
Convex-Concave Min-Max Stackelberg Games

Denizalp Goktas

Department of Computer Science
Brown University
Providence, RI 02912
denizalp_goktas@brown.edu

Amy Greenwald

Department of Computer Science
Brown University
Providence, RI 02912
amy_greenwald@brown.edu

Abstract

Min-max optimization problems (i.e., min-max games) have been attracting a great deal of attention because of their applicability to a wide range of machine learning problems. Although significant progress has been made recently, the literature to date has focused on games with independent strategy sets; little is known about solving games with dependent strategy sets, which can be characterized as min-max Stackelberg games. We introduce two first-order methods that solve a large class of convex-concave min-max Stackelberg games, and show that our methods converge in polynomial time. Min-max Stackelberg games were first studied by Wald, under the posthumous name of Wald’s maximin model, a variant of which is the main paradigm used in robust optimization, which means that our methods can likewise solve many convex robust optimization problems. We observe that the computation of competitive equilibria in Fisher markets also comprises a min-max Stackelberg game. Further, we demonstrate the efficacy and efficiency of our algorithms in practice by computing competitive equilibria in Fisher markets with varying utility structures. Our experiments suggest potential ways to extend our theoretical results, by demonstrating how different smoothness properties can affect the convergence rate of our algorithms.

1 Introduction

Min-max optimization problems have attracted a great deal of attention recently because of their applicability to a wide range of machine learning problems. Examples of settings in which min-max optimization problems arise include, but are not limited to, reinforcement learning [15], generative adversarial networks [66], fairness in machine learning [14, 23, 48, 68, 82], adversarial learning [72], generative adversarial imitation learning [11, 32], and statistical learning (e.g., learning parameters of exponential families) [14]. These applications often require solving a **constrained min-max optimization problem** (with independent feasible sets): i.e., $\min_{x \in X} \max_{y \in Y} f(x, y)$, where $f : X \times Y \rightarrow \mathbb{R}$ is continuous, and $X \subset \mathbb{R}^n$ and $Y \subset \mathbb{R}^m$ are non-empty and compact.

A **convex-concave** constrained min-max optimization problem is one in which f is convex in x and concave in y . In the special case of convex-concave objective functions, the seminal minimax theorem holds: i.e., $\min_{x \in X} \max_{y \in Y} f(x, y) = \max_{y \in Y} \min_{x \in X} f(x, y)$ [55]. This theorem guarantees the existence of a **saddle point** (i.e., a point that is simultaneously a minimum of f in the x -direction and a maximum of f in the y -direction), and allows us to interpret the optimization problem as a simultaneous-move, zero-sum game, where y^* (resp. x^*) is a best-response of the outer (resp. inner) player to their opponent’s strategy x^* (resp. y^*), in which case a saddle point is also called a minimax point or a **Nash equilibrium**.

In this paper, we show that the computation of competitive equilibria in Fisher markets [10], a central solution concept for one of the most well-studied market models in algorithmic game theory [57], can be understood as a convex-concave min-max optimization problem, albeit with *dependent*

feasible sets. We define a **constrained min-max optimization problem with dependent feasible sets** as an optimization problem of the following form: $\min_{\mathbf{x} \in X} \max_{\mathbf{y} \in Y: \mathbf{g}(\mathbf{x}, \mathbf{y}) \geq \mathbf{0}} f(\mathbf{x}, \mathbf{y})$, where $f : X \times Y \rightarrow \mathbb{R}$ is continuous, $X \subset \mathbb{R}^n$ and $Y \subset \mathbb{R}^m$ are non-empty and compact, and $\mathbf{g}(\mathbf{x}, \mathbf{y}) = (g_1(\mathbf{x}, \mathbf{y}), \dots, g_K(\mathbf{x}, \mathbf{y}))^T$ with $g_k : X \times Y \rightarrow \mathbb{R}$.¹ Unfortunately, certain desirable properties of standard convex-concave min-max optimization problems with independent feasible sets do not carry over to the more general dependent setting. First and foremost, the minimax theorem [55] does not hold, which in turn precludes the existence of Nash equilibrium in general:

Example 1.1. Consider the following constrained min-max optimization problem with dependent feasible sets: $\min_{x \in [-1, 1]} \max_{y \in [-1, 1]: x+y \leq 0} x^2 + y + 1$. The optimum is $x^* = 1/2, y^* = -1/2$, with value $3/4$. Now, consider the same problem, with the order of the min and the max reversed: $\max_{y \in [-1, 1]} \min_{x \in [-1, 1]: x+y \leq 0} x^2 + y + 1$. The optimum is now $x^* = -1, y^* = 1$, with value 3.

Without a minimax theorem, constrained min-max optimization problems with dependent feasible sets are more appropriately viewed as two-player sequential, i.e., Stackelberg, games where the outer player chooses $\mathbf{x} \in X$ before the inner player responds with their choice of $\mathbf{y}(\mathbf{x}) \in Y$ s.t. $\mathbf{g}(\mathbf{x}, \mathbf{y}(\mathbf{x})) \geq \mathbf{0}$. The relevant equilibrium concept is then a Stackelberg [78] equilibrium,² in which the outer player optimizes their choice assuming the inner player will best-respond: i.e., optimize their choice in turn. We thus refer to constrained min-max optimization problems with dependent feasible sets as **min-max Stackelberg games**. In such games, the **value function**³ $V : X \rightarrow \mathbb{R}$ is defined as $V(\mathbf{x}) = \max_{\mathbf{y} \in Y: \mathbf{g}(\mathbf{x}, \mathbf{y}) \geq \mathbf{0}} f(\mathbf{x}, \mathbf{y})$. This function represents the outer player's loss, assuming the inner player chooses a feasible best-response, so it is the function the outer player seeks to minimize. The inner player seeks only to maximize the objective function. The first-order necessary and sufficient conditions for a tuple $(\mathbf{x}^*, \mathbf{y}^*) \in X \times Y$ to be a Stackelberg equilibrium in a convex-concave min-max Stackelberg game are given by the KKT stationarity conditions for the two players' optimization problems, namely $\min_{\mathbf{x} \in X} V(\mathbf{x})$ for the outer player, and $\max_{\mathbf{y} \in Y: \mathbf{g}(\mathbf{x}, \mathbf{y}) \geq \mathbf{0}} f(\mathbf{x}, \mathbf{y})$ for the inner player.

