

Cooperative data-driven distributionally robust optimization

Ashish Cherukuri Jorge Cortés

Abstract—This paper studies a class of multiagent stochastic optimization problems where the objective is to minimize the expected value of a function which depends on a random variable. The probability distribution of the random variable is unknown to the agents, so each one gathers samples of it. The agents aim to cooperatively find, using their data, a solution to the optimization problem with guaranteed out-of-sample performance. The approach is to formulate a data-driven distributionally robust optimization problem using Wasserstein ambiguity sets, which turns out to be equivalent to a convex program. We reformulate the latter as a distributed optimization problem and identify a convex-concave augmented Lagrangian function whose saddle points are in correspondence with the optimizers provided a min-max interchangeability criteria is met. Our distributed algorithm design then consists of the saddle-point dynamics associated to the augmented Lagrangian. We formally establish that the trajectories of the dynamics converge asymptotically to a saddle point and hence an optimizer of the problem. Finally, we provide a class of functions that meet the min-max interchangeability criteria. Simulations illustrate our results.

I. INTRODUCTION

Stochastic optimization in the context of multiagent systems has numerous applications, such as target tracking, distributed estimation, and cooperative planning and learning. Solving these problems, in an exact sense, requires the knowledge of the probability distribution of the random variables. Often this information is unavailable and instead, agents gather samples and use the data to find a solution to the stochastic optimization in an approximate sense. If the available dataset is large, machine learning algorithms are able to find the optimizer of the original problem with arbitrary precision. However, when the dataset is small, these algorithms fail to provide guarantees on the output obtained from the procedure. Scenarios with small datasets appear in applications where acquiring samples is expensive due to the size and complexity of the system or when decisions must be taken in real time, leaving less room for gathering many samples. Distributionally robust optimization (DRO), instead, uses finite datasets to provide a solution with desirable out-of-sample performance guarantees. Motivated by this, we consider here the task for a group of agents to collaboratively find a data-driven solution for a stochastic optimization problem using the tools provided by the DRO framework. Instead of breaking the problem in two separate subproblems (model and estimate first the

uncertainty using data-fusion algorithms and then use the estimates to solve the stochastic optimization problem), the DRO method tackles them jointly, providing approximations of the optimizers tailored to the quality and size of the gathered data.

Literature review: Stochastic optimization is a classical topic [2]. To the large set of methods available to solve this type of problems, a recent addition is data-driven distributionally robust optimization, see e.g. [3], [4], [5] and references therein. In this setup, the distribution of the random variable is unknown and so, a worst-case optimization is carried over a set of distributions (termed ambiguity set) that contains the true distribution with high probability. This worst-case optimization provides probabilistic performance bounds for the original stochastic optimization. One way of designing the ambiguity sets is to consider the set of distributions that are close (in some distance metric over the space of distributions) to some reference distribution constructed from the available data. Depending on the metric, one gets different ambiguity sets with different performance bounds. Some popular metrics are ϕ -divergence [6], Prokhorov metric [7], and Wasserstein distance [3]. Here, we consider ambiguity sets defined using the Wasserstein metric. Tractable reformulations for the data-driven DRO methods have been well studied. However, designing coordination algorithms to find a data-driven solution when the data is gathered in a distributed way by a network of agents has not been investigated. This is the focus of this paper. Our work has connections with the growing body of literature on distribution optimization [8], [9], [10] problems and agreement-based algorithms to solve them, see e.g., [11], [12], [13], [14], [15], [16] and references therein. Our emphasis on guarantees with small datasets is in contrast with typical setups of distributed machine learning, see e.g. [17], which assume the availability of large datasets and provide asymptotic guarantees on the learning algorithms. Nonetheless, the coordination aspect in these works is similar in spirit to what we emphasize on here.

Statement of contributions: Our starting point is a multiagent stochastic optimization problem involving the minimization of the expected value of an objective function with a decision variable and a random variable as arguments. The probability distribution of the random variable is unknown and instead, agents collect a finite set of samples of it. Given this data, each agent can individually find a data-driven solution of the stochastic optimization. However, agents wish to cooperate to leverage on the data collected by everyone in the group. Our approach consists of formulating a distributionally robust optimization problem over ambiguity sets defined as neigh-

A preliminary version of this work appeared at the 2017 Allerton Conference on Communication, Control, and Computing, Monticello, Illinois as [1].

A. Cherukuri, J. Cortés are with the Dept. of Mechanical and Aerospace Engineering, UC San Diego, {acheruku,cortes}@ucsd.edu.

borhoods of the empirical distribution under the Wasserstein metric. The solution of this problem has guaranteed out-of-sample performance for the stochastic optimization. Our first contribution is the reformulation of the DRO problem to display a structure amenable to distributed algorithm design. We achieve this by augmenting the decision variables to yield a convex optimization whose objective function is the aggregate of individual objectives and whose constraints involve consensus among neighboring agents. Building on an augmented version of the associated Lagrangian function, we identify a convex-concave function which under a min-max interchangeability condition has the property that its saddle-points are in one-to-one correspondence with the optimizers of the reformulated problem. Our second contribution is the design of the saddle-point dynamics for the identified convex-concave Lagrangian function. We show that the proposed dynamics is distributed and provably correct, in the sense that its trajectories asymptotically converge to a solution of the original stochastic optimization problem. Our third contribution is the identification of two broad class of objective functions for which the min-max interchangeability holds. The first class is the set of functions that are convex-concave in the decision and the random variable, respectively. The second class is where functions are convex-convex and have some additional structure: they are either quadratic in the random variable or they correspond to the loss function of the least-squares problem. Finally, we illustrate our results in simulation.

II. PRELIMINARIES

This section introduces notation and basic notions on graph theory, convex analysis, and stability of discontinuous dynamical systems.

1) *Notation*: Let \mathbb{R} , $\mathbb{R}_{\geq 0}$, and $\mathbb{Z}_{\geq 1}$ denote the set of real, nonnegative real, and positive integer numbers. The extended reals are denoted as $\overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty, \infty\}$. We let $\|\cdot\|$ denote the 2-norm on \mathbb{R}^n . We use the notation $B_\delta(x) := \{y \in \mathbb{R}^n \mid \|x - y\| < \delta\}$. Given $x, y \in \mathbb{R}^n$, x_i denotes the i -th component of x , and $x \leq y$ denotes $x_i \leq y_i$ for $i \in [n]$. For vectors $u \in \mathbb{R}^n$ and $w \in \mathbb{R}^m$, the vector $(u; w) \in \mathbb{R}^{n+m}$ denotes their concatenation. We use the shorthand notation $\mathbf{0}_n = (0, \dots, 0) \in \mathbb{R}^n$, $\mathbf{1}_n = (1, \dots, 1) \in \mathbb{R}^n$, and $\mathbf{I}_n \in \mathbb{R}^{n \times n}$ for the identity matrix. For $A \in \mathbb{R}^{n_1 \times n_2}$ and $B \in \mathbb{R}^{m_1 \times m_2}$, $A \otimes B \in \mathbb{R}^{n_1 m_1 \times n_2 m_2}$ is the Kronecker product. The Cartesian product of $\{\mathcal{S}_i\}_{i=1}^n$ is denoted by $\prod_{i=1}^n \mathcal{S}_i := \mathcal{S}_1 \times \dots \times \mathcal{S}_n$. The interior of a set $\mathcal{S} \subset \mathbb{R}^n$ is denoted by $\text{int}(\mathcal{S})$. For a function $f : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$, $(x, \xi) \mapsto f(x, \xi)$, we denote the partial derivative of f with respect to the first argument by $\nabla_x f$ and with respect to the second argument by $\nabla_\xi f$. The higher-order derivatives follow the convention $\nabla_{x\xi} f = \frac{\partial^2 f}{\partial x \partial \xi}$, $\nabla_{xx} f = \frac{\partial^2 f}{\partial x^2}$, and so on. Given $V : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$, we denote the δ -sublevel set as $V^{-1}(\leq \delta) := \{x \in \mathcal{X} \mid V(x) \leq \delta\}$.

2) *Graph theory*: Following [18], an *undirected graph*, or simply a *graph*, is a pair $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = [n]$ is the vertex set and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the edge set with the property that $(i, j) \in \mathcal{E}$ if and only if $(j, i) \in \mathcal{E}$. A path is an ordered sequence such that any ordered pair of vertices

appearing consecutively is an edge. A graph is *connected* if there is a path between any pair of distinct vertices. Let $\mathcal{N}_i \subseteq \mathcal{V}$ denote the set of neighbors of vertex $i \in \mathcal{V}$, i.e., $\mathcal{N}_i = \{j \in \mathcal{V} \mid (i, j) \in \mathcal{E}\}$. A *weighted graph* is a triplet $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{A})$, where $(\mathcal{V}, \mathcal{E})$ is a digraph and $\mathbf{A} \in \mathbb{R}_{\geq 0}^{n \times n}$ is the (symmetric) *adjacency matrix* of \mathcal{G} , with the property that $a_{ij} > 0$ if $(i, j) \in \mathcal{E}$ and $a_{ij} = 0$, otherwise. The *weighted degree* of $i \in [n]$ is $w_i = \sum_{j=1}^n a_{ij}$. The *weighted degree matrix* \mathbf{D} is the diagonal matrix defined by $(\mathbf{D})_{ii} = w_i$, for all $i \in [n]$. The *Laplacian matrix* is $\mathbf{L} = \mathbf{D} - \mathbf{A}$. Note that $\mathbf{L} = \mathbf{L}^\top$ and $\mathbf{L}\mathbf{1}_n = \mathbf{0}$. If \mathcal{G} is connected, then zero is a simple eigenvalue of \mathbf{L} .

3) *Convex analysis*: Here we introduce elements from convex analysis following [19]. A set $C \subset \mathbb{R}^n$ is *convex* if $(1 - \lambda)x + \lambda y \in C$ whenever $x \in C$, $y \in C$, and $\lambda \in (0, 1)$. A vector $\varphi \in \mathbb{R}^n$ is *normal* to a convex set C at a point $x \in C$ if $(y - x)^\top \varphi \leq 0$ for all $y \in C$. The set of all vectors normal to C at x , denoted $N_C(x)$, is the *normal cone* to C at x . The *affine hull* of $S \subset \mathbb{R}^n$ is the smallest affine space containing S ,

$$\text{aff}(S) := \left\{ \sum_{i=1}^k \lambda_i x_i \mid k \in \mathbb{Z}_{\geq 1}, x_i \in S, \lambda_i \in \mathbb{R}, \sum_{i=1}^k \lambda_i = 1 \right\}.$$

The *relative interior* of a convex set C is the interior of C relative to the affine hull of C . Formally,

$$\text{ri}(C) := \{x \in \text{aff}(C) \mid \exists \epsilon > 0, (x + \epsilon B_1(0)) \cap (\text{aff}(C)) \subset C\}.$$

Given a convex set C , a vector d is a *direction of recession* of C if $x + \alpha d \in C$ for all $x \in C$ and $\alpha \geq 0$.

A convex function $f : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ is *proper* if there exists $x \in \mathbb{R}^n$ such that $f(x) < +\infty$ and f does not take the value $-\infty$ anywhere in \mathbb{R}^n . The *epigraph* of f is the set

$$\text{epi} f := \{(x, \lambda) \in (\mathbb{R}^n \times \overline{\mathbb{R}}) \mid \lambda \geq f(x)\}.$$

A function f is closed if $\text{epi} f$ is a closed set. The function f is convex if and only if $\text{epi} f$ is convex. For a closed proper convex function f , a vector d is a *direction of recession* of f if $(d, 0)$ is a direction of recession of the set $\text{epi} f$. Intuitively, it is the direction along which f is monotonically non-increasing. If $f(x) \rightarrow +\infty$ whenever $\|x\| \rightarrow +\infty$, then f does not have a direction of recession.

A function $F : \mathcal{X} \times \mathcal{Y} \rightarrow \overline{\mathbb{R}}$ is *convex-concave* (on $\mathcal{X} \times \mathcal{Y}$) if, given any point $(\tilde{x}, \tilde{y}) \in \mathcal{X} \times \mathcal{Y}$, $x \mapsto F(x, \tilde{y})$ is convex and $y \mapsto F(\tilde{x}, y)$ is concave. When the space $\mathcal{X} \times \mathcal{Y}$ is clear from the context, we refer to this property as F being convex-concave in (x, y) . A point $(x_*, y_*) \in \mathcal{X} \times \mathcal{Y}$ is a *saddle point* of F over the set $\mathcal{X} \times \mathcal{Y}$ if $F(x_*, y) \leq F(x_*, y_*) \leq F(x, y_*)$, for all $x \in \mathcal{X}$ and $y \in \mathcal{Y}$. The set of saddle points of a convex-concave function F is convex. Each saddle point is a critical point of F , i.e., if F is differentiable, then $\nabla_x F(x_*, z_*) = 0$ and $\nabla_z F(x_*, z_*) = 0$. Additionally, if F is twice differentiable, then $\nabla_{xx} F(x_*, z_*) \preceq 0$ and $\nabla_{zz} F(x_*, z_*) \succeq 0$. Given a convex-concave function $F : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \overline{\mathbb{R}}$, define

$$\begin{aligned} \mathcal{X} &:= \{x \in \mathbb{R}^n \mid F(x, y) < +\infty \text{ for all } y \in \mathbb{R}^m\}, \\ \mathcal{Y} &:= \{y \in \mathbb{R}^m \mid F(x, y) > -\infty \text{ for all } x \in \mathbb{R}^n\}. \end{aligned}$$

The product set $\mathcal{X} \times \mathcal{Y}$ is called the *effective domain* of F . The sets \mathcal{X} , \mathcal{Y} and so, $\mathcal{X} \times \mathcal{Y}$ are convex. Note that F is finite

on $\mathcal{X} \times \mathcal{Y}$. If $\mathcal{X} \times \mathcal{Y}$ is nonempty, then F is called *proper*. If the following equality holds

$$\sup_{y \in \mathcal{Y}} \inf_{x \in \mathcal{X}} F(x, y) = \inf_{x \in \mathcal{X}} \sup_{y \in \mathcal{Y}} F(x, y),$$

then this common value is called the *saddle value* of F . The function F is *closed* if for any $(x, y) \in \mathcal{X} \times \mathcal{Y}$, the functions $x \mapsto F(x, y)$ and $y \mapsto -F(x, y)$ are closed.

Theorem II.1. (Existence of finite saddle value and saddle point [19, Theorem 37.3 & 37.6]): Let $F : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \overline{\mathbb{R}}$ be a closed proper convex-concave function with effective domain $\mathcal{X} \times \mathcal{Y} \subset \mathbb{R}^n \times \mathbb{R}^m$. If the following conditions hold,

- (i) The convex functions $x \mapsto F(x, y)$ for $y \in \text{ri}(\mathcal{Y})$ have no common direction of recession;
- (ii) The convex functions $y \mapsto -F(x, y)$ for $x \in \text{ri}(\mathcal{X})$ have no common direction of recession;

then the saddle value must be finite, there exists a saddle point of F in the effective domain $\mathcal{X} \times \mathcal{Y}$ and the saddle value is attained at the saddle point.

