsilon^k)$ and $\alpha_k^2 \mathbb{E}[\|\Delta^k\|_{\Phi}^2 \mid \mathcal{F}_k]$ vanish as $k \rightarrow \infty$ (Assumptions 7 and 8) and $c(\alpha^k, \varepsilon^k) \in [0, 1]$ is such that [27, Lemma 4.7] can be applied to conclude that $\lim_{k \rightarrow \infty} \|\mathbf{u}^{k+1} - \mathbf{y}^k\|^2 = 0$. Then, the statements recalled in Remark 1 guarantee convergence of $(\mathbf{x}^k)_{k \in \mathbb{N}}$ to a v-SGNE of the SGNEP in (6). \blacksquare

Remark 3: Although our proof follows similar steps to [5, Prop. 1 and 2], note that their component-wise approach can be used basing on the fact that in SNEPs the feasible set reduces to the Cartesian product of the agents' local sets Ω_i , $i \in \mathcal{I}$. However, in our generalized setting, the nature of the coupling constraints in (7) makes this assumption not necessarily true. Moreover, this means that our proof applies to the general case of finding a zero of a monotone inclusion $0 \in \bar{\mathcal{A}}(\mathbf{u}) + \bar{\mathcal{B}}(\mathbf{u})$, independently from the fact that such an inclusion comes from the KKT conditions of the game in (11). Finally, this also explains why we resort to the step size matrix Φ_k with structure as defined in (17). \square

B. Discussion on the variable step size sequence

We traditionally identify two main approaches to perform stochastic approximations: either we take one sample as considered in (14), or we take the average over an increasing (possibly *infinite*) number of realizations. The idea behind having a large number of samples is that the variance of the stochastic error disappears with the number of iterations [6], [7]. By following classic results in convergence analysis [2], [3], [8], however, it turns out that this latter approach is computationally expensive, or even unrealistic in some cases, and hence taking just *one* or a *finite* number of samples is preferable in practice. This practical simplification comes at the price of choosing a vanishing step size to control the approximation error, which in our case corresponds to the time-varying matrix Φ_k , thus possibly involving a time-varying metric for the convergence analysis. Although convergence can be guaranteed in such cases [27], [28], some additional assumptions on the metric should be satisfied. Specifically, the matrix Φ_k should be chosen so that

$$\sup_{k \in \mathbb{N}} \|\Phi_k^{-1}\| < \infty \text{ and } \forall k > 0 \ (1 + \eta_k) \Phi_{k+1}^{-1} \succcurlyeq \Phi_k^{-1}, \quad (18)$$

where $(\eta^k)_{k \in \mathbb{N}}$ is a nonnegative sequence such that $\sum_k \eta^k < \infty$. Unfortunately, this contradicts the fact that the step size sequence should be decreasing. Loosely speaking, the motivation for (18) stems from the fact that, with a variable metric, it is hard to prove whether the algorithm converges

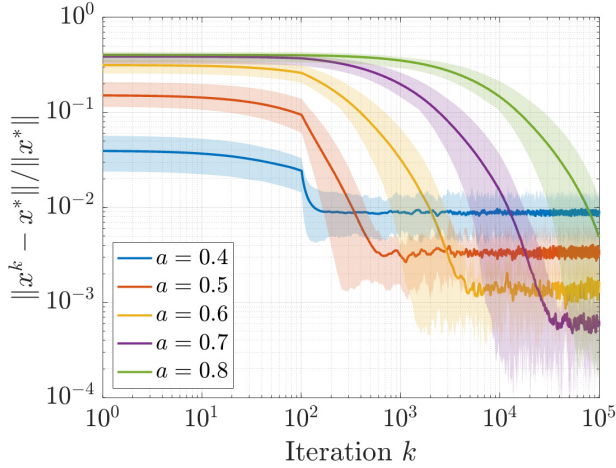


Fig. 1: Effect of the parameter η of step size, according to Assumption 7, on the convergence rate.




Fig. 2: Effect of the parameter η of step size, according to Assumption 7, on the convergence rate.

to a zero of the mapping or to a zero of the step sequences. On the other hand, given the specific structure of our matrix Φ_k , we overcome this issue by considering a fixed matrix Φ that we can use as a metric, and then by pre-multiplying Φ with a vanishing step as in (17). Note that this formulation allows us to preserve the distributed nature of the algorithm.

We also note that a “separable” matrix as Φ_k in (17) cannot be used for every iterative distributed algorithm. In the standard SFB [3], for instance, the matrix Φ_k serves as a preconditioning matrix that has non-zero off-diagonal entries. Pre-multiplying such entries for a quantity (even fixed) would compromise the corresponding KKT conditions, thus making impossible to produce distributed iterations.

V. NUMERICAL SIMULATIONS

We now validate both the model in §II and Algorithm 1 numerically. Specifically, we consider an instance of the competition among $N = 5$ on-demand ride-hailing firms over $|\mathcal{H}| = 10$ areas with main parameters in Table II.