In the independent strategy setting, Danskin's theorem [17] states that $\nabla_{\mathbf{x}} V(\mathbf{x}) = \nabla_{\mathbf{x}} f(\mathbf{x}, \mathbf{y}^*)$, where $\mathbf{y}^* \in \arg \max_{\mathbf{y} \in Y} f(\mathbf{x}, \mathbf{y})$. In other words, when there is no dependence among the players' strategy sets, the gradient of the value function coincides with that of the objective function. The first-order necessary and sufficient conditions for a tuple $(\mathbf{x}^*, \mathbf{y}^*)$ to be an (interior) saddle point are for it to be a stationary point of f (i.e., $\nabla_{\mathbf{x}} f(\mathbf{x}^*, \mathbf{y}^*) = \nabla_{\mathbf{y}} f(\mathbf{x}^*, \mathbf{y}^*) = \mathbf{0}$). It is therefore logical for players in the independent/simultaneous-move setting to follow the gradient of the objective function. In the dependent/sequential setting, however, the direction of steepest descent (resp. ascent) for the outer (resp. inner) player is the gradient of their *value* function.

Example 1.2. Consider, once again, the problem posed in Example [1.1] and recall Jin, Netrapalli, and Jordan's gradient descent with max-oracle algorithm [41]: $\mathbf{x}^{(t+1)} = \mathbf{x}^{(t)} - \eta \nabla_{\mathbf{x}} f(\mathbf{x}^{(t)}, \mathbf{y}^*(\mathbf{x}^{(t)}))$, where $\mathbf{y}^*(\mathbf{x}^{(t)}) \in \arg \max_{\mathbf{y} \in Y: \mathbf{g}(\mathbf{x}^{(t)}, \mathbf{y}) \geq \mathbf{0}} f(\mathbf{x}^{(t)}, \mathbf{y})$ for $\eta > 0$. Applied to this sample problem, with $\eta = 1$, this algorithm yields the following update rule: $x^{(t+1)} = x^{(t)} - 2x^{(t)} = -x^{(t)}$. Thus, letting $x^{(0)}$ equal any feasible x , the output cycles between x and $-x$, so that the average of the iterates converges to $x^* = 0$ (with $y^* = 0$), which is not a Stackelberg equilibrium, as the Stackelberg equilibrium of this game is $x^* = 1/2, y^* = -1/2$.

Now consider an algorithm that updates based not on gradient of the objective function, but of the value function, namely $\mathbf{x}^{(t+1)} = \mathbf{x}^{(t)} - \eta \nabla_{\mathbf{x}} V(\mathbf{x}^{(t)})$. The value function is $V(x) = \max_{y \in [-1, 1]: x+y \leq 0} x^2 + y + 1 = x^2 - x + 1$, with gradient $V'(x) = 2x - 1$. Thus, when $\eta = 1$, this algorithm yields the following update rule: $x^{(t+1)} = x^{(t)} - V'(x^{(t)}) = x^{(t)} - 2x^{(t)} + 1 = -x^{(t)} + 1$. If we run this algorithm from initial point $x^{(0)} = 1/8$, we get $x^{(1)} = -1/8 + 1 = 7/8, x^{(2)} = -7/8 + 1 = 1/8$, and so on. The average of the iterates $1/8, 7/8, 1/8, \dots$ converges to $x^* = 1/2(1/8 + 7/8) = 1/2$, and correspondingly $y^* = -1/2$, which is indeed the Stackelberg equilibrium.

¹Although notationally similar, we note that this model is more general than that of Daskalakis, Skoulakis, and Zampetakis [18]. The authors state that under the assumptions of the minimax theorem [55], a Nash equilibrium is guaranteed to exist in their model. When there is dependence among the players' feasible sets, however, this claim does not hold (see Example [1.1]).

²In constrained convex-concave min-max optimization problems with independent feasible sets, the notions of Nash and Stackelberg equilibria coincide. They also coincide in the dependent setting, when Nash equilibria exist, but they do not coincide in general, since Nash equilibria need not exist.

³Note that this use of the term value function in economics is distinct from its use in reinforcement learning.

In this paper, we introduce two first-order (subgradient) methods that solve min-max Stackelberg games—to our knowledge the first such methods. Our approach relies on a new generalization of a series of fundamental results in mathematical economics known as **envelope theorems** [1, 50]. Envelope theorems generalize aspects of Danskin’s theorem, by providing explicit formulas for the gradient of the value function in dependent strategy settings, *when a derivative is guaranteed to exist*. To sidestep the differentiability issue, we introduce a generalized envelope theorem that gives an explicit formula for the subdifferential of the value function in dependent strategy settings.

Our first algorithm follows Jin, Netrapalli, and Jordan [41], assuming access to a max-oracle that returns $\mathbf{y}^* \in \arg \max_{\mathbf{y} \in Y: g(\mathbf{x}, \mathbf{y}) \geq 0} f(\mathbf{x}, \mathbf{y})$, given $\mathbf{x} \in X$. Hence, our first algorithm solves only for an optimal \mathbf{x}^* , while our second algorithm explicitly solves for both \mathbf{x}^* and \mathbf{y}^* . We show that both algorithms converge in polynomial time to a Stackelberg equilibrium of any convex-concave min-max Stackelberg game with nonempty, compact and convex dependent strategy sets. In Table 1 we summarize the iteration complexities of our algorithms, i.e., the number of iterations required to achieve an ε -approximate equilibrium, where ε is the desired precision.