4) *Discontinuous dynamical systems*: Here we present notions of discontinuous and projected dynamical systems from [20], [21], [22]. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a Lebesgue measurable and locally bounded function, and consider

$$\dot{x} = f(x). \quad (1)$$

A map $\gamma : [0, T] \rightarrow \mathbb{R}^n$ is a (*Caratheodory*) *solution* of (1) on the interval $[0, T]$ if it is absolutely continuous on $[0, T]$ and satisfies $\dot{\gamma}(t) = f(\gamma(t))$ almost everywhere in $[0, T]$. We use the terms solution and trajectory interchangeably. A set $\mathcal{S} \subset \mathbb{R}^n$ is *invariant* under (1) if every solution starting in \mathcal{S} remains in \mathcal{S} . For a solution γ of (1) defined on the time interval $[0, \infty)$, the *omega-limit* set $\Omega(\gamma)$ is defined by

$$\Omega(\gamma) = \{y \in \mathbb{R}^n \mid \text{there exists } \{t_k\}_{k=1}^{\infty} \subset [0, \infty) \text{ with} \\ \lim_{k \rightarrow \infty} t_k = \infty \text{ and } \lim_{k \rightarrow \infty} \gamma(t_k) = y\}.$$

If the solution γ is bounded, then $\Omega(\gamma) \neq \emptyset$ by the Bolzano-Weierstrass theorem [23, p. 33]. Given a continuously differentiable function $V : \mathbb{R}^n \rightarrow \mathbb{R}$, the *Lie derivative of V along (1)* at $x \in \mathbb{R}^n$ is $\mathcal{L}_f V(x) = \nabla V(x)^\top f(x)$. The next result is a simplified version of [20, Proposition 3].

Proposition II.2. (*Invariance principle for discontinuous Caratheodory systems*): Let $\mathcal{S} \subset \mathbb{R}^n$ be compact and invariant. Assume that, for each point $x_0 \in \mathcal{S}$, there exists a unique solution of (1) starting at x_0 and that its omega-limit set is invariant too. Let $V : \mathbb{R}^n \rightarrow \mathbb{R}$ be a continuously differentiable map such that $\mathcal{L}_f V(x) \leq 0$ for all $x \in \mathcal{S}$. Then, any solution of (1) starting at \mathcal{S} converges to the largest invariant set in $\text{cl}(\{x \in \mathcal{S} \mid \mathcal{L}_f V(x) = 0\})$.

Projected dynamical systems are a particular class of discontinuous dynamical systems. Let $\mathcal{K} \subset \mathbb{R}^n$ be a closed convex set. Given a point $y \in \mathbb{R}^n$, the (point) projection of y onto \mathcal{K} is $\text{proj}_{\mathcal{K}}(y) = \arg\min_{z \in \mathcal{K}} \|z - y\|$. Note that $\text{proj}_{\mathcal{K}}(y)$ is a singleton and the map $\text{proj}_{\mathcal{K}}$ is Lipschitz on \mathbb{R}^n with constant

$L = 1$ [24, Proposition 2.4.1]. Given $x \in \mathcal{K}$ and $v \in \mathbb{R}^n$, the (vector) projection of v at x with respect to \mathcal{K} is

$$\Pi_{\mathcal{K}}(x, v) = \lim_{\delta \rightarrow 0^+} \frac{\text{proj}_{\mathcal{K}}(x + \delta v) - x}{\delta}.$$

Given a vector field $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and a closed convex polyhedron $\mathcal{K} \subset \mathbb{R}^n$, the associated projected dynamical system is

$$\dot{x} = \Pi_{\mathcal{K}}(x, f(x)), \quad x(0) \in \mathcal{K}, \quad (2)$$

One can verify easily that for any $x \in \mathcal{K}$, there exists an element φ_x belonging to the normal cone $N_{\mathcal{K}}(x)$ such that $\Pi_{\mathcal{K}}(x, f(x)) = f(x) - \varphi_x$. In particular, if x is in the interior of \mathcal{K} , then this element is the zero vector and we have $\Pi_{\mathcal{K}}(x, f(x)) = f(x)$. At any boundary point of \mathcal{K} , the projection operator restricts the flow of the vector field f such that the solutions of (2) remain in \mathcal{K} . Due to the projection, the dynamics (2) is in general discontinuous.

III. DATA-DRIVEN STOCHASTIC OPTIMIZATION

This section sets the stage for the formulation of our approach to deal with data-driven optimization in a distributed manner. The following material on data-driven stochastic optimization is taken from [3] and included here to provide a self-contained exposition. The reader familiar with these notions and tools can safely skip this section.

Let (Ω, \mathcal{F}, P) be a probability space and ξ be a random variable mapping this space to $(\mathbb{R}^m, B_{\sigma}(\mathbb{R}^m))$, where $B_{\sigma}(\mathbb{R}^m)$ is the Borel σ -algebra on \mathbb{R}^m . Let \mathbb{P} and $\Xi \subseteq \mathbb{R}^m$ be the distribution and the support of the random variable ξ . Assume that Ξ is closed and convex. Consider the stochastic optimization problem

$$\inf_{x \in \mathcal{X}} \mathbb{E}_{\mathbb{P}}[f(x, \xi)], \quad (3)$$

where $\mathcal{X} \subseteq \mathbb{R}^n$ is a closed convex set, $f : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$ is a continuous function, and $\mathbb{E}_{\mathbb{P}}[\cdot]$ is the expectation under the distribution \mathbb{P} . Assume that \mathbb{P} is unknown and so, solving (3) is not possible. However, we are given N independently drawn samples $\widehat{\Xi} := \{\widehat{\xi}^k\}_{k=1}^N \subset \Xi$ of the random variable ξ . Note that, until it is revealed, $\widehat{\Xi}$ is a random object with probability distribution $\mathbb{P}^N := \prod_{i=1}^N \mathbb{P}$ supported on $\Xi^N := \prod_{i=1}^N \Xi$. The objective is to find a *data-driven* solution of (3), denoted $\widehat{x}_N \in \mathcal{X}$, constructed using the dataset $\widehat{\Xi}$, that has desirable properties for the expected cost $\mathbb{E}_{\mathbb{P}}[f(\widehat{x}_N, \xi)]$ under a new sample. The property we are looking for is the *finite-sample guarantee* given by

$$\mathbb{P}^N \left(\mathbb{E}_{\mathbb{P}}[f(\widehat{x}_N, \xi)] \leq \widehat{J}_N \right) \geq 1 - \beta, \quad (4)$$

where \widehat{J}_N might also depend on the training dataset and $\beta \in (0, 1)$ is the parameter which governs \widehat{x}_N and \widehat{J}_N . The quantities \widehat{J}_N and $1 - \beta$ are referred to as the *certificate* and the *reliability* of the performance of \widehat{x}_N . The goal is to find a data-driven solution with a low certificate and a high reliability. To do so, we use the available information $\widehat{\Xi}_N$. The strategy is to determine a set $\widehat{\mathcal{P}}_N$ of probability distributions supported on Ξ that contain the true distribution \mathbb{P} with high confidence. The set $\widehat{\mathcal{P}}_N$ is referred to as the *ambiguity* set. Once such a set

is designed, the certificate \widehat{J}_N is defined as the optimal value of the following *distributionally robust optimization* problem

$$\widehat{J}_N := \inf_{x \in \mathcal{X}} \sup_{\mathbb{Q} \in \widehat{\mathcal{P}}_N} \mathbb{E}_{\mathbb{Q}}[f(x, \xi)]. \quad (5)$$

This is the worst-case optimal value considering all distributions in $\widehat{\mathcal{P}}_N$. A good candidate for $\widehat{\mathcal{P}}_N$ is the set of distributions that are close (under a certain metric) to the uniform distribution on $\widehat{\Xi}_N$, termed the *empirical distribution*. Formally, the empirical distribution is

$$\widehat{\mathbb{P}}_N := \frac{1}{N} \sum_{k=1}^N \delta_{\widehat{\xi}^k}, \quad (6)$$

where $\delta_{\widehat{\xi}^k}$ is the unit point mass at $\widehat{\xi}^k$. Let $\mathcal{M}(\Xi)$ be the space of probability distributions \mathbb{Q} supported on Ξ with finite first moment, i.e., $\mathbb{E}_{\mathbb{Q}}[\|\xi\|] = \int_{\Xi} \|\xi\| \mathbb{Q}(d\xi) < +\infty$. The 2-Wasserstein metric¹ $d_{W_2} : \mathcal{M}(\Xi) \times \mathcal{M}(\Xi) \rightarrow \mathbb{R}_{\geq 0}$ is

$$d_{W_2}(\mathbb{Q}_1, \mathbb{Q}_2) = \left(\inf \left\{ \int_{\Xi^2} \|\xi_1 - \xi_2\|^2 \Pi(d\xi_1, d\xi_2) \mid \Pi \in \mathcal{H}(\mathbb{Q}_1, \mathbb{Q}_2) \right\} \right)^{\frac{1}{2}}, \quad (7)$$

where $\mathcal{H}(\mathbb{Q}_1, \mathbb{Q}_2)$ is the set of all distributions on $\Xi \times \Xi$ with marginals \mathbb{Q}_1 and \mathbb{Q}_2 . Given $\epsilon \geq 0$, we use the notation

$$\mathcal{B}_{\epsilon}(\widehat{\mathbb{P}}_N) := \{\mathbb{Q} \in \mathcal{M}(\Xi) \mid d_{W_2}(\widehat{\mathbb{P}}_N, \mathbb{Q}) \leq \epsilon\} \quad (8)$$

to define the set of distributions that are ϵ -close to $\widehat{\mathbb{P}}_N$ under the defined metric. Let $\mathcal{M}_{\text{lt}}(\Xi) \subset \mathcal{M}(\Xi)$ be the set of *light-tailed* distributions $\mathbb{P} \in \mathcal{M}_{\text{lt}}(\Xi)$ for which there exists an exponent $a > 2$ such that

$$A := \mathbb{E}_{\mathbb{P}}[\exp(\|\xi\|^a)] = \int_{\Xi} \exp(\|\xi\|^a) \mathbb{P}(d\xi) < +\infty.$$

The next result gives a lower bound on the probability with which the true distribution \mathbb{P} is ϵ -close to $\widehat{\mathbb{P}}_N$.

Theorem III.1. (*Finite-sample guarantee of \mathbb{P} belonging to the Wasserstein ambiguity set*): Let $\mathbb{P} \in \mathcal{M}_{\text{lt}}(\Xi)$. Then,

$$\mathbb{P}^N(d_{W_2}(\mathbb{P}, \widehat{\mathbb{P}}_N) \geq \epsilon) \leq \begin{cases} c_1 e^{-c_2 N \epsilon^{\max\{4, m\}}}, & \text{if } \epsilon \leq 1, \\ c_1 e^{-c_2 N \epsilon^a}, & \text{if } \epsilon > 1, \end{cases} \quad (9)$$

for all $N \geq 1$, $m \neq 4$, and $\epsilon > 0$, where c_1, c_2 are positive constants that only depend on a, A , and m .

The proof is a direct application of [25, Theorem 2]. This result gives a method to construct the ambiguity set $\widehat{\mathcal{P}}_N$. Equating the right-hand side of (9) with the chosen $\beta \in (0, 1)$, we define for each $N \in \mathbb{Z}_{\geq 1}$,

$$\epsilon_N(\beta) := \begin{cases} \left(\frac{\log(c_1 \beta^{-1})}{c_2 N} \right)^{1/\max\{4, m\}}, & \text{if } N \geq \frac{\log(c_1 \beta^{-1})}{c_2}, \\ \left(\frac{\log(c_1 \beta^{-1})}{c_2 N} \right)^{1/a}, & \text{if } N < \frac{\log(c_1 \beta^{-1})}{c_2}. \end{cases} \quad (10)$$

¹We note that [3] employs the 1-Wasserstein metric instead of the 2-Wasserstein metric considered here.

Plugging this value in (9) yields $\mathbb{P}^N(d_{W_2}(\mathbb{P}, \widehat{\mathbb{P}}_N) \geq \epsilon_N(\beta)) \leq \beta$. That is, if we let $\widehat{\mathcal{P}}_N := \mathcal{B}_{\epsilon_N(\beta)}(\widehat{\mathbb{P}}_N)$, then

$$\mathbb{P}^N(\mathbb{P} \in \widehat{\mathcal{P}}_N) \geq 1 - \beta, \quad (11)$$

i.e., the true distribution belongs to the ambiguity set with probability at least $1 - \beta$. This leads us to the following result.

Theorem III.2. (*Finite-sample guarantee of (5) with $\widehat{\mathcal{P}}_N = \mathcal{B}_{\epsilon_N(\beta)}(\widehat{\mathbb{P}}_N)$*): For $\mathbb{P} \in \mathcal{M}_{\text{lt}}(\Xi)$ and $\beta \in (0, 1)$, let \widehat{J}_N and \widehat{x}_N be the optimal value and an optimizer of the distributionally robust optimization (5) with $\widehat{\mathcal{P}}_N = \mathcal{B}_{\epsilon_N(\beta)}(\widehat{\mathbb{P}}_N)$. Then, the finite-sample guarantee (4) holds.

The proof follows by using (5) and (11) to yield (4). We end this section by discussing the tractability of solving (5) with $\widehat{\mathcal{P}}_N = \mathcal{B}_{\epsilon_N(\beta)}(\widehat{\mathbb{P}}_N)$.

Theorem III.3. (*Tractable reformulation of (5)*): Assume that for all $\xi \in \Xi$, the map $x \mapsto f(x, \xi)$ is convex. Then, for any $\beta \in (0, 1)$ and $N \in \mathbb{Z}_{\geq 1}$, the optimal value of (5) with $\widehat{\mathcal{P}}_N = \mathcal{B}_{\epsilon_N(\beta)}(\widehat{\mathbb{P}}_N)$ is equal to the optimum of the following convex optimization problem

$$\inf_{\lambda \geq 0, x \in \mathcal{X}} \left\{ \lambda \epsilon_N^2(\beta) + \frac{1}{N} \sum_{k=1}^N \max_{\xi \in \widehat{\Xi}} \left(f(x, \xi) - \lambda \|\xi - \widehat{\xi}^k\|^2 \right) \right\}.$$

Theorem III.3 shows that under mild conditions on the objective function, one can reformulate the distributionally robust optimization problem as a convex optimization problem. This result plays a key role in our forthcoming discussion.