TABLE II: Simulation parameters

Parameter	Unit	Description	Value
\bar{p}	\$	Maximum service price	35
\bar{p}_h	\$	Area price cap	$\sim \mathcal{U}(0.65\bar{p}, 0.95\bar{p})$
β		Commission rate	0.9
w	\$	Wage lower bound	12
θ_i		Competition parameters	$\sim \mathcal{U}(0.6, 1)$
C_h		Area costumers' demand	$\sim \mathcal{U}(5, 12) \times 10^3$
$K_{i,h}$		Registered drivers	$\sim \mathcal{U}(0.5, 3) \times 10^3$




Fig. 3: Impact of the substitutability parameter on the averaged percentage of expected profit, for each firm $i \in \mathcal{I}$.




Fig. 4: Impact of the substitutability parameter on the averaged costumers' demand satisfaction, for each area $h \in \mathcal{H}$.

First, we test the effect of the step size sequence on the convergence of Algorithm 1 by recalling that, in view of Assumption 7, $\alpha^k = (k + \eta)^{-a}$. In particular, Fig. 1 shows how the rate of convergence is affected by the exponent a while in Fig. 2 we plot the effect of the base η . In the first example we choose $\eta = 10^8$ and $\varepsilon^k = (k + 10^6)^{-0.15}$, while for the second one $a = 0.7$ and $\varepsilon^k = (k + 10^6)^{-0.2}$. The thick lines indicate the average performance while the transparent areas are the variability over 10 runs of the algorithm.

Successively, we examine how the competition parameters θ_i , $i \in \mathcal{I}$, affect the equilibrium solution of the SGNEP in (5) against $K = 1000$ random realizations of $\delta_{i,h}$, i.e., the willingness of the driver to provide service. In view of (4), given some drivers' expectation level $\delta_{i,h}^{(j)} \sim \mathcal{U}(\underline{w}, \bar{w}_{i,j})$ and an equilibrium strategy profile $\mathbf{x}^* = \text{col}((x_i^*)_{i \in \mathcal{I}})$ with $x_i^* = \text{col}((p_{i,h}^*)_{h \in \mathcal{H}})$, if $\delta_{i,h}^{(j)} \leq \beta p_{i,h}^*$ then $k_{i,j}^{(j)} = K_{i,j}$. The number of drivers providing a service allows each firm to actually request a payment from the costumers, thus earning $\tilde{J}_i^{(j)} := \mathbb{E}_\xi[\sum_{h \in \mathcal{H}} p_{i,h}^* (d_{i,h}^{e,(j)}(\xi) - \beta k_{i,h}^{(j)})]$. This value, however, can be different (possibly smaller) from the expected profit $J_i := \mathbb{E}_\xi[\sum_{h \in \mathcal{H}} p_{i,h}^* (d_{i,h}^e(\xi) - \beta k_{i,h})]$, i.e., the value function associated to the equilibrium condition. Then, for each firm $i \in \mathcal{I}$, Fig. 3 shows how the ratio between these two quantities, resulting into a percentage of expected profit, changes with the level of competition in the ride-hailing market. Specifically, we note that when firms operate in an almost oligopoly regime, i.e., $\theta_i \in [0.2, 0.4]$ for all $i \in \mathcal{I}$, the averaged percentage is small and grows when the market becomes more competitive. For small values of θ_i , indeed, each company is almost independent and tends to select prices to match the lower bound \underline{w}/β – this also coincides with the strategy adopted by firms with only few registered drivers, thus explaining the behavior of firms 2 and 5 – whereas for larger values of θ_i the firms are entitled to significantly raise their prices, thus allowing to meet drivers' expectations with a higher probability. Such trend is confirmed by the numerical results shown in Fig. 4, which reports the costumers' demand satisfaction, measured as $\sum_{i \in \mathcal{I}} k_{i,h}^{(j)}/C_h$, for each area $h \in \mathcal{H}$. In almost oligopoly regimes the costumers' request is met few times only, while it grows significantly when also the competition increases.

VI. CONCLUSION

We have proposed a model for the ride-hailing market under a regulated pricing scenario involving several platforms that compete to offer mobility services. The model takes into account the natural uncertainty of the demand and the need for the platforms to set up suitable pricing strategies to attract riders and fair wages to secure a fleet of drivers. To optimize the operations of these interdependent problems, we have recasted the model as a stochastic Nash equilibrium problem for which we have proposed a distributed, Tikhonov regularization-based algorithm that enjoys convergence guarantees to a Nash equilibrium. In particular, the proposed equilibrium seeking method leverages only a finite number of samples of the uncertainty to perform the stochastic approximation, as well as it requires mere monotonicity of the pseudogradient mapping to establish convergence.

REFERENCES

- [1] NYC OpenData. (2022) For-hire vehicle base aggregate report. [Online] <https://opendata.cityofnewyork.us/>.
- [2] P. Yi and L. Pavel, "An operator splitting approach for distributed generalized Nash equilibria computation," *Automatica*, vol. 102, pp. 111–121, 2019.
- [3] B. Franci and S. Grammatico, "A distributed forward-backward algorithm for stochastic generalized Nash equilibrium seeking," *IEEE Transactions on Automatic Control*, 2020.