Table 1: Iteration complexities of Algorithms 1 and 2 for min-max Stackelberg games. Here, μ_x and μ_y are strong convexity/concavity parameters.

Properties of f	Iteration Complexity	
	Algorithm 1	Algorithm 2
μ_x -Strongly-Convex- μ_y -Strongly-Concave	$O(\varepsilon^{-1})$	$\tilde{O}(\varepsilon^{-1})$
μ_x -Strongly-Convex-Concave		$O(\varepsilon^{-2})$
Convex- μ_y -Strongly-Concave	$O(\varepsilon^{-2})$	$\tilde{O}(\varepsilon^{-2})$
Convex-Concave		$O(\varepsilon^{-3})$

Finally, we apply our results to the computation of competitive equilibria in Fisher markets. In this context, our method for solving a min-max Stackelberg game reduces to solving the market in a decentralized manner using a natural market dynamic called tâtonnement [80]. We demonstrate the efficacy and efficiency of our algorithms in practice by running a series of experiments in which we compute competitive equilibria in Fisher markets with varying utility structures—specifically, linear, Cobb-Douglas, and Leontief. Although our theoretical results do not apply to all these Fisher markets—Leontief utilities, in particular, are not differentiable—tâtonnement converges in all our experiments. That said, the rate of convergence does seem to depend on the smoothness characteristics of the utility structures; we observe slower convergence for Leontief utilities, and faster convergence than our theory predicts for Cobb-Douglas utilities, which are not only differentiable, but whose value function is also differentiable.

Related Work Our model of min-max Stackelberg games seems to have first been studied by Wald, under the posthumous name of Wald’s maximin model [79]. A variant of Wald’s maximin model is the main paradigm used in robust optimization, a fundamental framework in operations research for which many methods have been proposed [6, 56, 61]. Shimizu and Aiyoshi [70, 71] proposed the first algorithm to solve min-max Stackelberg games via a relaxation to a constrained optimization problem with infinitely many constraints, which nonetheless seems to perform well in practice. More recently, Segundo, Krohling, and Cosme [69] proposed an evolutionary algorithm for these games, but they provided no guarantees. As pointed out by Postek and Shtern, all prior methods either require oracles and are stochastic in nature [6], or rely on a binary search for the optimal value, which can be computationally complex [56]. The algorithms we propose in this paper circumvent the aforementioned issues and can be used to solve a large class of convex robust optimization problems in a simple and efficient manner.

Extensive-form games in which players’ strategy sets can depend on other players’ actions have been studied by Davis, Waugh, and Bowling [19] assuming payoffs are bilinear, and by Farina, Kroer, and Sandholm [28] for another specific class of convex-concave payoffs. Fabiani et al. [25] and Kebriaei and Iannelli [43] study more general settings than ours, namely non-zero-sum Stackelberg games with more than two players. Both sets of authors derive convergence guarantees assuming specific payoff structures, but their algorithms do not converge in polynomial time.

Min-max Stackelberg games naturally model various economic settings. They are related to the abstract economies first studied by Arrow and Debreu [4]; however, the solution concept that has been the focus of this literature is generalized Nash equilibrium [26, 27], which, like Stackelberg, is a weaker solution concept than Nash, but which makes the arguably unreasonable assumption that the players move simultaneously and nonetheless satisfy the constraint dependencies on their strategies imposed by one another’s moves. See Appendix A for a more detailed discussion of generalized Nash equilibria versus Stackelberg equilibria.

Duetting et al. [22] study optimal auction design problems. They propose a neural network architecture called RegretNet that represents optimal auctions, and train their networks using Algorithm 3.4. Optimal auction design problems can be seen as min-max Stackelberg games; however, as their objectives are non-convex-concave in general, our guarantees do not apply.

In this paper, we observe that solving for the competitive equilibrium of a Fisher market can also be seen as solving a (convex-concave) min-max Stackelberg game. The study of the computation of competitive equilibria in Fisher markets was initiated by Devanur et al. [20], who provided a polynomial-time method for the case of linear utilities. Jain, Vazirani, and Ye [40] subsequently showed that a large class of Fisher markets could be solved in polynomial-time using interior point methods. Recently, Gao and Kroer [29] studied an alternative family of first-order methods for solving Fisher markets (only; not min-max Stackelberg games more generally), assuming linear, quasilinear, and Leontief utilities, as such methods can be more efficient when markets are large.

See Appendix G for a detailed discussion of recent progress on solving min-max Stackelberg games, both in the convex-concave case and the non-convex-concave case.

2 Preliminaries

Notation We use Roman uppercase letters to denote sets (e.g., X), bold uppercase letters to denote matrices (e.g., \mathbf{X}), bold lowercase letters to denote vectors (e.g., \mathbf{p}), and Roman lowercase letters to denote scalar quantities, (e.g., c). We denote the i th row vector of a matrix (e.g., \mathbf{X}) by the corresponding bold lowercase letter with subscript i (e.g., \mathbf{x}_i). Similarly, we denote the j th entry of a vector (e.g., \mathbf{p} or \mathbf{x}_i) by the corresponding Roman lowercase letter with subscript j (e.g., p_j or x_{ij}). We denote the set of integers $\{1, \dots, n\}$ by $[n]$, the set of natural numbers by \mathbb{N} , the set of real numbers by \mathbb{R} , the set of non-negative real numbers by \mathbb{R}_+ . We denote by Π_Y the Euclidean projection operator onto the set $Y \subset \mathbb{R}^m$: i.e., $\Pi_Y(\mathbf{y}) = \arg \min_{\mathbf{z} \in Y} \|\mathbf{y} - \mathbf{z}\|_2$.