IV. PROBLEM STATEMENT

Consider $n \in \mathbb{Z}_{\geq 1}$ agents communicating over an undirected weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{A})$. The set of vertices are enumerated as $\mathcal{V} := [n]$. Each agent $i \in [n]$ can send and receive information from its neighbors \mathcal{N}_i in \mathcal{G} . Let $f : \mathbb{R}^d \times \mathbb{R}^m \rightarrow \mathbb{R}$, $(x, \xi) \mapsto f(x, \xi)$, be a continuously differentiable objective function. Assume that for any $\xi \in \mathbb{R}^m$, the map $x \mapsto f(x, \xi)$ is convex and that for any $x \in \mathbb{R}^d$, the map $\xi \mapsto f(x, \xi)$ is either convex or concave. Suppose that the set of $\xi \in \mathbb{R}^m$ for which $\mathbf{1}_n$ and $-\mathbf{1}_n$ are not a direction of recession for the convex function $x \mapsto f(x, \xi)$ is dense in \mathbb{R}^m . As we progress, we stipulate additional conditions on f as necessary. Assume that all agents know the objective function f . Given a random variable $\xi \in \mathbb{R}^m$ with support \mathbb{R}^m and distribution \mathbb{P} , the original objective for the agents is to solve the following stochastic optimization problem

$$\inf_{x \in \mathbb{R}^d} \mathbb{E}_{\mathbb{P}}[f(x, \xi)]. \quad (12)$$

However, agents cannot solve this problem as they do not know \mathbb{P} . Instead, each agent has a certain number (at least one) of independent and identically distributed realizations of the random variable ξ . We denote the data available to agent i by $\widehat{\Xi}_i$ that is assumed to be nonempty. Assume that $\widehat{\Xi}_i \cap \widehat{\Xi}_j = \emptyset$ for all $i, j \in [n]$ and let $\widehat{\Xi} = \cup_{i=1}^n \widehat{\Xi}_i$ containing N samples be the available data set.

The goal for the agents is then to collectively find, in a distributed manner, a data-driven solution $\widehat{x}_N \in \mathbb{R}^d$ to

approximate the optimizer of (12) with guaranteed performance bounds. To achieve this, we rely on the framework of distributionally robust optimization, cf. Section II-2. From Theorem III.3, a data-driven solution for (12) can be obtained by solving the following convex optimization problem

$$\inf_{\lambda \geq 0, x} \left\{ \lambda \epsilon_N^2(\beta) + \frac{1}{N} \sum_{k=1}^N \max_{\xi \in \mathbb{R}^m} \left(f(x, \xi) - \lambda \|\xi - \widehat{\xi}^k\|^2 \right) \right\} \quad (13)$$

where $\beta \in (0, 1)$ and $\epsilon_N(\beta)$ is given in (10). We assume that there exists a finite optimizer (x^*, λ^*) of (13), e.g., one of the conditions for existence of finite optimizers given in [26] is met. This optimizer has the following out-of-sample performance guarantee

$$\mathbb{P}^N \left(\mathbb{E}_{\mathbb{P}}[f(x^*, \xi)] \leq \widehat{J}_N \right) \geq 1 - \beta,$$

where \widehat{J}_N is the optimum value (13). Note that each agent can individually find a data-driven solution to (12) by using only the data available to it in the convex formulation (13). However, such a solution in general will have an inferior out-of-sample guarantee as compared to the one obtained collectively. In the cooperative setting, agents aim to solve (13) in a distributed manner, that is

- (i) each agent i has the information

$$\mathcal{I}_i := \{\widehat{\Xi}_i, f, a, c_1, c_2, A, \beta, n, N\}, \quad (14)$$

where a, c_1, c_2 , and A are parameters associated with the distribution \mathbb{P} , as defined in Section II-2 and $\beta \in (0, 1)$ is a parameter that agents agree upon beforehand,

- (ii) each agent i can only communicate with its neighbors \mathcal{N}_i in the graph \mathcal{G} ,
 (iii) each agent i does not share with its neighbors any element of the dataset $\widehat{\Xi}_i$ available to it, and
 (iv) there is no central coordinator or leader that can communicate with all agents.

The challenge in solving (13) in a distributed manner lies in the fact that the data is distributed over the network and the optimizer x^* depends on it all. Moreover, the inner maximization can be a nonconvex problem, in general. One way of solving (13) in a cooperative fashion is to let agents share their data with everyone in the network via some sort of flooding mechanism. This violates item (iii) of our definition of distributed algorithm given above. We specifically keep such methods out of scope due to two reasons. First, the data would not be private anymore, creating a possibility of adversarial action. Second, the communication burden of such a strategy is higher than our proposed distributed strategy when the size of the network and the dataset grows along the execution of the algorithm.

Our strategy to tackle the problem is organized as follows: in Section V we reformulate the problem (13) to obtain a structure which allows us in Section VI to propose our distributed algorithm. Section VII discusses a class of objective functions f for which the distributed algorithm provably converges.

V. DISTRIBUTED PROBLEM FORMULATION AND SADDLE POINTS

This section studies the structure of the optimization problem presented in Section IV with the ulterior goal of facilitating the design of a distributed algorithmic solution. Our first step is a reformulation of (13) that, by augmenting the decision variables of the agents, yields an optimization where the objective function is the aggregate of individual functions (that can be independently evaluated by the agents) and constraints which display a distributed structure. Our second step is the identification of a convex-concave function whose saddle points are the primal-dual optimizers of the reformulated problem under suitable conditions on the objective function f . This opens the way to consider the associated saddle-point dynamics as our candidate distributed algorithm. The structure of the original optimization problem makes this step particularly nontrivial.

A. Reformulation as distributed optimization problem

We have each agent $i \in [n]$ maintain a copy of λ and x , denoted by $\lambda^i \in \mathbb{R}$ and $x^i \in \mathbb{R}^d$, respectively. Thus, the decision variables for i are (x^i, λ^i) . For notational ease, let the concatenated vectors be $\lambda_v := (\lambda^1; \dots; \lambda^n)$, and $x_v := (x^1; \dots; x^n)$. Let $v_k \in [n]$ be the agent that holds the k -th sample $\widehat{\xi}^k$ of the dataset. Consider the following convex optimization problem

$$\min_{x_v, \lambda_v \geq \mathbf{0}_n} h(\lambda_v) + \frac{1}{N} \sum_{k=1}^N \max_{\xi \in \mathbb{R}^m} g_k(x^{v_k}, \lambda^{v_k}, \xi) \quad (15a)$$

$$\text{subject to } \mathbf{L}\lambda_v = \mathbf{0}_n, \quad (15b)$$

$$(\mathbf{L} \otimes \mathbf{I}_d)x_v = \mathbf{0}_{nd}. \quad (15c)$$

where $\mathbf{L} \in \mathbb{R}^{n \times n}$ is the Laplacian of the graph \mathcal{G} and we have used the shorthand notation $h: \mathbb{R}^n \rightarrow \mathbb{R}$ for

$$h(\lambda_v) := \frac{\epsilon_N^2(\beta)(\mathbf{1}_n^\top \lambda_v)}{n}$$

and, for each $k \in [N]$, $g_k: \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^m \rightarrow \mathbb{R}$ for

$$g_k(x, \lambda, \xi) := f(x, \xi) - \lambda \|\xi - \widehat{\xi}^k\|^2.$$

The following result establishes the correspondence between the optimizers of (13) and (15), respectively.

Lemma V.1. (One-to-one correspondence between optimizers of (13) and (15)): *The following holds:*

- (i) *If (x^*, λ^*) is an optimizer of (13), then $(\mathbf{1}_n \otimes x^*, \lambda^* \mathbf{1}_n)$ is an optimizer of (15).*
 (ii) *If (x_v^*, λ_v^*) is an optimizer of (15), then there exists an optimizer (x^*, λ^*) of (13) such that $x_v^* = \mathbf{1}_n \otimes x^*$ and $\lambda_v^* = \lambda^* \mathbf{1}_n$.*

Proof: The proof follows by noting that \mathcal{G} is connected and hence, (i) $\mathbf{L}\lambda_v = \mathbf{0}_n$ if and only if $\lambda_v = \alpha \mathbf{1}_n$, $\alpha \in \mathbb{R}$; and (ii) $(\mathbf{L} \otimes \mathbf{I}_d)x_v = \mathbf{0}_{nd}$ if and only if $x_v = \mathbf{1}_n \otimes x$, $x \in \mathbb{R}^d$. ■

Note that constraints (15b) and (15c) force agreement and that each of their components is computable by an agent of the

network using only local information. Moreover, the objective function (15a) can be written as $\sum_{i=1}^n J_i(x^i, \lambda^i, \widehat{\Xi}_i)$, where

$$J_i(x^i, \lambda^i, \widehat{\Xi}_i) := \frac{\epsilon_N^2(\beta)\lambda^i}{n} + \frac{1}{N} \sum_{k: \widehat{\xi}^k \in \widehat{\Xi}_i} \max_{\xi \in \mathbb{R}^m} g_k(x^i, \lambda^i, \xi),$$

for all $i \in [n]$. Therefore, the problem (15) has the adequate structure from a distributed optimization viewpoint: an aggregate objective function and locally computable constraints.

B. Augmented Lagrangian and saddle points

Our next step is to identify an appropriate variant of the Lagrangian function of (15) with the following two properties: (i) it does not consist of an inner maximization, unlike the objective in (15a), and (ii) the primal-dual optimizers of (15) are saddle points of the newly introduced function. The availability of these two facts sets the stage for our ensuing algorithm design.

To proceed further, we first denote for convenience the objective function (15a) with $F : \mathbb{R}^{nd} \times \mathbb{R}_{\geq 0}^n \rightarrow \mathbb{R}$,

$$F(x_v, \lambda_v) := h(\lambda_v) + \frac{1}{N} \sum_{k=1}^N \max_{\xi \in \mathbb{R}^m} g_k(x^{v_k}, \lambda^{v_k}, \xi). \quad (16)$$

Note that the Lagrangian of (15) is $L : \mathbb{R}^{nd} \times \mathbb{R}_{\geq 0}^n \times \mathbb{R}^n \times \mathbb{R}^{nd} \rightarrow \mathbb{R}$,

$$L(x_v, \lambda_v, \nu, \eta) := F(x_v, \lambda_v) + \nu^\top L \lambda_v + \eta^\top (L \otimes I_d) x_v, \quad (17)$$

where $\nu \in \mathbb{R}^n$ and $\eta \in \mathbb{R}^{nd}$ are dual variables corresponding to the equality constraints (15b) and (15c), respectively. L is convex-concave in $((x_v, \lambda_v), (\nu, \eta))$ on the domain $\lambda_v \geq \mathbf{0}_n$. The next result establishes important properties of the Lagrangian giving a correspondence between its saddle points and the optimizers of (15). We provide its proof, rather than establishing it as a consequence of results from the literature, because the objective (15a) takes values in the extended reals and in general might not be closed.

Lemma V.2. (Min-max equality for L): *The set of saddle points of L over the domain $(\mathbb{R}^{nd} \times \mathbb{R}_{\geq 0}^n) \times (\mathbb{R}^n \times \mathbb{R}^{nd})$ is nonempty and*

$$\inf_{x_v, \lambda_v \geq \mathbf{0}_n} \sup_{\nu, \eta} L(x_v, \lambda_v, \nu, \eta) = \sup_{\nu, \eta} \inf_{x_v, \lambda_v \geq \mathbf{0}_n} L(x_v, \lambda_v, \nu, \eta). \quad (18)$$

Furthermore, the following holds:

- (i) *If $(\bar{x}_v, \bar{\lambda}_v, \bar{\nu}, \bar{\eta})$ is a saddle point of L over $(\mathbb{R}^{nd} \times \mathbb{R}_{\geq 0}^n) \times (\mathbb{R}^n \times \mathbb{R}^{nd})$, then $(\bar{x}_v, \bar{\lambda}_v)$ is an optimizer of (15).*
- (ii) *If $(\bar{x}_v, \bar{\lambda}_v)$ is an optimizer of (15), then there exists $(\bar{\nu}, \bar{\eta})$ such that $(\bar{x}_v, \bar{\lambda}_v, \bar{\nu}, \bar{\eta})$ is a saddle point of L over $(\mathbb{R}^{nd} \times \mathbb{R}_{\geq 0}^n) \times (\mathbb{R}^n \times \mathbb{R}^{nd})$.*

Proof: Consider the left-hand side of (18). For a fixed x_v and $\lambda_v \geq \mathbf{0}_n$, if $L \lambda_v \neq \mathbf{0}_n$ or $(L \otimes I_d) x_v \neq \mathbf{0}_{nd}$, then $\sup_{\nu, \eta} L(x_v, \lambda_v, \nu, \eta) = +\infty$. Otherwise, $\sup_{\nu, \eta} L(x_v, \lambda_v, \nu, \eta) = F(x_v, \lambda_v)$. That is,

$$\sup_{\nu, \eta} L(x_v, \lambda_v, \nu, \eta) = \begin{cases} F(x_v, \lambda_v), & \text{if } L \lambda_v = \mathbf{0}_n \text{ and} \\ & (L \otimes I_d) x_v = \mathbf{0}_{nd}, \\ +\infty, & \text{otherwise.} \end{cases}$$

By the definition of L and the assumption that there exists a finite optimizer of (13) and consequently that of (15), we get

$$\inf_{x_v, \lambda_v \geq \mathbf{0}_n} \sup_{\nu, \eta} L(x_v, \lambda_v, \nu, \eta) = F(\mathbf{1}_n \otimes x^*, \lambda^* \mathbf{1}_n), \quad (19)$$

where (x^*, λ^*) recall is an optimizer of (13). Now pick $(\bar{\nu}, \bar{\eta})$ such that $L \bar{\nu} = \mathbf{0}_n$ and $(L \otimes I_d) \bar{\eta} = \mathbf{0}_{nd}$. Then, one can show that $\inf_{x_v, \lambda_v \geq \mathbf{0}_n} L(x_v, \lambda_v, \bar{\nu}, \bar{\eta}) = F(\mathbf{1}_n \otimes x^*, \lambda^* \mathbf{1}_n)$. Thus, we deduce that

$$\sup_{\nu, \eta} \inf_{x_v, \lambda_v \geq \mathbf{0}_n} L(x_v, \lambda_v, \nu, \eta) \geq F(\mathbf{1}_n \otimes x^*, \lambda^* \mathbf{1}_n). \quad (20)$$

Note that the following inequality holds (cf. [19, Lemma 36.1]),

$$\inf_{x_v, \lambda_v \geq \mathbf{0}_n} \sup_{\nu, \eta} L(x_v, \lambda_v, \nu, \eta) \geq \sup_{\nu, \eta} \inf_{x_v, \lambda_v \geq \mathbf{0}_n} L(x_v, \lambda_v, \nu, \eta).$$

Using the above inequality along with (19) and (20), we conclude that (18) holds. One can verify that the point $(\mathbf{1}_n \otimes x^*, \lambda^* \mathbf{1}_n, \bar{\nu}, \bar{\eta})$ with $L \bar{\nu} = \mathbf{0}_n$ and $(L \otimes I_d) \bar{\eta} = \mathbf{0}_{nd}$ is a saddle point of L over the domain $(\mathbb{R}^{nd} \times \mathbb{R}_{\geq 0}^n) \times (\mathbb{R}^n \times \mathbb{R}^{nd})$. Hence, the set of saddle points is nonempty.