- [4] H. Robbins and S. Monro, "A stochastic approximation method," *The Annals of Mathematical Statistics*, pp. 400–407, 1951.
- [5] J. Koshal, A. Nedic, and U. V. Shanbhag, "Regularized iterative stochastic approximation methods for stochastic variational inequality problems," *IEEE Transactions on Automatic Control*, vol. 58, no. 3, pp. 594–609, 2013.
- [6] A. Iusem, A. Jofré, R. I. Oliveira, and P. Thompson, "Extragradient method with variance reduction for stochastic variational inequalities," *SIAM Journal on Optimization*, vol. 27, no. 2, pp. 686–724, 2017.
- [7] R. I. Boj, P. Mertikopoulos, M. Staudigl, and P. T. Vuong, "Minibatch forward-backward-forward methods for solving stochastic variational inequalities," *Stochastic Systems*, vol. 11, no. 2, pp. 112–139, 2021.
- [8] B. Franci and S. Grammatico, "Stochastic generalized Nash equilibrium seeking in merely monotone games," *IEEE Transactions on Automatic Control*, 2021.
- [9] —, "Distributed projected–reflected–gradient algorithms for stochastic generalized Nash equilibrium problems," in *2021 European Control Conference (ECC)*. IEEE, 2021, pp. 369–374.
- [10] F. Facchinei and J.-S. Pang, *Finite-dimensional variational inequalities and complementarity problems*. Springer Science & Business Media, 2007.
- [11] A. Kannan and U. V. Shanbhag, "Distributed computation of equilibria in monotone Nash games via iterative regularization techniques," *SIAM Journal on Optimization*, vol. 22, no. 4, pp. 1177–1205, 2012.
- [12] F. Facchinei and C. Kanzow, "Generalized Nash equilibrium problems," *Annals of Operations Research*, vol. 175, no. 1, pp. 177–211, 2010.
- [13] Y. Zhong, T. Yang, B. Cao, and T. Cheng, "On-demand ride-hailing platforms in competition with the taxi industry: Pricing strategies and government supervision," *International Journal of Production Economics*, vol. 243, p. 108301, 2022.
- [14] E. J. He, S. Savin, J. Goh, and C.-P. Teo, "Off-platform threats in on-demand services," *Available at SSRN 3550646*, 2020.
- [15] K. Bimpikis, O. Candogan, and D. Saban, "Spatial pricing in ride-sharing networks," *Operations Research*, vol. 67, no. 3, pp. 744–769, 2019.
- [16] R. Beer, C. Brakewood, S. Rahman, and J. Viscardi, "Qualitative analysis of ride-hailing regulations in major American cities," *Transportation Research Record*, vol. 2650, no. 1, pp. 84–91, 2017.
- [17] T. W. McGuire and R. Staelin, "An industry equilibrium analysis of downstream vertical integration," *Marketing Science*, vol. 2, no. 2, pp. 161–191, 1983.
- [18] V. Pandey, J. Monteil, C. Gambella, and A. Simonetto, "On the needs for MaaS platforms to handle competition in ridesharing mobility," *Transportation Research Part C: Emerging Technologies*, vol. 108, pp. 269–288, 2019.
- [19] F. Fabiani, A. Simonetto, and P. J. Goulart, "Personalized incentives as feedback design in generalized Nash equilibrium problems," *IEEE Transactions on Automatic Control*, 2021, (Under review – available at arxiv.org/abs/2203.12948).
- [20] G. Belgioioso and S. Grammatico, "Semi-decentralized Nash equilibrium seeking in aggregative games with separable coupling constraints and non-differentiable cost functions," *IEEE Control Systems Letters*, vol. 1, no. 2, pp. 400–405, 2017.
- [21] U. Ravat and U. V. Shanbhag, "On the characterization of solution sets of smooth and nonsmooth convex stochastic Nash games," *SIAM Journal on Optimization*, vol. 21, no. 3, pp. 1168–1199, 2011.
- [22] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on stochastic programming: modeling and theory*. SIAM, 2021.
- [23] D. P. Palomar and Y. C. Eldar, *Convex optimization in signal processing and communications*. Cambridge university press, 2010.
- [24] F. Facchinei, A. Fischer, and V. Piccialli, "On generalized Nash games and variational inequalities," *Operations Research Letters*, vol. 35, no. 2, pp. 159–164, 2007.
- [25] A. Auslender and M. Teboulle, "Lagrangian duality and related multiplier methods for variational inequality problems," *SIAM Journal on Optimization*, vol. 10, no. 4, pp. 1097–1115, 2000.
- [26] B. Franci and S. Grammatico, "A damped forward–backward algorithm for stochastic generalized Nash equilibrium seeking," in *2020 European Control Conference (ECC)*. IEEE, 2020, pp. 1117–1122.
- [27] —, "Convergence of sequences: A survey," *Annual Reviews in Control*, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1367578822000037>
- [28] P. L. Combettes and B. C. Vũ, "Variable metric forward–backward splitting with applications to monotone inclusions in duality," *Optimization*, vol. 63, no. 9, pp. 1289–1318, 2014.