A min-max Stackelberg game, denoted (X, Y, f, \mathbf{g}) , is a two-player, zero-sum game, where one player, who we call the \mathbf{x} -player (resp. the \mathbf{y} -player), is trying to minimize their loss (resp. maximize their gain), defined by a continuous **objective function** $f : X \times Y \rightarrow \mathbb{R}$, by choosing a strategy from a compact **strategy set** $X \subset \mathbb{R}^n$ (resp. $Y \subset \mathbb{R}^m$), and s.t. $\mathbf{g}(\mathbf{x}, \mathbf{y}) \geq \mathbf{0}$ where $\mathbf{g}(\mathbf{x}, \mathbf{y}) = (g_1(\mathbf{x}, \mathbf{y}), \dots, g_K(\mathbf{x}, \mathbf{y}))^T$ with $g_k : X \times Y \rightarrow \mathbb{R}$. A strategy profile $(\mathbf{x}, \mathbf{y}) \in X \times Y$ is said to be **feasible** iff for all $k \in [K]$, $g_k(\mathbf{x}, \mathbf{y}) \geq 0$. The function f maps a pair of feasible strategies taken by the players $(\mathbf{x}, \mathbf{y}) \in X \times Y$ to a real value (i.e., a payoff), which represents the loss (resp. the gain) of the \mathbf{x} -player (resp. \mathbf{y} -player). A min-max game is said to be convex-concave if the objective function f is convex-concave.

The relevant solution concept for Stackelberg games is the **Stackelberg equilibrium**:

Definition 2.1 (Stackelberg Equilibrium). *Consider the min-max Stackelberg game (X, Y, f, \mathbf{g}) . A strategy profile $(\mathbf{x}^*, \mathbf{y}^*) \in X \times Y$ such that $\mathbf{g}(\mathbf{x}^*, \mathbf{y}^*) \geq \mathbf{0}$ is a (ε, δ) -Stackelberg equilibrium if $\max_{\mathbf{y} \in Y: \mathbf{g}(\mathbf{x}^*, \mathbf{y}) \geq \mathbf{0}} f(\mathbf{x}^*, \mathbf{y}) - \delta \leq f(\mathbf{x}^*, \mathbf{y}^*) \leq \min_{\mathbf{x} \in X} \max_{\mathbf{y} \in Y: \mathbf{g}(\mathbf{x}, \mathbf{y}) \geq \mathbf{0}} f(\mathbf{x}, \mathbf{y}) + \varepsilon$.*

Intuitively, a (ε, δ) -Stackelberg equilibrium is a point at which the \mathbf{x} -player’s (resp. \mathbf{y} -player’s) payoff is no more than ε (resp. δ) away from its optimum. By Proposition B.2, a $(0, 0)$ -Stackelberg equilibrium is guaranteed to exist in min-max Stackelberg, its value is unique, and by Proposition B.3 the set of Stackelberg equilibria is compact and convex.

Finally, we define several concepts that are needed in our convergence proofs. Given $A \subset \mathbb{R}^n$, the function $f : A \rightarrow \mathbb{R}$ is said to be ℓ_f -**Lipschitz-continuous** iff $\forall \mathbf{x}_1, \mathbf{x}_2 \in A, \|f(\mathbf{x}_1) - f(\mathbf{x}_2)\| \leq \ell_f \|\mathbf{x}_1 - \mathbf{x}_2\|$. If the gradient of f , ∇f , is $\ell_{\nabla f}$ -Lipschitz-continuous, we then refer to f as $\ell_{\nabla f}$ -**Lipschitz-smooth**. A function $f : A \rightarrow \mathbb{R}$ is μ -**strongly convex** if $f(\mathbf{x}_1) \geq f(\mathbf{x}_2) + \langle \nabla_x f(\mathbf{x}_2), \mathbf{x}_1 - \mathbf{x}_2 \rangle + \mu/2 \|\mathbf{x}_1 - \mathbf{x}_2\|^2$, and μ -**strongly concave** if $-f$ is μ -strongly convex.

3 First-Order Methods via an Envelope Theorem

The envelope theorems, popular tools in mathematical economics, allow for explicit formulas for the gradient of the value function in min-max games, even when the strategy sets are dependent. Afriat [1] appears to have been the first make use of the Lagrangian to differentiate the value function, though his conclusion was later obtained under weaker assumptions by Milgrom and Segal [50]. Our envelope theorem relies on the following assumptions:

Assumption 3.1. 1. f, g_1, \dots, g_K are continuous and convex-concave; 2. $\nabla_x f, \nabla_x g_1, \dots, \nabla_x g_K$ are continuous; and 3. (Slater's condition) $\forall \mathbf{x} \in X, \exists \hat{\mathbf{y}} \in Y$ s.t. $g_k(\mathbf{x}, \hat{\mathbf{y}}) > 0$, for all $k = 1, \dots, K$.