We proceed to establish (i) and (ii). Assume $(\bar{x}_v, \bar{\lambda}_v, \bar{\nu}, \bar{\eta})$ to be a saddle point of L over the considered domain. To arrive at a contradiction, assume that $(\bar{x}_v, \bar{\lambda}_v)$ is not an optimizer of (13). Recall that there exists a finite optimizer (x^*, λ^*) of (13). This means that

$$F(\mathbf{1}_n \otimes x^*, \lambda^* \mathbf{1}_n) < F(\bar{x}_v, \bar{\lambda}_v). \quad (21)$$

Note that $(L \otimes I_d) \bar{x}_v = \mathbf{0}_{nd}$ and $L \bar{\lambda}_v = \mathbf{0}_n$ because otherwise $\sup_{\nu, \eta} L(\bar{x}_v, \bar{\lambda}_v, \nu, \eta) = +\infty$, contradicting the fact that $(\bar{x}_v, \bar{\lambda}_v, \bar{\nu}, \bar{\eta})$ is a saddle point. Therefore, using this property of $(\bar{x}_v, \bar{\lambda}_v)$ in (21), we obtain

$$L(\mathbf{1}_n \otimes x^*, \lambda^* \mathbf{1}_n, \bar{\nu}, \bar{\eta}) < L(\bar{x}_v, \bar{\lambda}_v, \bar{\nu}, \bar{\eta}).$$

This is a contradiction as $(\bar{x}_v, \bar{\lambda}_v, \bar{\nu}, \bar{\eta})$ is a saddle point. Finally, assume $(\bar{x}_v, \bar{\lambda}_v)$ to be an optimizer of (13). Again, pick $(\bar{\nu}, \bar{\eta})$ such that $L \bar{\nu} = \mathbf{0}_n$ and $(L \otimes I_d) \bar{\eta} = \mathbf{0}_{nd}$. One can show that $(\bar{x}_v, \bar{\lambda}_v, \bar{\nu}, \bar{\eta})$ is a saddle point of L over the specified domain, completing the proof. ■

Owing to the above result, one could potentially write a saddle-point dynamics for the Lagrangian L as a distributed algorithm to find the optimizers. However, without strict or strong convexity assumptions on the objective function, the resulting dynamics is in general not guaranteed to converge, see e.g., [27]. To overcome this hurdle, we augment the Lagrangian with quadratic terms in the primal variables. Let the augmented Lagrangian $L_{\text{aug}} : \mathbb{R}^{nd} \times \mathbb{R}_{\geq 0}^n \times \mathbb{R}^n \times \mathbb{R}^{nd} \rightarrow \mathbb{R}$ be

$$L_{\text{aug}}(x_v, \lambda_v, \nu, \eta) := L(x_v, \lambda_v, \nu, \eta) + \frac{1}{2} x_v^\top (L \otimes I_d) x_v + \frac{1}{2} \lambda_v^\top L \lambda_v.$$

Note that L_{aug} is also convex-concave in $((x_v, \lambda_v), (\nu, \eta))$ on the domain $\lambda_v \geq \mathbf{0}_n$. The next result guarantees that this augmentation step does not change the saddle points.

Lemma V.3. (Saddle points of L and L_{aug} are the same): *A point $(x_v^*, \lambda_v^*, \nu^*, \eta^*)$ is a saddle point of L over $(\mathbb{R}^{nd} \times \mathbb{R}_{\geq 0}^n) \times (\mathbb{R}^n \times \mathbb{R}^{nd})$ if and only if it is a saddle point of L_{aug} over the same domain.*

The proof follows by using the convexity property of the objective function in [28, Theorem 1.1]. The above result implies that finding the saddle points of L_{aug} would take us to the primal-dual optimizers of (15). However, a final roadblock remaining is writing a gradient-based dynamics for L_{aug} , given that this function involves a set of maximizations in its definition and so the gradient of L_{aug} with respect to x_v is undefined for $\lambda_v = 0$. Thus, our next task is to get rid of these internal optimization routines and identify a function for which the saddle-point dynamics is well defined over the feasible domain. Note that

$$L_{\text{aug}}(x_v, \lambda_v, \nu, \eta) = \max_{\{\xi^k\}} \tilde{L}_{\text{aug}}(x_v, \lambda_v, \nu, \eta, \{\xi^k\}), \quad (22)$$

where

$$\begin{aligned} \tilde{L}_{\text{aug}}(x_v, \lambda_v, \nu, \eta, \{\xi^k\}) &:= h(\lambda_v) + \frac{1}{N} \sum_{k=1}^N g_k(x^{v_k}, \lambda^{v_k}, \xi^k) \\ &+ \nu^\top L \lambda_v + \eta^\top (L \otimes I_d) x_v + \frac{1}{2} x_v^\top (L \otimes I_d) x_v + \frac{1}{2} \lambda_v^\top L \lambda_v. \end{aligned} \quad (23)$$

The following result shows that, under appropriate conditions, \tilde{L}_{aug} is the function we are looking for.

Proposition V.4. (Saddle points of \tilde{L}_{aug} and correspondence with optimizers of (15)): *Let $\mathcal{C} \subset \mathbb{R}^{nd} \times \mathbb{R}_{\geq 0}^n$ with $\text{int}(\mathcal{C}) \neq \emptyset$ be a closed, convex set such that*

- (i) *the saddle points of L_{aug} over the domain $(\mathbb{R}^{nd} \times \mathbb{R}_{\geq 0}^n) \times (\mathbb{R}^n \times \mathbb{R}^{nd})$ are contained in the set $\mathcal{C} \times (\mathbb{R}^n \times \mathbb{R}^{nd})$;*
- (ii) *\tilde{L}_{aug} is convex-concave on $\mathcal{C} \times (\mathbb{R}^n \times \mathbb{R}^{nd} \times \mathbb{R}^{mN})$;*
- (iii) *for any (ν, η) ,*

$$\begin{aligned} &\min_{(x_v, \lambda_v) \in \mathcal{C}} \max_{\{\xi^k\}} \tilde{L}_{\text{aug}}(x_v, \lambda_v, \nu, \eta, \{\xi^k\}) \\ &= \max_{\{\xi^k\}} \min_{(x_v, \lambda_v) \in \mathcal{C}} \tilde{L}_{\text{aug}}(x_v, \lambda_v, \nu, \eta, \{\xi^k\}). \end{aligned} \quad (24)$$

Then, the following holds

- (i) *The set of saddle points of \tilde{L}_{aug} over the domain $\mathcal{C} \times (\mathbb{R}^n \times \mathbb{R}^{nd} \times \mathbb{R}^{mN})$ is nonempty, convex, and closed.*
- (ii) *If $(\bar{x}_v, \bar{\lambda}_v, \bar{\nu}, \bar{\eta}, \{(\bar{\xi})^k\})$ is a saddle point of \tilde{L}_{aug} over $\mathcal{C} \times (\mathbb{R}^n \times \mathbb{R}^{nd} \times \mathbb{R}^{mN})$, then $(\bar{x}_v, \bar{\lambda}_v)$ is a optimizer of (15).*
- (iii) *If $(\bar{x}_v, \bar{\lambda}_v) \in \mathcal{C}$ is an optimizer of (15), then there exists $(\bar{\nu}, \bar{\eta}, \{(\bar{\xi})^k\})$ such that $(\bar{x}_v, \bar{\lambda}_v, \bar{\nu}, \bar{\eta}, \{(\bar{\xi})^k\})$ is a saddle point of \tilde{L}_{aug} over $\mathcal{C} \times (\mathbb{R}^n \times \mathbb{R}^{nd} \times \mathbb{R}^{mN})$.*

Proof: If saddle points of L_{aug} belong to $\mathcal{C} \times (\mathbb{R}^n \times \mathbb{R}^{nd})$, then according to [19, Lemma 36.2], we have

$$\begin{aligned} &\min_{(x_v, \lambda_v) \in \mathcal{C}} \max_{\nu, \eta} L_{\text{aug}}(x_v, \lambda_v, \nu, \eta) \\ &= \max_{\nu, \eta} \min_{(x_v, \lambda_v) \in \mathcal{C}} L_{\text{aug}}(x_v, \lambda_v, \nu, \eta). \end{aligned}$$

Using the definition (22) of L_{aug} in the above equality, we get

$$\begin{aligned} &\min_{(x_v, \lambda_v) \in \mathcal{C}} \max_{\nu, \eta} \max_{\{\xi^k\}} \tilde{L}_{\text{aug}}(x_v, \lambda_v, \nu, \eta, \{\xi^k\}) \\ &= \max_{\nu, \eta} \min_{(x_v, \lambda_v) \in \mathcal{C}} \max_{\{\xi^k\}} \tilde{L}_{\text{aug}}(x_v, \lambda_v, \nu, \eta, \{\xi^k\}). \end{aligned} \quad (25)$$

Using (24) on the right-hand side of the above expression gives

$$\min_{(x_v, \lambda_v) \in \mathcal{C}} \max_{\nu, \eta, \{\xi^k\}} \tilde{L}_{\text{aug}}(x_v, \lambda_v, \nu, \eta, \{\xi^k\})$$

$$= \max_{\nu, \eta, \{\xi^k\}} \min_{(x_v, \lambda_v) \in \mathcal{C}} \tilde{L}_{\text{aug}}(x_v, \lambda_v, \nu, \eta, \{\xi^k\}).$$

From the above equality and the fact that \tilde{L}_{aug} is convex-concave and finite-valued, we conclude from [19, Lemma 36.2] that the set of saddle points of \tilde{L}_{aug} over the domain $\mathcal{C} \times (\mathbb{R}^n \times \mathbb{R}^{nd} \times \mathbb{R}^{mN})$ is nonempty. Further, this set is closed and convex again due to convexity-concavity of \tilde{L}_{aug} . Finally, parts (ii) and (iii) follow from combining Lemmas V.2 and V.3 with the following two facts. First, from (25), if $(\bar{x}_v, \bar{\lambda}_v, \bar{\nu}, \bar{\eta}, \{(\bar{\xi})^k\})$ is a saddle point of \tilde{L}_{aug} , then $(\bar{x}_v, \bar{\lambda}_v, \bar{\nu}, \bar{\eta})$ is a saddle point of L_{aug} . Second, if $(\bar{x}_v, \bar{\lambda}_v, \bar{\nu}, \bar{\eta})$ is a saddle point of L_{aug} , then there exists $\{(\bar{\xi})^k\}$, which is the maximizer of $\{\xi^k\} \rightarrow \tilde{L}_{\text{aug}}(\bar{x}_v, \bar{\lambda}_v, \bar{\nu}, \bar{\eta}, \{\xi^k\})$, such that $(\bar{x}_v, \bar{\lambda}_v, \bar{\nu}, \bar{\eta}, \{(\bar{\xi})^k\})$ is a saddle point of \tilde{L}_{aug} , completing the proof. ■

Section VII describes classes of objective functions for which the hypotheses of Proposition V.4 are met. We have introduced in Proposition V.4 the set \mathcal{C} to increase the level of generality in preparation for the exposition of our algorithm that follows next. Specifically, since f is not necessarily convex-concave, the function \tilde{L}_{aug} might not be convex-concave over the entire domain $(\mathbb{R}^{nd} \times \mathbb{R}_{\geq 0}^n) \times (\mathbb{R}^n \times \mathbb{R}^{nd} \times \mathbb{R}^{mN})$. For such cases, one can restrict the attention to the set $\mathcal{C} \times (\mathbb{R}^n \times \mathbb{R}^{nd} \times \mathbb{R}^{mN})$ provided the hypotheses of the above result are satisfied. As we show later, when the objective function f is convex-concave, one can employ the set $\mathcal{C} = \mathbb{R}^{nd} \times \mathbb{R}_{\geq 0}^n$.

VI. DISTRIBUTED ALGORITHM DESIGN AND CONVERGENCE ANALYSIS

Here we design and analyze our distributed algorithm to find the solutions of the optimization problem (13). Given the results of Section V, and specifically Proposition V.4, our algorithm seeks to find the saddle points of \tilde{L}_{aug} over the domain $\mathcal{C} \times (\mathbb{R}^n \times \mathbb{R}^{nd} \times \mathbb{R}^{mN})$. The dynamics consists of (projected) gradient-descent of \tilde{L}_{aug} in the convex variables and gradient-ascent in the concave ones. This is popularly termed as the saddle-point or the primal-dual dynamics [27], [29].

Given a closed, convex set $\mathcal{C} \subset \mathbb{R}^{nd} \times \mathbb{R}_{\geq 0}^n$, the saddle-point dynamics for \tilde{L}_{aug} is

$$\left[\begin{array}{c} \frac{dx_v}{dt} \\ \frac{d\lambda_v}{dt} \end{array} \right] = \Pi_{\mathcal{C}} \left((x_v, \lambda_v), \left[\begin{array}{c} -\nabla_{x_v} \tilde{L}_{\text{aug}}(x_v, \lambda_v, \nu, \eta, \{\xi^k\}) \\ -\nabla_{\lambda_v} \tilde{L}_{\text{aug}}(x_v, \lambda_v, \nu, \eta, \{\xi^k\}) \end{array} \right] \right), \quad (26a)$$

$$\frac{d\nu}{dt} = \nabla_{\nu} \tilde{L}_{\text{aug}}(x_v, \lambda_v, \nu, \eta, \{\xi^k\}), \quad (26b)$$

$$\frac{d\eta}{dt} = \nabla_{\eta} \tilde{L}_{\text{aug}}(x_v, \lambda_v, \nu, \eta, \{\xi^k\}), \quad (26c)$$

$$\frac{d\xi^k}{dt} = \nabla_{\xi^k} \tilde{L}_{\text{aug}}(x_v, \lambda_v, \nu, \eta, \{\xi^k\}), \quad \forall k \in [N]. \quad (26d)$$

For convenience, denote (26) by the vector field $X_{\text{sp}}: \mathbb{R}^{nd} \times \mathbb{R}_{\geq 0}^n \times \mathbb{R}^{nd+n+mN} \rightarrow \mathbb{R}^{nd} \times \mathbb{R}_{\geq 0}^n \times \mathbb{R}^{nd+n+mN}$. In this notation, the first, second, and third components correspond to the dynamics of x_v , λ_v , and $(\nu, \eta, \{\xi^k\})$, respectively.