Milgrom and Segal's [50] envelope theorem provides an explicit formula for the gradient of the value function. When strategy sets are independent, under mild assumptions, this function is guaranteed to be differentiable [58]. When strategy sets are dependent, however, it is not necessarily differentiable, as seen in Example C.2. As a remedy, we present a subdifferential envelope theorem for non-differentiable value functions:

Theorem 3.2 (Subdifferential Envelope Theorem). *Consider the value function $V(\mathbf{x}) = \max_{\mathbf{y} \in Y: g(\mathbf{x}, \mathbf{y}) \geq 0} f(\mathbf{x}, \mathbf{y})$. Let $Y^*(\mathbf{x}) = \arg \max_{\mathbf{y} \in Y: g(\mathbf{x}, \mathbf{y}) \geq 0} f(\mathbf{x}, \mathbf{y})$ and suppose Assumption 3.1 holds. Then, at any point $\hat{\mathbf{x}} \in X$, $\partial_x V(\hat{\mathbf{x}}) =$*

$$\text{conv} \left(\bigcup_{\mathbf{y}^*(\hat{\mathbf{x}}) \in Y^*(\hat{\mathbf{x}})} \bigcup_{\lambda_k^*(\hat{\mathbf{x}}, \mathbf{y}^*(\hat{\mathbf{x}})) \in \Lambda^*(\hat{\mathbf{x}}, \mathbf{y}^*(\hat{\mathbf{x}}))} \left\{ \nabla_x f(\hat{\mathbf{x}}, \mathbf{y}^*(\hat{\mathbf{x}})) + \sum_{k=1}^K \lambda_k^*(\hat{\mathbf{x}}, \mathbf{y}^*(\hat{\mathbf{x}})) \nabla_x g_k(\hat{\mathbf{x}}, \mathbf{y}^*(\hat{\mathbf{x}})) \right\} \right), \quad (1)$$

where conv is the convex hull operator and $\lambda^*(\hat{\mathbf{x}}, \mathbf{y}^*(\hat{\mathbf{x}})) = (\lambda_1^*(\hat{\mathbf{x}}, \mathbf{y}^*(\hat{\mathbf{x}})), \dots, \lambda_K^*(\hat{\mathbf{x}}, \mathbf{y}^*(\hat{\mathbf{x}})))^T \in \Lambda^*(\hat{\mathbf{x}}, \mathbf{y}^*(\hat{\mathbf{x}}))$ are the optimal KKT multipliers associated with $\mathbf{y}^*(\hat{\mathbf{x}}) \in Y^*(\hat{\mathbf{x}})$.

The envelope theorem states that the gradient of a differentiable value function is the gradient of the Lagrangian evaluated at the optimal solution. Generalizing this fact, our subdifferential envelope theorem states that every subgradient of the value function, $V(\mathbf{x}) = \max_{\mathbf{y} \in Y: g(\mathbf{x}, \mathbf{y}) \geq 0} f(\mathbf{x}, \mathbf{y})$ is a convex combination of the values of the gradient of the Lagrangian evaluated at the optimal solutions $(\mathbf{y}^*(\mathbf{x}), \lambda^*(\mathbf{x}, \mathbf{y}^*(\mathbf{x}))) \in Y^*(\mathbf{x}) \times \Lambda^*(\mathbf{x}, \mathbf{y}^*(\mathbf{x}))$.

With our envelope theorem in hand, we are now ready to present two gradient-descent/ascent-type algorithms for min-max Stackelberg games, which follow the gradient of the value function.

Our first algorithm, **max-oracle gradient-descent**, following Jin, Netrapalli, and Jordan [41], assumes access to a max-oracle, which given $\mathbf{x} \in X$, returns a δ -maximum of the inner player's value function. That is, for all $\mathbf{x} \in X$, the max-oracle returns $\hat{\mathbf{y}} \in Y$ s.t. $g(\mathbf{x}, \hat{\mathbf{y}}) \geq 0$ and $f(\mathbf{x}, \hat{\mathbf{y}}) \geq \max_{\mathbf{y} \in Y: g(\mathbf{x}, \mathbf{y}) \geq 0} f(\mathbf{x}, \mathbf{y}) - \delta$. It then runs (sub)gradient descent on the outer player's value function, using Theorem 3.2 to compute the requisite subgradients. Inspired by the multi-step gradient-descent algorithm of Nouiheed et al. [58], our second algorithm, **nested gradient-descent/ascent** (see Appendix E), computes both \mathbf{x}^* and \mathbf{y}^* explicitly, without oracle access. We simply replace the max-oracle in our max-oracle gradient-descent algorithm by a projected gradient-ascent procedure, which again computes a δ -maximum of the inner player's value function.

Once $\hat{\mathbf{y}}$ is found at iteration t , one can compute optimal KKT multipliers $\lambda_1^*(\mathbf{x}^{(t)}, \hat{\mathbf{y}}(\mathbf{x}^{(t)})), \dots, \lambda_K^*(\mathbf{x}^{(t)}, \hat{\mathbf{y}}(\mathbf{x}^{(t)}))$ for the outer player's value function, either via a system of linear equations using the complementary slackness conditions and the value of the objective function at the optimal, namely $(\mathbf{x}^{(t)}, \hat{\mathbf{y}}(\mathbf{x}^{(t)}))$, or by running gradient descent on the Lagrangian for the dual variables. Additionally, most algorithms solving convex programs will return $\lambda^*(\mathbf{x}^{(t)}, \hat{\mathbf{y}}(\mathbf{x}^{(t)})) = (\lambda_1^*(\mathbf{x}^{(t)}, \hat{\mathbf{y}}(\mathbf{x}^{(t)})), \dots, \lambda_K^*(\mathbf{x}^{(t)}, \hat{\mathbf{y}}(\mathbf{x}^{(t)})))$ together with the optimal $\hat{\mathbf{y}}(\mathbf{x}^{(t)})$ without incurring any additional computational expense. As a result, we assume that the optimal KKT multipliers $\lambda^*(\mathbf{x}^{(t)}, \hat{\mathbf{y}}(\mathbf{x}^{(t)}))$ associated with a solution $\hat{\mathbf{y}}(\mathbf{x}^{(t)})$ can be computed in constant time.

Having explained our two procedures, our next task is to derive their convergence rates. It turns out that under very mild assumptions, i.e., when Assumption 3.1 holds, the outer player's value function is Lipschitz continuous in \mathbf{x} . More precisely, the objective function f is ℓ_f -Lipschitz in \mathbf{x} , where $\ell_f = \max_{(\hat{\mathbf{x}}, \hat{\mathbf{y}}) \in X \times Y} \|\nabla_x f(\hat{\mathbf{x}}, \hat{\mathbf{y}})\|$ ⁴. As f is ℓ_f -Lipschitz continuous in \mathbf{x} , the value function is

⁴This max norm is well-defined since $\nabla_x f$ is continuous and the constraint set is non-empty and compact.

Algorithm 1 Max-Oracle Gradient Descent

Inputs: $X, Y, f, g, \eta, T, \mathbf{x}^{(0)}$
Output: $(\mathbf{x}^*, \mathbf{y}^*)$

- 1: **for** $t = 1, \dots, T$ **do**
 - 2: Find $\mathbf{y}^{(t-1)} \in Y$ s.t. $f(\mathbf{x}^{(t-1)}, \mathbf{y}^{(t-1)}) \geq V(\mathbf{x}^{(t-1)}) - \delta$ & $\mathbf{g}(\mathbf{x}^{(t-1)}, \mathbf{y}^{(t-1)}) \geq \mathbf{0}$
 - 3: Set $\boldsymbol{\lambda}^{(t-1)} = \boldsymbol{\lambda}^*(\mathbf{x}^{(t-1)}, \mathbf{y}^{(t-1)})$
 - 4: Set $\mathbf{x}^{(t)} = \Pi_X \left(\mathbf{x}^{(t-1)} - \eta_t \left[\nabla_{\mathbf{x}} f(\mathbf{x}^{(t-1)}, \mathbf{y}^{(t-1)}) + \sum_{k=1}^K \lambda_k^{(t-1)} \nabla_{\mathbf{x}} g_k(\mathbf{x}^{(t-1)}, \mathbf{y}^{(t-1)}) \right] \right)$
 - 5: **end for**
 - 6: Find $\mathbf{y}^{(T)} \in Y$ s.t. $f(\mathbf{x}^{(T)}, \mathbf{y}^{(T)}) \geq V(\mathbf{x}^{(T)}) - \delta$ & $\mathbf{g}(\mathbf{x}^{(T)}, \mathbf{y}^{(T)}) \geq \mathbf{0}$
 - 7: **return** $(\mathbf{x}^{(T)}, \mathbf{y}^{(T)})$
-

also ℓ_f -Lipschitz.⁵ This fact suggests that an $(\varepsilon, \varepsilon)$ -Stackelberg equilibrium should be computable in $O(\varepsilon^{-2})$ iterations by our max-oracle gradient descent algorithm (Algorithm 1), since our method is a subgradient method.

Theorem 3.3. Consider a convex-concave min-max Stackelberg game, (X, Y, f, g) , where X is convex, and $(\mathbf{x}^*, \mathbf{y}^*)$ is a $(0, \delta)$ -Stackelberg equilibrium, and suppose Assumption 3.1 holds. If Algorithm 1 is run for T iterations, with step sizes η_t s.t. $\sum_{k=1}^T \eta_k^2 \leq \infty$, $\sum_{k=1}^T \eta_k = \infty$, and if $(\mathbf{x}_{\text{best}}^{(t)}, \mathbf{y}_{\text{best}}^{(t)}) \in \arg \min_{(\mathbf{x}^{(k)}, \mathbf{y}^{(k)}) : k \in [t]} f(\mathbf{x}^{(k)}, \mathbf{y}^{(k)})$, it holds that $\lim_{T \rightarrow \infty} f(\mathbf{x}_{\text{best}}^{(T)}, \mathbf{y}_{\text{best}}^{(T)}) = f(\mathbf{x}^*, \mathbf{y}^*)$. Furthermore, for $\varepsilon \in (0, 1)$, if we choose $T \geq N_T(\varepsilon) \in O(\varepsilon^{-2})$, then there exists an iteration $T^* \leq T$ s.t. $(\mathbf{x}_{\text{best}}^{(T^*)}, \mathbf{y}_{\text{best}}^{(T^*)})$ is an (ε, δ) -Stackelberg equilibrium.

As is expected, the $O(\varepsilon^{-2})$ iteration complexity can be improved to $O(\varepsilon^{-1})$, if additionally, f is strongly convex in \mathbf{x} . (See Appendix E, Theorem E.2). Combining the convergence results for our max-oracle gradient descent algorithm with convergence results for gradient descent [8], we obtain the following convergence rates for the nested gradient descent-ascent algorithm (Algorithm 2, Appendix E). We include the formal proof and statement for the case when Assumption 3.1 holds and f is Lipschitz-smooth in Appendix E (Theorem E.3). The other results follow similarly.

Theorem 3.4. Consider a convex-concave min-max Stackelberg game, (X, Y, f, g) , with X and Y convex. Suppose that Assumption 3.1 holds. Then, under standard assumptions on the step sizes, the iteration complexities given below hold for the computation of a $(\varepsilon, \varepsilon)$ -Stackelberg equilibrium:

	f is $\ell_{\nabla f}$ -smooth	f is $\ell_{\nabla f}$ -smooth + f Strongly Concave in \mathbf{y}
Assumption 3.1	$O(\varepsilon^{-3})$	$O(\varepsilon^{-2} \log(\varepsilon^{-1}))$
Assumption 3.1 + f Strongly Convex in \mathbf{x}	$O(\varepsilon^{-2})$	$O(\varepsilon^{-1} \log(\varepsilon^{-1}))$

Since the value function in the convex-concave dependent setting is not guaranteed to be differentiable (see Example C.2), we cannot ensure that the objective function is Lipschitz-smooth in general. Thus, unlike previous results for the independent setting that required this latter assumption to achieve faster convergence (e.g., [58]), in our analysis of Algorithm 1, we assume only that the objective function is continuously differentiable, which leads to a more widely applicable, albeit slower, convergence rate. Note, however, that we assume Lipschitz-smoothness in our analysis of Algorithm 2 as it allows for faster convergence to the inner player's optimal strategy, but this assumption could also be done away with again, at the cost of a slower convergence rate.