Remark VI.1. (Distributed implementation of X_{sp}): Here we discuss the distributed character of the dynamics (26). For

this, we rely on the set \mathcal{C} being decomposable into constraints on individual agent's decision variables, i.e., $\mathcal{C} := \prod_{i=1}^n \mathcal{C}_i$ with $\mathcal{C}_i \subset \mathbb{R}_{\geq 0} \times \mathbb{R}^d$. This allows agents to perform the projection in (26a) in a distributed way (we show later that the set \mathcal{C} enjoys this structure for a broad class of objective functions f). Denote the components of the dual variables η and ν by $\eta = (\eta^1; \eta^2; \dots; \eta^n)$ and $\nu = (\nu^1; \nu^2; \dots; \nu^n)$, so that agent $i \in [n]$ maintains $\eta^i \in \mathbb{R}^d$ and $\nu^i \in \mathbb{R}$. Further, let $\mathcal{K}_i \subset [N]$ be the set of indices representing the samples held by i ($k \in \mathcal{K}_i$ if and only if $\hat{\xi}^k \in \hat{\Xi}_i$). For implementing X_{sp} , we assume that each agent i maintains and updates the variables $(x^i, \lambda^i, \nu^i, \eta^i, \{\xi^k\}_{k \in \mathcal{K}_i})$. The collection of these variables for all $i \in [n]$ forms $(x_v, \lambda_v, \nu, \eta, \{\xi^k\})$. From (26), the dynamics of variables maintained by i is

$$\begin{aligned} \left(\frac{dx^i}{dt}; \frac{d\lambda^i}{dt} \right) &= \Pi_{\mathcal{C}_i} \left(-\frac{1}{N} \sum_{k \in \mathcal{K}_i} \nabla_x g_k(x^i, \lambda^i, \xi^k) \right. \\ &\quad \left. - \sum_{j \in \mathcal{N}_i} a_{ij} ((\eta^i - \eta^j) + (x^i - x^j)); \right. \\ &\quad \left. - \frac{\epsilon_N^2(\beta)}{n} - \frac{1}{N} \sum_{k \in \mathcal{K}_i} \nabla_{\lambda} g_k(x^i, \lambda^i, \xi^k) \right. \\ &\quad \left. - \sum_{j \in \mathcal{N}_i} a_{ij} ((\nu^i - \nu^j) + (\lambda^i - \lambda^j)) \right), \\ \frac{d\nu^i}{dt} &= \sum_{j \in \mathcal{N}_i} a_{ij} (\lambda^i - \lambda^j), \\ \frac{d\eta^i}{dt} &= \sum_{j \in \mathcal{N}_i} a_{ij} (x^i - x^j), \\ \frac{d\xi^k}{dt} &= \frac{1}{N} \nabla_{\xi} g_k(x^i, \lambda^i, \xi^k), \quad \forall k \in \mathcal{K}_i. \end{aligned}$$

Observe that the right-hand side of the above dynamics is computable by agent i using the variables that it maintains and information collected from its neighbors. Hence, X_{sp} can be implemented in a distributed manner. Note that the number of variables in the set $\{\xi^k\}$, grows with the size of the data, whereas the size of all other variables is independent of the number of samples. Further, for any agent i , $\{\xi^k\}_{k \in \mathcal{K}_i}$ can be interpreted as its internal state that is not communicated to its neighbors. •

The following result establishes the convergence of the dynamics X_{sp} to the saddle points of \tilde{L}_{aug} . In our previous work [27], we have extensively analyzed the convergence properties of saddle-point dynamics associated to convex-concave functions. However, those results do not apply directly to infer convergence for X_{sp} because projection operators are involved in our algorithm design and the function \tilde{L}_{aug} is not linear in the concave variable $\{\xi^k\}$ and not necessarily strictly convex in (x_v, λ_v) . Nonetheless, we borrow much insight from our previous analysis to prove the following result.

Theorem VI.2. (Convergence of trajectories of X_{sp} to the optimizers of (15)): *Suppose the hypotheses of Proposition V.4 hold. Assume further that there exists a saddle point $(x_v^*, \lambda_v^*, \nu^*, \eta^*, \{(\xi^k)^*\})$ of \tilde{L}_{aug} satisfying $(x_v^*, \lambda_v^*) \in \text{int}(\mathcal{C})$. Then, the trajectories of (26) starting in $\mathcal{C} \times \mathbb{R}^n \times \mathbb{R}^{nd} \times \mathbb{R}^{mN}$ remain in this set and converge asymptotically to a saddle*

point of \tilde{L}_{aug} . As a consequence, the (x_v, λ_v) component of the trajectory converges to an optimizer of (15).

Proof: We understand the trajectories of (26) in the Caratheodory sense, cf. Section II-4. Note that by definition of the projection operator, any solution $t \mapsto (x_v(t), \lambda_v(t), \nu(t), \eta(t), \{\xi^k(t)\})$ of (26) starting with $(x_v(0), \lambda_v(0)) \in \mathcal{C}$ satisfies $(x_v(t), \lambda_v(t)) \in \mathcal{C}$ for all $t \geq 0$.

LaSalle function. Let $(x_v^*, \lambda_v^*, \nu^*, \eta^*, \{(\xi^k)^*\})$ be the equilibrium point of \tilde{L}_{aug} satisfying $(x_v^*, \lambda_v^*) \in \text{int}(\mathcal{C})$. Using the definition of equilibrium point in (26b) and (26c), we get

$$(\mathbf{L} \otimes \mathbf{I}_d)x_v^* = \mathbf{0}_{nd} \text{ and } \mathbf{L}\lambda_v^* = \mathbf{0}_n. \quad (27)$$

Consider the function $V : \mathcal{C} \times \mathbb{R}^{nd+n+Nm} \rightarrow \mathbb{R}_{\geq 0}$,

$$V(x_v, \lambda_v, \zeta) := \frac{1}{2} (\|x_v - x_v^*\|^2 + \|\lambda_v - \lambda_v^*\|^2 + \|\zeta - \zeta^*\|^2),$$

where, for convenience, we use $\zeta := (\nu, \eta, \{\xi^k\})$ and, likewise, $\zeta^* := (\nu^*, \eta^*, \{(\xi^k)^*\})$. The Lie derivative of V along the dynamics (26) can be written as

$$\mathcal{L}_{(26)}V(x_v, \lambda_v, \zeta) = M_1 + M_2,$$

where

$$\begin{aligned} M_1 &= -(x_v - x_v^*)^\top \nabla_{x_v} \tilde{L}_{\text{aug}}(x_v, \lambda_v, \zeta) \\ &\quad - (\lambda_v - \lambda_v^*)^\top \nabla_{\lambda_v} \tilde{L}_{\text{aug}}(x_v, \lambda_v, \zeta) \\ &\quad + (\zeta - \zeta^*)^\top \nabla_{\zeta} \tilde{L}_{\text{aug}}(x_v, \lambda_v, \zeta), \\ M_2 &= -(x_v - x_v^*)^\top \varphi_{x_v} - (\lambda_v - \lambda_v^*)^\top \varphi_{\lambda_v}, \end{aligned}$$

and $(\varphi_{x_v}, \varphi_{\lambda_v})$ is an element of the normal cone $N_{\mathcal{C}}(x_v, \lambda_v)$ (cf. Section II-3). This representation allows us to analyze the projection in the (x_v, λ_v) -dynamics separately. By definition of the normal cone, we get $M_2 \leq 0$. Recall that \tilde{L}_{aug} is convex-concave in $((x_v, \lambda_v), \zeta)$. Thus, using the first-order convexity and concavity condition for these maps and following the computation as in [29, Proof of Lemma 4.1], we get

$$\begin{aligned} M_1 &\leq \tilde{L}_{\text{aug}}(x_v^*, \lambda_v^*, \zeta) - \tilde{L}_{\text{aug}}(x_v^*, \lambda_v^*, \zeta^*) \\ &\quad + \tilde{L}_{\text{aug}}(x_v^*, \lambda_v^*, \zeta^*) - \tilde{L}_{\text{aug}}(x_v, \lambda_v, \zeta^*). \end{aligned}$$

From the definition of saddle point, the sum of the first two terms of the right-hand side are nonpositive and so is the sum of the last two. Thus, $M_1 \leq 0$. Therefore, we conclude

$$\mathcal{L}_{(26)}V(x_v, \lambda_v, \zeta) \leq 0. \quad (28)$$

Application of LaSalle invariance principle. Using the property (28), we deduce two facts. First, given $\delta \geq 0$, any trajectory of (26) starting in $\mathcal{S}_\delta := V^{-1}(\leq \delta) \cap (\mathcal{C} \times \mathbb{R}^{n+nd+mN})$ remains in \mathcal{S}_δ at all times. In particular, every equilibrium point is stable under the dynamics. Second, the omega-limit set of each trajectory of (26) starting in \mathcal{S}_δ is invariant under the dynamics. Thus, from the invariance principle for discontinuous dynamical systems, cf. Proposition II.2, any solution of (26) converges to the largest invariant set

$$\mathcal{M} \subset \{(x_v, \lambda_v, \zeta) \mid \mathcal{L}_{(26)}V(x_v, \lambda_v, \zeta) = 0, (x_v, \lambda_v) \in \mathcal{C}\}.$$

Properties of the largest invariant set. Let $(x_v, \lambda_v, \zeta) \in \mathcal{M}$. Then, from $\mathcal{L}_{(26)}V(x_v, \lambda_v, \zeta) = 0$ and our bounding above of M_1 , we get

$$\tilde{L}_{\text{aug}}(x_v^*, \lambda_v^*, \zeta) \stackrel{(a)}{=} \tilde{L}_{\text{aug}}(x_v^*, \lambda_v^*, \zeta^*) \stackrel{(b)}{=} \tilde{L}_{\text{aug}}(x_v, \lambda_v, \zeta^*). \quad (29)$$

Expanding the equality (a) and using (27), we obtain

$$\begin{aligned} & \sum_{k=1}^N g_k((x_v^*)^{v_k}, (\lambda_v^*)^{v_k}, \xi^k) \\ &= \sum_{k=1}^N g_k((x_v^*)^{v_k}, (\lambda_v^*)^{v_k}, (\xi^*)^k). \end{aligned} \quad (30)$$

Since $(x_v^*, \lambda_v^*) \in \text{int}(\mathcal{C})$ and \mathcal{C} is a closed subset of $\mathbb{R}^{nd} \times \mathbb{R}_{\geq 0}^n$, we have $\lambda_v^* > \mathbf{0}_n$. Further, from the saddle-point property, $\{(\xi^*)^k\}$ maximizes the function $\{\xi^k\} \mapsto \sum_{k=1}^N g_k((x^*)^{v_k}, (\lambda^*)^{v_k}, \xi^k)$. This map is strongly concave as $\lambda_v^* > \mathbf{0}_n$ (which can be verified by checking that the Hessian with respect to $\{\xi^k\}$ is negative definite). Therefore, (30) yields

$$\xi^k = (\xi^*)^k, \text{ for all } k \in [N].$$

Expanding the equality (b) in (29) and using (27), we get

$$\begin{aligned} & h(\lambda_v^*) + \frac{1}{N} \sum_{k=1}^N g_k((x^*)^{v_k}, (\lambda^*)^{v_k}, (\xi^*)^k) = h(\lambda_v) \\ & + \frac{1}{N} \sum_{k=1}^N g_k(x^{v_k}, \lambda^{v_k}, (\xi^*)^k) + (\nu^*)^\top \mathbf{L} \lambda_v \\ & + (\eta^*)^\top (\mathbf{L} \otimes \mathbf{I}_d) x_v + \frac{1}{2} x_v^\top (\mathbf{L} \otimes \mathbf{I}_d) x_v + \frac{1}{2} \lambda_v^\top \mathbf{L} \lambda_v. \end{aligned} \quad (31)$$

For ease of notation, let $y_v := (x_v; \lambda_v)$, $y_v^* := (x_v^*; \lambda_v^*)$, and

$$F(y_v) := h(\lambda_v) + \frac{1}{N} \sum_{k=1}^N g_k(x^{v_k}, \lambda^{v_k}, (\xi^*)^k).$$

Then, the expression (31) can be written as

$$\begin{aligned} F(y_v^*) &= F(y_v) + (\nu^*)^\top \mathbf{L} \lambda_v + (\eta^*)^\top (\mathbf{L} \otimes \mathbf{I}_d) x_v \\ & + \frac{1}{2} y_v^\top (\mathbf{L} \otimes \mathbf{I}_{d+1}) y_v. \end{aligned} \quad (32)$$

From the definition of saddle point, (x_v^*, λ_v^*) minimizes the function $(x_v, \lambda_v) \mapsto \tilde{L}_{\text{aug}}(x_v, \lambda_v, \zeta^*)$ over the domain \mathcal{C} . Moreover, by assumption (x_v^*, λ_v^*) lies in the interior of \mathcal{C} . Thus,

$$\nabla_{x_v} \tilde{L}_{\text{aug}}(x_v^*, \lambda_v^*, \zeta^*) = \mathbf{0}_{nd}, \quad (33a)$$

$$\nabla_{\lambda_v} \tilde{L}_{\text{aug}}(x_v^*, \lambda_v^*, \zeta^*) = \mathbf{0}_n. \quad (33b)$$

The first of the above equalities yield

$$(\mathbf{L} \otimes \mathbf{I}_d) \eta^* = -\nabla_{x_v} F(y_v^*).$$

Plugging this equality in (32) and rearranging terms gives

$$\begin{aligned} & \frac{1}{2} y_v^\top (\mathbf{L} \otimes \mathbf{I}_{d+1}) y_v = F(y_v^*) - F(y_v) \\ & - (\nu^*)^\top \mathbf{L} \lambda_v + x_v^\top \nabla_{x_v} F(y_v^*). \end{aligned} \quad (34)$$

Note that $(x_v^*)^\top \nabla_{x_v} F(y_v^*) = (x_v^*)^\top (\nabla_{x_v} F(y_v^*) + (\mathbf{L} \otimes \mathbf{I}_d) \eta^* + (\mathbf{L} \otimes \mathbf{I}_d) x_v^*) = (x_v^*)^\top \nabla_{x_v} \tilde{L}_{\text{aug}}(x_v^*, \lambda_v^*, \zeta^*)$, where we have used (27). This in turn equals 0 because of (33a). Thus, we can rewrite (34) as

$$\begin{aligned} & \frac{1}{2} y_v^\top (\mathbf{L} \otimes \mathbf{I}_{d+1}) y_v = F(y_v^*) - F(y_v) \\ & - (\nu^*)^\top \mathbf{L} \lambda_v + (x_v - x_v^*)^\top \nabla_{x_v} F(y_v^*) \end{aligned} \quad (35)$$

Expanding (33b) gives

$$\nabla_{\lambda_v} F(y_v^*) + \mathbf{L} \nu^* + \frac{1}{2} \mathbf{L} \lambda_v^* = 0. \quad (36)$$

Pre-multiplying the above equation with $(\lambda_v^*)^\top$ and using (27), we get $(\lambda_v^*)^\top \nabla_{\lambda_v} F(y_v^*) = 0$ and we can further rewrite (35) as

$$\begin{aligned} & \frac{1}{2} y_v^\top (\mathbf{L} \otimes \mathbf{I}_{d+1}) y_v = F(y_v^*) - F(y_v) - (\nu^*)^\top \mathbf{L} \lambda_v \\ & + (x_v - x_v^*)^\top \nabla_{x_v} F(y_v^*) - (\lambda_v^*)^\top \nabla_{\lambda_v} F(y_v^*) \end{aligned} \quad (37)$$

Using (27) in (36) yields $\nabla_{\lambda_v} F(y_v^*) = -\mathbf{L} \nu^*$. That is, $\lambda_v^\top \nabla_{\lambda_v} F(y_v^*) = -\lambda_v^\top \mathbf{L} \nu^*$ which then replaced in (37) gives

$$\frac{1}{2} y_v^\top (\mathbf{L} \otimes \mathbf{I}_{d+1}) y_v = F(y_v^*) - F(y_v) + (y_v - y_v^*)^\top \nabla_{y_v} F(y_v^*).$$

The first-order convexity condition for F takes the form

$$F(y_v) \geq F(y_v^*) + (y_v - y_v^*)^\top \nabla_{y_v} F(y_v^*)$$

Using the previous two expressions, we obtain

$$y_v^\top (\mathbf{L} \otimes \mathbf{I}_{d+1}) y_v \leq 0.$$

This is only possible if this expression is zero because $\mathbf{L} \otimes \mathbf{I}_{d+1}$ is positive semidefinite. Equating it to zero, we get $x_v = \mathbf{1}_n \otimes x$ and $\lambda_v = \lambda \mathbf{1}_n$ for some (x, λ) and $(x_v, \lambda_v) \in \mathcal{C}$. Collecting our derivations so far, we have that if $(x_v, \lambda_v, \zeta) \in \mathcal{M}$, then

$$\xi^k = (\xi^*)^k, \forall k \in [N], \quad x_v = \mathbf{1}_n \otimes x, \quad (38a)$$

$$\lambda_v = \lambda \mathbf{1}_n, \quad (x_v, \lambda_v) \in \mathcal{C}. \quad (38b)$$

Identification of the largest invariant set. Consider a trajectory $t \mapsto (x_v(t), \lambda_v(t), \zeta(t))$ of (26) starting at $(x_v(0), \lambda_v(0), \zeta(0)) \in \mathcal{M}$ and remaining in \mathcal{M} at all times (recall that \mathcal{M} is invariant). Then, the trajectory must satisfy (38) for all $t \geq 0$, that is, there exists $t \mapsto (x(t), \lambda(t))$ such that

$$\xi^k(t) = (\xi^*)^k, \forall k \in [N], \quad x_v(t) = \mathbf{1}_n \otimes x(t), \quad (39a)$$

$$\lambda_v(t) = \lambda(t) \mathbf{1}_n, \quad (x_v(t), \lambda_v(t)) \in \mathcal{C}, \quad (39b)$$

for all $t \geq 0$. Plugging (39) in (26), we obtain that for all $t \geq 0$, along the considered trajectory, we have $\dot{v}(t) = \mathbf{0}_n$, $\dot{\eta}(t) = \mathbf{0}_{nd}$, and $\dot{\xi}(t) = \mathbf{0}_{mN}$. This implies that the considered trajectory satisfies the following for all $t \geq 0$,

$$\left[\frac{dx_v(t)}{dt} \right] = \Pi_{\mathcal{C}} \left((x_v(t), \lambda_v(t), \zeta(0)), \left[\begin{array}{c} -\nabla_{x_v} \tilde{L}_{\text{aug}}(x_v(t), \lambda_v(t), \zeta(0)) \\ -\nabla_{\lambda_v} \tilde{L}_{\text{aug}}(x_v(t), \lambda_v(t), \zeta(0)) \end{array} \right] \right)$$

which is a gradient descent dynamics of the convex function $(x_v, \lambda_v) \mapsto \tilde{L}_{\text{aug}}(x_v, \lambda_v, \zeta(0))$ projected over the set \mathcal{C} . Thus, either $t \mapsto \tilde{L}_{\text{aug}}(x_v(t), \lambda_v(t), \zeta(0))$ decreases at some t or the

right-hand side of the above dynamics is zero at all times. Note that for all $t \geq 0$,

$$\begin{aligned} \tilde{L}_{\text{aug}}(x_v(t), \lambda_v(t), \zeta(0)) &\stackrel{(a)}{=} \tilde{L}_{\text{aug}}(\mathbf{1}_n \otimes x(t), \lambda(t)\mathbf{1}_n, \zeta(0)) \\ &\stackrel{(b)}{=} h(\lambda(t)\mathbf{1}_n) + \frac{1}{N} \sum_{k=1}^N g_k(\mathbf{1}_n \otimes x(t), \lambda(t)\mathbf{1}_n, (\xi^*)^k) \\ &\stackrel{(c)}{=} \tilde{L}_{\text{aug}}(\mathbf{1}_n \otimes x(t), \lambda(t)\mathbf{1}_n, \zeta^*) \stackrel{(d)}{=} \tilde{L}_{\text{aug}}(x_v^*, \lambda_v^*, \zeta^*). \end{aligned}$$

In the above set of expressions, equalities (a), (b), and (c) follow from conditions (39) and the definition of \tilde{L}_{aug} . Equality (d) follows from (29), which holds from every point in \mathcal{M} . The above implies that $t \mapsto \tilde{L}_{\text{aug}}(x_v(t), \lambda_v(t), \zeta(0))$ is a constant map. As a consequence, we conclude that $(x_v(0), \lambda_v(0), \zeta(0))$ is an equilibrium point of (26). Therefore, we have proved that the set \mathcal{M} is entirely composed of the equilibrium points of the dynamics (26). Convergence to an equilibrium point in the set of saddle points for each trajectory follows from this and the fact that each equilibrium point is stable, cf. [30]. ■

Remark VI.3. (*Convergence of algorithm for nonsmooth objective functions*): Let f satisfy all assumptions outlined in Section IV except the differentiability and instead assume it is locally Lipschitz. This implies that the gradient of \tilde{L}_{aug} with respect to variables x_v and $\{\xi^k\}$ need not exist everywhere. However, the generalized gradients exist, see e.g., [21] for the definition. Therefore, one can replace gradients in (26a) and (26d) with the generalized counterparts and end up with a differential inclusion for the $\{\xi^k\}$ dynamics and a projected differential inclusion for the x_v dynamics. Although we do not explore it here, we believe that, using analysis tools of nonsmooth dynamical systems, see [21] and references therein, one can show that the trajectories of the resulting nonsmooth dynamical system retain the convergence properties of Theorem VI.2. A promising route to establish this is to follow the exposition of [31], which studies saddle-point dynamics for a general class of functions. •

Remark VI.4. (*Constrained stochastic optimization*): Certain constrained stochastic optimization problems can be cast in the form (12) and are therefore amenable to the distributed algorithmic solution techniques developed here. Given $\delta \in (0, 1)$ and a measurable map $F : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$, consider the following constrained stochastic optimization problem

$$\inf_{x \in \{\mathbb{R}^d \mid \mathbb{P}(F(x, \xi) \leq 0) \geq 1 - \delta\}} \mathbb{E}_{\mathbb{P}}[f(x, \xi)]. \quad (40)$$

The constraint is probabilistic in nature and so is commonly referred to as *chance constraint* [2]. One approach to solve this problem is to remove the constraint and add a convex function to the objective that penalizes its violation. Conditional value-at-risk (CVaR) is one such penalizing function. Formally, the CVaR of $\xi \mapsto F(x, \xi)$ at level δ is

$$\text{CVaR}_{\delta}(F(x, \xi)) := \inf_{\tau \in \mathbb{R}} \mathbb{E}_{\mathbb{P}} \left[\tau + \frac{1}{\delta} \max\{F(x, \xi) - \tau, 0\} \right].$$

Roughly speaking, this value represents the expectation of $\xi \mapsto F(x, \xi)$ over the set of ξ that has measure δ and that contain the highest values of this function. Note the fact [2, Chapter

6] that $\text{CVaR}_{\delta}(F(x, \xi)) \leq 0$ implies $\mathbb{P}(F(x, \xi) \leq 0) \geq 1 - \delta$. Thus, using CVaR, problem (40) can be approximated by

$$\inf_{x \in \mathbb{R}^d} \mathbb{E}_{\mathbb{P}}[f(x, \xi)] + \rho \text{CVaR}_{\delta}(F(x, \xi)),$$

where $\rho > 0$ determines the trade-off between the two goals: minimizing the objective and satisfying the constraint. By invoking the definition of CVaR, the above problem can be written compactly as

$$\inf_{x \in \mathbb{R}^d, \tau \in \mathbb{R}} \mathbb{E}_{\mathbb{P}} \left[f(x, \xi) + \rho \left(\tau + \frac{1}{\delta} \max\{F(x, \xi) - \tau, 0\} \right) \right].$$

This can be further recast as a stochastic optimization of the form (12). Therefore, under appropriate conditions on the function F , one can solve a chance-constrained problem in a distributed way under the data-driven optimization paradigm using the algorithm design introduced here. •

VII. OBJECTIVE FUNCTIONS THAT MEET THE ALGORITHM CONVERGENCE CRITERIA

In this section we report on two broad classes of objective functions f for which the hypotheses of Proposition V.4 hold. For both cases, we justify how the dynamics (26) serves as the distributed algorithm for solving (15).

A. Convex-concave functions

Here we focus on objective functions that are convex-concave in (x, ξ) . That is, in addition to $x \mapsto f(x, \xi)$ being convex for each $\xi \in \mathbb{R}^m$, the function $\xi \mapsto f(x, \xi)$ is concave for each $x \in \mathbb{R}^n$. We proceed to check the hypotheses of Theorem VI.2. To this end, let $\mathcal{C} = \mathbb{R}^{nd} \times \mathbb{R}_{\geq 0}^n$, which is closed, convex set with $\text{int}(\mathcal{C}) \neq \emptyset$. Note that \tilde{L}_{aug} is convex-concave on $\mathcal{C} \times (\mathbb{R}^n \times \mathbb{R}^{nd} \times \mathbb{R}^{mN})$ as f is convex-concave. The following result shows that (24) holds.

Lemma VII.1. (*Min-max operators can be interchanged for \tilde{L}_{aug}*): Let f be convex-concave in (x, ξ) . Then, for any $(\nu, \eta) \in \mathbb{R}^n \times \mathbb{R}^{nd}$, the following holds

$$\begin{aligned} \min_{x_v, \lambda_v \geq \mathbf{0}_n} \max_{\{\xi^k\}} \tilde{L}_{\text{aug}}(x_v, \lambda_v, \nu, \eta, \{\xi^k\}) \\ = \max_{\{\xi^k\}} \min_{x_v, \lambda_v \geq \mathbf{0}_n} \tilde{L}_{\text{aug}}(x_v, \lambda_v, \nu, \eta, \{\xi^k\}). \end{aligned} \quad (41)$$

Proof: Given any (ν, η) , denote the function $(x_v, \lambda_v, \{\xi^k\}) \mapsto \tilde{L}_{\text{aug}}(x_v, \lambda_v, \nu, \eta, \{\xi^k\})$ by $\tilde{L}_{\text{aug}}^{(\nu, \eta)}$. Since f is convex-concave, so is $\tilde{L}_{\text{aug}}^{(\nu, \eta)}$ in the variables $((x_v, \lambda_v), \{\xi^k\})$. We use Theorem II.1 to prove the result. To do so, let us extend $\tilde{L}_{\text{aug}}^{(\nu, \eta)}$ over the entire domain $(\mathbb{R}^{nd} \times \mathbb{R}^n) \times (\mathbb{R}^{mN})$ as

$$\bar{L}_{\text{aug}}^{(\nu, \eta)}(x_v, \lambda_v, \{\xi^k\}) = \begin{cases} \tilde{L}_{\text{aug}}^{(\nu, \eta)}(x_v, \lambda_v, \{\xi^k\}), & \text{if } \lambda_v \geq \mathbf{0}_n, \\ +\infty, & \text{otherwise.} \end{cases}$$

One can see that $\bar{L}_{\text{aug}}^{(\nu, \eta)}$ is closed, proper, and convex-concave (cf. Section II for definitions). Further, following [19, Theorem 36.3], the equality (41) holds if and only if the following holds

$$\min_{x_v, \lambda_v} \max_{\{\xi^k\}} \bar{L}_{\text{aug}}^{(\nu, \eta)}(x_v, \lambda_v, \{\xi^k\})$$

$$= \max_{\{\xi^k\}} \min_{x_v, \lambda_v} \bar{L}_{\text{aug}}^{(\nu, \eta)}(x_v, \lambda_v, \{\xi^k\}).$$

The rest of the proof establishes the above condition by checking the hypotheses of Theorem II.1 for $\bar{L}_{\text{aug}}^{(\nu, \eta)}$. For showing Theorem II.1(i), it is enough to identify $\{\bar{\xi}^k\} \in \mathbb{R}^{mN}$ for which the function $(x_v, \lambda_v) \mapsto \bar{L}_{\text{aug}}^{(\nu, \eta)}(x_v, \lambda_v, \{\bar{\xi}^k\})$ does not have a direction of recession. By the assumptions on f , for each $k \in [N]$, there exists $\bar{\xi}^k \in B_{\epsilon_N(\beta)\sqrt{N}/\sqrt{2n}}(\hat{\xi}^k)$ such that $\mathbf{1}_n$ and $-\mathbf{1}_n$ are not directions of recession for the function $x \mapsto f(x, \bar{\xi}^k)$. Picking these values, one has $\|\bar{\xi}^k - \hat{\xi}^k\|^2 \leq \epsilon_N^2(\beta)N/2n$ for all $k \in [N]$. Thus,

$$\begin{aligned} \bar{L}_{\text{aug}}^{(\nu, \eta)}(x_v, \lambda_v, \{\bar{\xi}^k\}) &= \frac{\epsilon_N^2(\beta)(z^t \lambda_v)}{n} + \frac{1}{N} \sum_{k=1}^N f(x^{v_k}, \bar{\xi}^k) \\ &+ \nu^\top L \lambda_v + \eta^\top (L \otimes I_d) x_v + \frac{1}{2} \lambda_v^\top (L \otimes I_d) x_v + \frac{1}{2} \lambda_v^\top L \lambda_v, \end{aligned}$$

where $z \in \mathbb{R}^n$ with $z_i > 0$ for all $i \in [n]$. One can show that the right-hand side of the above expression as a function of (x_v, λ_v) does not have a direction of recession, that is, Theorem II.1(i) holds. Next, we check Theorem II.1(ii). We show that there exists $(\bar{x}_v, \bar{\lambda}_v) \in \text{ri}(\mathbb{R}^{nd} \times \mathbb{R}_{\geq 0}^n)$ such that the function $\{\xi^k\} \mapsto -\bar{L}_{\text{aug}}^{(\nu, \eta)}(\bar{x}_v, \bar{\lambda}_v, \{\xi^k\})$ does not have a direction of recession. To this end, pick $\bar{x}_v = \mathbf{1}_{nd}$ and $\bar{\lambda}_v = \mathbf{1}_n$. Then,

$$\bar{L}_{\text{aug}}^{(\nu, \eta)}(\bar{x}_v, \bar{\lambda}_v, \{\xi^k\}) = \epsilon_N^2(\beta) + \frac{1}{N} \sum_{k=1}^N f(\mathbf{1}_d, \xi^k) - \|\xi^k - \hat{\xi}^k\|^2.$$

Recall that for any $x \in \mathbb{R}^d$, $\xi \mapsto f(x, \xi)$ is concave. Hence, we deduce from the above expression that $\bar{L}_{\text{aug}}^{(\nu, \eta)}(\bar{x}_v, \bar{\lambda}_v, \{\xi^k\}) \rightarrow -\infty$ as $\|\{\xi^k\}\| \rightarrow \infty$. Therefore, $\{\xi^k\} \mapsto -\bar{L}_{\text{aug}}^{(\nu, \eta)}(\bar{x}_v, \bar{\lambda}_v, \{\xi^k\})$ does have a direction of recession, completing the proof. ■

As a consequence of the above discussion, we conclude that the hypotheses of Proposition V.4 hold true for the considered class of objective functions, and we can state, invoking Theorem VI.2, the following convergence result.