4 An Economic Application: Fisher Markets

The Fisher market model, attributed to Irving Fisher [10], has received a great deal of attention recently, in particular by computer scientists, as its applications to fair division and mechanism design have proven useful for the design of automated markets in many online marketplaces. In this section, we argue that a competitive equilibrium in Fisher markets can be understood a Stackelberg

⁵A formal statement of this claim can be found in Appendix E (Lemma E.1).

equilibrium of a convex-concave min-max Stackelberg game. We then apply our first-order methods to compute these equilibria in various Fisher markets.

A **Fisher market** consists of n buyers and m divisible goods [10]. Each buyer $i \in [n]$ has a budget $b_i \in \mathbb{R}_+$ and a utility function $u_i : \mathbb{R}_+^m \rightarrow \mathbb{R}$. As is standard in the literature, we assume that there is one divisible unit of each good in the market [57]. An instance of a Fisher market is given by a tuple (n, m, U, \mathbf{b}) , where $U = \{u_1, \dots, u_n\}$ is a set of utility functions, one per buyer, and $\mathbf{b} \in \mathbb{R}_+^n$ is the vector of buyer budgets. We abbreviate as (U, \mathbf{b}) when n and m are clear from context.

An **allocation** $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^T \in \mathbb{R}_+^{n \times m}$ is a map from goods to buyers, represented as a matrix, s.t. $x_{ij} \geq 0$ denotes the amount of good $j \in [m]$ allocated to buyer $i \in [n]$. Goods are assigned **prices** $\mathbf{p} = (p_1, \dots, p_m)^T \in \mathbb{R}_+^m$. A tuple $(\mathbf{p}^*, \mathbf{X}^*)$ is said to be a **competitive (or Walrasian) equilibrium** of Fisher market (U, \mathbf{b}) if 1. buyers are utility maximizing, constrained by their budget, i.e., $\forall i \in [n], \mathbf{x}_i^* \in \arg \max_{\mathbf{x}: \mathbf{x} \cdot \mathbf{p}^* \leq b_i} u_i(\mathbf{x})$; and 2. the market clears, i.e., $\forall j \in [m], p_j^* > 0 \Rightarrow \sum_{i \in [n]} x_{ij}^* = 1$ and $p_j^* = 0 \Rightarrow \sum_{i \in [n]} x_{ij}^* \leq 1$.

We now formulate the problem of computing a competitive equilibrium $(\mathbf{p}^*, \mathbf{X}^*)$ of a Fisher market (U, \mathbf{b}) , where U is a set of continuous, concave, and homogeneous⁶ utility functions, as a convex-concave min-max Stackelberg game, a perspective which has not been taken before. Fisher markets can be solved via the Eisenberg-Gale convex program [24]. Recently, Cole and Tao [13] derived a convex program, which is intimately related to the dual of the Eisenberg-Gale program [31], namely

$$\min_{\mathbf{p} \geq \mathbf{0}} \sum_{j \in [m]} p_j + \sum_{i \in [n]} b_i \log \left(\max_{\mathbf{x}_i \geq \mathbf{0}: \mathbf{x}_i \cdot \mathbf{p} \leq b_i} u_i(\mathbf{x}_i) \right). \quad (2)$$

Rearranging, we obtain the following convex-concave min-max Stackelberg game:

$$\min_{\mathbf{p} \geq \mathbf{0}} \max_{\mathbf{X} \geq \mathbf{0}: \mathbf{X} \mathbf{p} \leq \mathbf{b}} \sum_{j \in [m]} p_j + \sum_{i \in [n]} b_i \log (u_i(\mathbf{x}_i)). \quad (3)$$

This min-max game is played by a fictitious (Walrasian) auctioneer and a set of buyers, who effectively play as a team. The objective function in this game is then the sum of the auctioneer's welfare (i.e., the sum of the prices) and the Nash social welfare of buyers (i.e., the second summation). As the buyer's strategy set is dependent on the price vector \mathbf{p} selected by the auctioneer, we cannot use existing first-order methods to solve this problem. However, we can use Algorithms 1 and 2

Starting from Equation (3), define the auctioneer's value function $V(\mathbf{p}) = \max_{\mathbf{X} \geq \mathbf{0}: \mathbf{X} \mathbf{p} \leq \mathbf{b}} \sum_{j \in [m]} p_j + \sum_{i \in [n]} b_i \log (u_i(\mathbf{x}_i))$, and buyer i 's demand set $X_i^*(\mathbf{p}, \mathbf{b}) = \arg \max_{\mathbf{x}_i \geq \mathbf{0}: \mathbf{x}_i \cdot \mathbf{p} \leq b_i} u_i(\mathbf{x}_i)$. Theorem 3.2 then provides the relevant subgradients so that we can run Algorithms 1 and 2, namely $\partial_{\mathbf{p}} V(\mathbf{p}) = \mathbf{1}_m - \sum_{i \in [n]} X_i^*(\mathbf{p}, \mathbf{b})$ and $\nabla_{\mathbf{x}_i} \left(\sum_{j \in [m]} p_j + \sum_{i \in [n]} b_i \log (u_i(\mathbf{x}_i)) \right) = \frac{b_i}{u_i(\mathbf{x}_i)} \nabla_{\mathbf{x}_i} u_i(\mathbf{x}_i)$, using the Minkowski sum to add set-valued quantities, where $\mathbf{1}_m$ is the vector of ones of size m .⁷