Corollary VII.2. (Convergence of trajectories of X_{sp} for convex-concave f): Let f be convex-concave in (x, ξ) and $\mathcal{C} = \mathbb{R}^{nd} \times \mathbb{R}_{\geq 0}$. Assume further that there exists a saddle point $(x_v^*, \lambda_v^*, \nu^*, \eta^*, \{(\xi^k)^*\})$ of \tilde{L}_{aug} satisfying $\lambda_v^* > \mathbf{0}_n$. Then, the trajectories of (26) starting in $\mathcal{C} \times \mathbb{R}^n \times \mathbb{R}^{nd} \times \mathbb{R}^{mN}$ remain in this set and converge asymptotically to a saddle point of \tilde{L}_{aug} . As a consequence, the (x_v, λ_v) component of the trajectory converges to an optimizer of (15).

It is important to note that $\mathcal{C} = \prod_{i=1}^n (\mathbb{R}^d \times \mathbb{R}_{\geq 0})$ and thus the projection in (26a) can be executed by individual agents. Following Remark VI.1, the dynamics (26) is implementable in a distributed way.

B. Convex-convex function

Here we focus on objective functions for which both $x \mapsto f(x, \xi)$ and $\xi \mapsto f(x, \xi)$ are convex maps for all $x \in \mathbb{R}^d$ and $\xi \in \mathbb{R}^m$. Note that f need not be jointly convex in x and ξ . We further divide this classification into two.

1) *Quadratic function in ξ* : Assume additionally that the function f is of the form

$$f(x, \xi) := \xi^\top Q \xi + x^\top R \xi + \ell(x), \quad (42)$$

where $Q \in \mathbb{R}^{m \times m}$ is positive definite, $R \in \mathbb{R}^{d \times m}$, and ℓ is a continuously differentiable convex function. Our next result is useful in identifying a domain that contains the saddle points of L_{aug} over $(\mathbb{R}^{nd} \times \mathbb{R}_{\geq 0}^n) \times (\mathbb{R}^n \times \mathbb{R}^{nd})$.

Lemma VII.3. (Characterizing where the objective function of (15) is finite-valued): Assume f is of the form (42). Then, the function F defined in (16) is finite-valued if and only if $\lambda^i \geq \lambda_{\max}(Q)$ for all $i \in [n]$.

Proof: Assume there exists $\tilde{i} \in [n]$ such that $\lambda^{\tilde{i}} < \lambda_{\max}(Q)$. We wish to show that $F(x_v, \lambda_v) = +\infty$ in this case. For any k such that $\hat{\xi}^k \in \hat{\Xi}_{\tilde{i}}$, we have

$$\begin{aligned} g_k(x^{\tilde{i}}, \lambda^{\tilde{i}}, \xi) &= \xi^\top (Q - \lambda^{\tilde{i}} I_m) \xi + (x^{\tilde{i}})^\top R \xi + 2\lambda^{\tilde{i}} (\hat{\xi}^k)^\top \xi \\ &+ \ell(x^{\tilde{i}}) - \lambda^{\tilde{i}} \|\hat{\xi}^k\|^2. \end{aligned}$$

Let $w_{\max}(Q) \in \mathbb{R}^m$ be an eigenvector of Q corresponding to the eigenvalue $\lambda_{\max}(Q)$. Parametrizing $\xi = \alpha w_{\max}(Q)$, we obtain

$$\begin{aligned} g_k(x^{\tilde{i}}, \lambda^{\tilde{i}}, \alpha w_{\max}(Q)) &= \alpha^2 (\lambda_{\max}(Q) - \lambda^{\tilde{i}}) \|w_{\max}(Q)\|^2 \\ &+ \alpha \left((x^{\tilde{i}})^\top R + 2\lambda^{\tilde{i}} (\hat{\xi}^k)^\top \right) w_{\max}(Q) + \ell(x^{\tilde{i}}) - \lambda^{\tilde{i}} \|\hat{\xi}^k\|^2. \end{aligned}$$

Thus, we get $\max_{\alpha} g_k(x^{\tilde{i}}, \lambda^{\tilde{i}}, \alpha w_{\max}(Q)) = +\infty$ and so $\max_{\xi} g_k(x^{\tilde{i}}, \lambda^{\tilde{i}}, \xi) = +\infty$. Further note that for any i and k with $\hat{\xi}^k \in \hat{\Xi}_i$, $\max_{\xi} g_k(x^i, \lambda^i, \xi) > -\infty$. This implies that $\sum_{k=1}^N \max_{\xi} g_k(x^{v_k}, \lambda^{v_k}, \xi) = +\infty$ and so $F(x_v, \lambda_v) = +\infty$.

Now assume that $\lambda^i \geq \lambda_{\max}(Q)$ for all $i \in [n]$. Then, for each k , evaluating the Hessian of g_k with respect to ξ leads to the conclusion that g_k is a concave function of ξ . Therefore, $\max_{\xi} g_k(x^{v_k}, \lambda^{v_k}, \xi)$ is finite-valued for each k . Hence, $F(x_v, \lambda_v) < +\infty$, completing the proof. ■

The above result implies that the optimizers of (15) for objective functions of the form (42) belong to the domain

$$\mathcal{C} := \mathbb{R}^{nd} \times \{\lambda_v \in \mathbb{R}_{\geq 0}^n \mid \lambda_v \geq \lambda_{\max}(Q) \mathbf{1}_n\}. \quad (43)$$

Therefore, the saddle points of L_{aug} over the domain $(\mathbb{R}^{nd} \times \mathbb{R}_{\geq 0}^n) \times (\mathbb{R}^n \times \mathbb{R}^{nd})$ are contained in the set $\mathcal{C} \times (\mathbb{R}^n \times \mathbb{R}^{nd})$. Note that \mathcal{C} is closed, convex with a nonempty interior. Furthermore, following the proof of Lemma VII.3, one can show that \tilde{L}_{aug} is convex-concave on $\mathcal{C} \times (\mathbb{R}^n \times \mathbb{R}^{nd} \times \mathbb{R}^{mN})$. An easy way to validate this fact is by noting that the Hessian of \tilde{L}_{aug} with respect to the convex (concave) variables is positive (negative) semidefinite. Finally, repeating the proof of Lemma VII.1, we arrive at the equality (24). Using these facts in Theorem VI.2 yields the following result.

Corollary VII.4. (Convergence of trajectories of X_{sp} for quadratic f): Let f be of the form (42) and \mathcal{C} be given in (43). Assume further that there exists a saddle point $(x_v^*, \lambda_v^*, \nu^*, \eta^*, \{(\xi^k)^*\})$ of \tilde{L}_{aug} satisfying $\lambda_v^* > \lambda_{\max}(Q) \mathbf{1}_n$. Then, the trajectories of (26) starting in $\mathcal{C} \times \mathbb{R}^n \times \mathbb{R}^{nd} \times \mathbb{R}^{mN}$ remain in this set and converge asymptotically to a saddle point of \tilde{L}_{aug} . As a consequence, the (x_v, λ_v) component of the trajectory converges to an optimizer of (15).

Note that \mathcal{C} given in (43) can be written as $\mathcal{C} = \prod_{i=1}^n (\mathbb{R}^d \times \{\lambda \in \mathbb{R}_{\geq 0} \mid \lambda \geq \lambda_{\max}(Q)\})$. Thus, following Remark VI.1, the dynamics (26) for this case can be implemented in a distributed manner.

2) *Least-squares problem:* Let $d = m$ and assume additionally that the function f is of the form

$$f(x, \xi) := a(\xi_m - (\xi_{1:m-1}; 1)^\top x)^2, \quad (44)$$

where $a > 0$ and $\xi_{1:m-1}$ denotes the vector ξ without the last component ξ_m . Note that f corresponds to the objective function for a least-squares problem. Further, note that it cannot be written in the form (42), as can be seen from its equivalent expression

$$f(x, \xi) = a \left(\xi^\top (-x_{1:m-1}; 1) (-x_{1:m-1}; 1)^\top \xi - 2x_m (-x_{1:m-1}; 1)^\top \xi + x_m^2 \right).$$

Our first step is to characterize, similarly to Lemma VII.3, the set over which the objective function (16) takes finite values.

Lemma VII.5. (*Characterizing where the objective function of (15) is finite-valued:*) Assume f is of the form (44). Then, the function F defined in (16) is finite-valued if and only if $\lambda^i \geq a \|(x_{1:m-1}^i; 1)\|^2$ for all $i \in [n]$.

Proof: The proof mimics the steps of the proof of Lemma VII.3. Assume there exists $\tilde{i} \in [n]$ such that $\lambda^{\tilde{i}} < a \|(x_{1:m-1}^{\tilde{i}}; 1)\|^2$. For any k such that $\hat{\xi}^k \in \tilde{\Xi}_{\tilde{i}}$, we have

$$g_k(x^{\tilde{i}}, \lambda^{\tilde{i}}, \xi) = a(\xi_m - (\xi_{1:m-1}; 1)^\top x^{\tilde{i}})^2 - \lambda^{\tilde{i}} \|\xi - \hat{\xi}^k\|^2.$$

Parametrizing $\xi = \alpha(-x_{1:m-1}^{\tilde{i}}; 1)$, we obtain

$$g_k(x^{\tilde{i}}, \lambda^{\tilde{i}}, \alpha(-x_{1:m-1}^{\tilde{i}}; 1)) = a \left(\alpha \left(\| -x_{1:m-1}^{\tilde{i}}; 1 \|^2 - x_m^{\tilde{i}} \right)^2 - \lambda^{\tilde{i}} \|\alpha(-x_{1:m-1}^{\tilde{i}}; 1) - \hat{\xi}^k\|^2 \right).$$

This is a quadratic function in the parameter α and the coefficient of the second-order term is $\|(-x_{1:m-1}^{\tilde{i}}; 1)\|^2 (a \|(-x_{1:m-1}^{\tilde{i}}; 1)\|^2 - \lambda^{\tilde{i}})$. This coefficient is positive by the assumption stipulated above. Therefore, $\max_{\xi} g_k(x^{\tilde{i}}, \lambda^{\tilde{i}}, \xi) = +\infty$. Since $\max_{\xi} g_k(x^{v_k}, \lambda^{v_k}, \xi) > -\infty$ for all $k \in [N]$, we conclude that $F(x_v, \lambda_v) = +\infty$.

On the other hand if $\lambda^i \geq a \|(x_{1:m-1}^i; 1)\|^2$ for all $i \in [n]$, then g_k is a concave function of ξ for all $k \in [N]$. Thus, $\max_{\xi} g_k(x^{v_k}, \lambda^{v_k}, \xi)$ is finite-valued for each k and so is $F(x_v, \lambda_v) < +\infty$. ■

Guided by the above result, let

$$\mathcal{C} := \mathbb{R}^{nd} \times \{\lambda_v \in \mathbb{R}_{\geq 0} \mid \lambda^i \geq a \|(x_{1:m-1}^i; 1)\|^2\}. \quad (45)$$

As a consequence of Lemma VII.5, the optimizers of (15) belong to \mathcal{C} and so, the saddle points of L_{aug} over the domain $(\mathbb{R}^{nd} \times \mathbb{R}_{\geq 0}^n) \times (\mathbb{R}^n \times \mathbb{R}^{nd})$ are contained in the set $\mathcal{C} \times (\mathbb{R}^n \times \mathbb{R}^{nd})$. Further, \mathcal{C} is closed, convex with a nonempty interior and the function \tilde{L}_{aug} is convex-concave on $\mathcal{C} \times (\mathbb{R}^n \times \mathbb{R}^{nd} \times \mathbb{R}^{mN})$. Finally, one can show that (24) holds in this case. Using these facts in Theorem VI.2 yields the following result.

Corollary VII.6. (*Convergence of trajectories of X_{sp} for least squares problem:*) Let f be of the form (44) and \mathcal{C} be

given in (45). Assume further that there exists a saddle point $(x_v^*, \lambda_v^*, \nu^*, \eta^*, \{(\xi^k)^*\})$ of \tilde{L}_{aug} satisfying $(x_v^*, \lambda_v^*) \in \text{int}(\mathcal{C})$. Then, the trajectories of (26) starting in $\mathcal{C} \times \mathbb{R}^n \times \mathbb{R}^{nd} \times \mathbb{R}^{mN}$ remain in this set and converge asymptotically to a saddle point of \tilde{L}_{aug} . As a consequence, the (x_v, λ_v) component of the trajectory converges to an optimizer of (15).

In this case too, the saddle-point dynamics (26) is amenable to distributed implementation, cf. Remark VI.1, as one can write $\mathcal{C} = \prod_{i=1}^n \{(x, \lambda) \in \mathbb{R}^d \times \mathbb{R}_{\geq 0} \mid \lambda \geq a \|(x_{1:m-1}; 1)\|^2\}$.

VIII. SIMULATIONS

Here we illustrate the application of the distributed algorithm (26) to find a data-driven solution for the regression problem with quadratic loss function and an affine predictor [32, Chapter 3], commonly termed as the least-squares problem. Assume $n = 10$ agents with communication topology defined by an undirected ring with additional edges $\{(1, 4), (2, 5), (3, 7), (6, 10)\}$. The weight of each edge is equal to one. We consider data points of the form $\hat{\xi}^k = (\hat{w}^k, \hat{y}^k)$ consisting of the input $\hat{w}^k \in \mathbb{R}^4$ and the output $\hat{y}^k \in \mathbb{R}$ pairs. The objective is to find an affine predictor $x \in \mathbb{R}^5$ using the dataset such that, ideally, for any new data point $\xi = (w, y)$, the predictor $x^\top(w; 1)$ is equal to y . One way of finding such a predictor x is to solve the following problem

$$\inf_x \mathbb{E}_{\mathbb{P}} [f(x, w, y)] \quad (46)$$

where \mathbb{P} is the probability distribution of the data (w, y) and $f : \mathbb{R}^5 \times \mathbb{R}^5 \rightarrow \mathbb{R}$ is the quadratic loss function, i.e., $f(x, w, y) = (x^\top(w; 1) - y)^2$, corresponding to the case considered in Section VII-B2.

To find the data-driven solution, we assume that each agent in the network has 30 i.i.d samples of (w, y) and hence $N = 300$ is the total number of samples. The dataset is generated by assuming the input vector w having a standard multivariate normal distribution, that is, zero mean and covariance as the identity matrix I_4 . The output y is assigned values $y = [1, 4, 3, 2] * w + v$ where v is a random variable, uniformly distributed over the interval $[-1, 1]$. This defines completely the distribution \mathbb{P} of (w, y) . Let $\beta \in (0, 1)$ such that $\epsilon_N(\beta) = 0.05$. This value is assumed to be computed by the agents beforehand. This defines completely the distributed optimization problem (15).

Figure 1 shows the execution of the distributed algorithm (26) that solves this problem. The trajectories converge to an equilibrium of the dynamics (26) whose (x_v, λ_v) component corresponds to an optimizer of (15), consistent with Corollary VII.6.

To evaluate the quality of the obtained solution, we compute the average value of the loss function f for a randomly generated validation dataset consisting of $N_{\text{val}} = 10^4$ data points $\{(\hat{w}_{\text{val}}^k, \hat{y}_{\text{val}}^k)\}_{k=1}^{N_{\text{val}}}$. These points are i.i.d with the same distribution as that of the training dataset generated above. Given the obtained solution $(\mathbf{1}_{10} \otimes x^*, \lambda^* \mathbf{1}_{10})$, see Figure 1,




Fig. 1. Illustration of the execution of the dynamics (26) to find a data-driven solution of the regression problem (46). Plots (a) and (b) depict the evolution of the primal variables of the distributed optimization problem (15) defined for (46) with $\epsilon_N(\beta) = 0.05$ (for the sake of simplicity we have not shown the dual variables). The number of agents is 10 and each agent collects 30 i.i.d samples of the random variable. The initial condition $(x_v(0), \lambda_v(0))$ is chosen randomly from the set $[0, 5]^{50} \times [30, 80]^{10}$ and $\nu(0) = \mathbf{0}_{10}$, $\eta(0) = \mathbf{0}_{50}$, and $\xi^k(0) = \mathbf{0}_5$ for all $k \in [N]$. The primal variables converge to $x_v = \mathbf{1}_{10} \otimes x^*$ and $\lambda_v = \lambda^* \mathbf{1}_{10}$ with $x^* = (0.9460; 3.9624; 2.9840; 1.9878; 0.0268)$ and $\lambda^* = 58.4100$. This is an equilibrium point of X_{sp} as well as an optimizer of (15). The agents reach consensus early in the execution. Therefore, we only observe 5 (resp. 1) curves in the plot of x_v (resp. λ_v), each corresponding to a component of x (resp. λ). To depict this consensus clearly, we provided the zoomed part of the initial 10 seconds for both x_v and λ_v plots.

we evaluate

$$f_{\text{val}}^{N_{\text{val}}}(x^*) = \frac{1}{N_{\text{val}}} \sum_{k=1}^{N_{\text{val}}} f(x^*, \hat{w}_{\text{val}}^k, \hat{y}_{\text{val}}^k) \quad (47)$$

and get $f_{\text{val}}^{N_{\text{val}}}(x^*) = 0.3387$. This is the average loss for the solution x^* obtained by the agents cooperating with each other, essentially fusing the information of the 300 data points. Note that each agent individually can also solve a data-driven solution with the samples gathered by it. However, the solution obtained in such a manner, in general, incurs a higher average loss. In the current setup, if agent 1 solves (13) only with the data available to it (and keeping other parameters equal), then it gets the optimizer as $x^{\text{opt},1} = (0.8548; 3.8933; 2.8623; 2.1317; 0.2227)$. Using the validation dataset, we obtain $f_{\text{val}}^{N_{\text{val}}}(x^{\text{opt},1}) = 0.4520$, which is significantly greater than $f_{\text{val}}^{N_{\text{val}}}(x^*)$. This shows the value of cooperation, that is, fusing the information contained in the data available to different agents leads to an optimizer with better out-of-sample performance. To highlight this fact further, Figure 2 shows the effect of the number of cooperating agents on the average loss incurred by the obtained solution to the data-driven optimization problem. As the plot shows,




Fig. 2. Relative benefit of cooperation versus isolation. The x-axis represents the number of agents cooperating in our example. The y-axis plots the function $R(i) := \frac{f_{\text{val}}^{N_{\text{val}}}(x^{\text{opt},1}) - f_{\text{val}}^{N_{\text{val}}}(x^{\text{opt},i})}{f_{\text{val}}^{N_{\text{val}}}(x^{\text{opt},1})} \times 100\%$. Here, the function $f_{\text{val}}^{N_{\text{val}}}$ is given in (47) and $x^{\text{opt},i}$ denotes the solution of (13) determined by i cooperating agents. That is, $R(i)$ for each $i \in [n]$ represents the relative decrease in the average loss when i agents cooperate, taking the base case as no cooperation ($i = 1$). Thus, as more agents cooperate, the average loss decreases. The error bars represent the standard deviation of $R(i)$ values as we carry out 100 runs of the simulation, randomly generating different validation data each time.

the improvement in performance due to coordination becomes more prominent as the size of the coordination agents grows.

IX. CONCLUSIONS

We have considered a cooperative stochastic optimization problem, where a group of agents rely on their individually collected data to collectively determine a data-driven solution with guaranteed out-of-sample performance. Our technical approach has proceeded by first developing a reformulation in the form of a distributed optimization problem, leading us to the identification of an augmented Lagrangian function whose saddle points have a one-to-one correspondence with the primal-dual optimizers. This characterization relies upon certain interchangeability properties between the min and max operators. Our discussion has identified several classes of objective functions for which these properties hold: convex-concave functions, convex-convex functions quadratic in the data, and convex-convex functions associated to least-squares problems. Building on the analytical results, we have designed a distributed saddle-point coordination algorithm where agents share their individual estimates about the solution, not the collected data. We have also formally established the asymptotic convergence of the algorithm to the solution of the cooperative stochastic optimization problem. Future work will explore the characterization of the algorithm convergence rate, the design of strategies capable of tracking the solution of the stochastic optimization problem when new data becomes available in an online fashion, and the analysis of scenarios with network chance constraints.

REFERENCES

- [1] A. Cherukuri and J. Cortés, “Data-driven distributed optimization using Wasserstein ambiguity sets,” in *Allerton Conf. on Communications, Control and Computing*, (Monticello, IL), 2017. To appear.
- [2] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on stochastic programming*. Philadelphia, PA: SIAM, 2014.

- [3] P. M. Esfahani and D. Kuhn, "Data-driven distributionally robust optimization using the Wasserstein metric: performance guarantees and tractable reformulations," 2015. Available at <https://arxiv.org/abs/1505.05116>.
- [4] R. Gao and A. J. Kleywegt, "Distributionally robust stochastic optimization with Wasserstein distance," 2016. Available at <https://arxiv.org/abs/1604.02199>.
- [5] F. Luo and S. Mehrotra, "Decomposition algorithm for distributionally robust optimization using Wasserstein metric," 2017. Available at <https://arxiv.org/abs/1704.03920>.
- [6] R. Jiang and Y. Guan, "Data-driven chance constrained stochastic program," *Mathematical Programming, Series A*, vol. 158, pp. 291–327, 2016.
- [7] E. Erdoğan and G. Iyengar, "Ambiguous chance constrained problems and robust optimization," *Mathematical Programming, Series B*, vol. 107, pp. 37–61, 2006.
- [8] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*. Athena Scientific, 1997.
- [9] M. G. Rabbat and R. D. Nowak, "Quantized incremental algorithms for distributed optimization," *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 4, pp. 798–808, 2005.
- [10] P. Wan and M. D. Lemmon, "Event-triggered distributed optimization in sensor networks," in *Symposium on Information Processing of Sensor Networks*, (San Francisco, CA), pp. 49–60, 2009.
- [11] A. Nedic and A. Ozdaglar, "Distributed subgradient methods for multi-agent optimization," *IEEE Transactions on Automatic Control*, vol. 54, no. 1, pp. 48–61, 2009.
- [12] B. Johansson, M. Rabi, and M. Johansson, "A randomized incremental subgradient method for distributed optimization in networked systems," *SIAM Journal on Control and Optimization*, vol. 20, no. 3, pp. 1157–1170, 2009.
- [13] M. Zhu and S. Martínez, "On distributed convex optimization under inequality and equality constraints," *IEEE Transactions on Automatic Control*, vol. 57, no. 1, pp. 151–164, 2012.
- [14] J. Wang and N. Elia, "A control perspective for centralized and distributed convex optimization," in *IEEE Conf. on Decision and Control*, (Orlando, Florida), pp. 3800–3805, 2011.
- [15] A. Nedić, "Distributed optimization," in *Encyclopedia of Systems and Control* (J. Baillieul and T. Samad, eds.), New York: Springer, 2015.
- [16] B. Ghahsifard and J. Cortés, "Distributed continuous-time convex optimization on weight-balanced digraphs," *IEEE Transactions on Automatic Control*, vol. 59, no. 3, pp. 781–786, 2014.
- [17] J. Konečný, H. B. McMahan, D. Ramage, and P. Richtárik, "Federated optimization: distributed machine learning for on-device intelligence," 2016. Available at <https://arxiv.org/abs/1610.02527>.
- [18] F. Bullo, J. Cortés, and S. Martínez, *Distributed Control of Robotic Networks*. Applied Mathematics Series, Princeton University Press, 2009. Electronically available at <http://coordinationbook.info>.
- [19] R. T. Rockafellar, *Convex Analysis*. Princeton Landmarks in Mathematics and Physics, Princeton, NJ: Princeton University Press, 1997. Reprint of 1970 edition.
- [20] A. Bacciotti and F. Ceragioli, "Nonpathological Lyapunov functions and discontinuous Caratheodory systems," *Automatica*, vol. 42, no. 3, pp. 453–458, 2006.
- [21] J. Cortés, "Discontinuous dynamical systems - a tutorial on solutions, nonsmooth analysis, and stability," *IEEE Control Systems Magazine*, vol. 28, no. 3, pp. 36–73, 2008.
- [22] A. Nagurny and D. Zhang, *Projected Dynamical Systems and Variational Inequalities with Applications*, vol. 2 of *International Series in Operations Research and Management Science*. Dordrecht, The Netherlands: Kluwer Academic Publishers, 1996.
- [23] S. Lang, *Real and Functional Analysis*. New York: Springer, 3 ed., 1993.
- [24] F. H. Clarke, *Optimization and Nonsmooth Analysis*. Canadian Mathematical Society Series of Monographs and Advanced Texts, Wiley, 1983.
- [25] N. Fournier and A. Guillin, "On the rate of convergence in Wasserstein distance of the empirical measure," *Probability Theory and Related Fields*, vol. 162, no. 3, pp. 707–738, 2015.
- [26] A. E. Ozdaglar and P. Tseng, "Existence of global minima for constrained optimization," *Journal of Optimization Theory & Applications*, vol. 128, no. 3, pp. 523–546, 2006.
- [27] A. Cherukuri, B. Ghahsifard, and J. Cortés, "Saddle-point dynamics: conditions for asymptotic stability of saddle points," *SIAM Journal on Control and Optimization*, vol. 55, no. 1, pp. 486–511, 2017.
- [28] X. L. Sun, D. Li, and K. I. M. McKinnon, "On saddle points of augmented Lagrangians for constrained nonconvex optimization," *SIAM Journal on Optimization*, vol. 15, no. 4, pp. 1128–1146, 2005.
- [29] A. Cherukuri, E. Mallada, and J. Cortés, "Asymptotic convergence of primal-dual dynamics," *Systems & Control Letters*, vol. 87, pp. 10–15, 2016.
- [30] S. P. Bhat and D. S. Bernstein, "Nontangency-based Lyapunov tests for convergence and stability in systems having a continuum of equilibria," *SIAM Journal on Control and Optimization*, vol. 42, no. 5, pp. 1745–1775, 2003.
- [31] R. Goebel, "Stability and robustness for saddle-point dynamics through monotone mappings," *Systems & Control Letters*, vol. 108, pp. 16–22, 2017.
- [32] B. Schölkopf and A. J. Smola, *Learning with kernels*. Cambridge, MA: MIT Press, 2002.



Ashish Cherukuri is a Postdoctoral Researcher at the Automatic Control Laboratory, ETH Zürich, Switzerland. He received the B.Tech degree in mechanical engineering from the Indian Institute of Technology, Delhi, India, in 2008, the M.Sc. degree in mechanical engineering from ETH, Zurich, Switzerland, in 2010, and the Ph.D. degree in mechanical and aerospace engineering at the University of California, San Diego, CA, USA, in 2017. His research interests include dynamical systems, game theory, distributed optimization, and power networks. He received the Excellence Scholarship in 2008 from ETH Zürich, the Focht-Powell Fellowship in 2012 from the University of California, San Diego, and the Robert E. Skelton Systems and Control Dissertation Award in 2017 from the University of California, San Diego. He was a finalist for Best Student Paper Award in the American Control Conference, 2015.



Jorge Cortés (M'02-SM'06-F'14) received the Licenciatura degree in mathematics from Universidad de Zaragoza, Zaragoza, Spain, in 1997, and the Ph.D. degree in engineering mathematics from Universidad Carlos III de Madrid, Madrid, Spain, in 2001. He held postdoctoral positions with the University of Twente, Twente, The Netherlands, and the University of Illinois at Urbana-Champaign, Urbana, IL, USA. He was an Assistant Professor with the Department of Applied Mathematics and Statistics, University of California, Santa Cruz, CA, USA, from

2004 to 2007. He is currently a Professor in the Department of Mechanical and Aerospace Engineering, University of California, San Diego, CA, USA. He is the author of *Geometric, Control and Numerical Aspects of Nonholonomic Systems* (Springer-Verlag, 2002) and co-author (together with F. Bullo and S. Martínez) of *Distributed Control of Robotic Networks* (Princeton University Press, 2009). He has been an IEEE Control Systems Society Distinguished Lecturer (2010–2014) and is an elected member for 2018–2020 of the Board of Governors of the IEEE Control Systems Society. His current research interests include distributed control and optimization, network science, reasoning under uncertainty, multi-agent coordination in power networks, robotics, and transportation, nonsmooth analysis, and geometric mechanics and control.