Cheung, Cole, and Devanur [12] observed that solving the dual of the Eisenberg-Gale program (Equation 2) via (sub)gradient descent [20] is equivalent to solving for a competitive equilibrium in a Fisher market using an auction-like economic price adjustment process named tâtonnement that was first proposed by Léon Walras in the 1870s [80]. The tâtonnement process increases the prices of goods that are overdemanded and decreases the prices of goods that are underdemanded. Mathematically, the (vanilla) tâtonnement process [5, 80] is defined as $\mathbf{p}(t) = \max \left\{ \mathbf{p}(t-1) + \eta_t \left(\sum_{i \in [n]} \mathbf{x}_i^*(\mathbf{p}(t), \mathbf{b}) - \mathbf{1} \right), 0 \right\}$ for $\mathbf{p}(0) \in \mathbb{R}_+^m$, where $\mathbf{x}_i^*(\mathbf{p}(t), b_i) \in X_i^*(\mathbf{p}(t), b_i) = \arg \max_{\mathbf{x}_i \geq \mathbf{0}: \mathbf{x}_i \cdot \mathbf{p}(t) \leq b_i} u_i(\mathbf{x}_i)$ is the demand set of buyer i . The max-oracle algorithm applied to Equation (3) is then equivalent to a tâtonnement process where the buyers report a δ -utility maximizing demand. Further, we have the following corollary of Theorem 3.3.

Corollary 4.1. *Let (U, \mathbf{b}) be a Fisher market with equilibrium price vector \mathbf{p}^* , where U is a set of continuous, concave, homogeneous, and continuously differentiable utility functions. Consider*

⁶A function $f : \mathbb{R}^m \rightarrow \mathbb{R}$ is said to be **homogeneous of degree k** if $\forall \mathbf{x} \in \mathbb{R}^m, \lambda > 0, f(\lambda \mathbf{x}) = \lambda^k f(\mathbf{x})$. Unless otherwise indicated, a homogeneous function is assumed to be homogeneous of degree 1.

⁷We include detailed descriptions of the algorithms applied to Fisher markets in Appendix F.

the entropic tâtonnement process [31].⁸ Assume that the step sizes η_t satisfy the usual conditions: $\sum_{k=1}^T \eta_k^2 \leq \infty$ and $\sum_{k=1}^T \eta_k = \infty$. If $\mathbf{p}_{\text{best}}^{(t)} \in \arg \min_{\mathbf{p}^{(k)}: k \in [t]} V(\mathbf{p}^{(k)})$, then $\lim_{k \rightarrow \infty} V(\mathbf{p}_{\text{best}}^{(k)}) = V(\mathbf{p}^*)$. Additionally, tâtonnement converges to an ε -competitive equilibrium in $O(\varepsilon^{-2})$ iterations.

If we also apply the nested gradient-descent-ascent algorithm to Equation (3), we arrive at an algorithm that is arguably more descriptive of market dynamics than tâtonnement itself, as it also includes the demand-side market dynamics of buyers optimizing their demands, potentially in a decentralized manner. The nested tâtonnement algorithm essentially describes a two-step trial-and-error (i.e., tâtonnement) process, where first the buyers try to discover their optimal demand by increasing their demand for goods in proportion to the marginal utility the goods provide, and then the seller/auctioneer adjusts market prices by decreasing the prices of goods that are underdemanded and increasing the prices of goods that are overdemanded. As buyers can calculate their demands in a decentralized fashion, the nested tâtonnement algorithm offers a more complete picture of market dynamics than the classic tâtonnement process.

Experiments In order to better understand, the iteration complexity of Algorithms 1 and 2 (Appendix E), we ran a series of experiments on Fisher markets with three different classes of utility functions.⁹ Each utility structure endows Equation (3) with different smoothness properties, which allows us to compare the efficiency of the algorithms under varying conditions.

Let $\mathbf{v}_i \in \mathbb{R}^m$ be a vector of parameters for the utility function of buyer $i \in [n]$. We have the following utility function classes: Linear: $u_i(\mathbf{x}_i) = \sum_{j \in [m]} v_{ij} x_{ij}$, Cobb-Douglas: $u_i(\mathbf{x}_i) = \prod_{j \in [m]} x_{ij}^{v_{ij}}$, Leontief: $u_i(\mathbf{x}_i) = \min_{j \in [m]} \left\{ \frac{x_{ij}}{v_{ij}} \right\}$. Equation (3) satisfies the smoothness properties listed in Table 2 when U is one of these three classes. Our goals are two-fold. First, we want to understand how the empirical convergence rates of Algorithms 1 and 2 (which, when applied to Equation (3) give rise to Algorithms 3 and 4 in Appendix F, respectively) compare to their theoretical guarantees under different utility structures. Second, we want to understand the extent to which the convergence rates of these two algorithms differ in practice. We include a more detailed description of our experimental setup in Appendix F.

Table 2: Smoothness properties satisfied by Equation (3) assuming different utility functions. Note that Assumption 3.1 does not hold for Leontief utilities, because they are not differentiable.

	V is differentiable	Assumption 3.1 holds
Linear	×	✓
Cobb-Douglas	✓	✓
Leontief	✓	×

Figure 1 describes the empirical convergence rates of Algorithms 1 and 2 for linear, Cobb-Douglas, and Leontief utilities. We observe that convergence is fastest in Fisher markets with Cobb-Douglas utilities, followed by linear, and then Leontief. We seem to obtain a tight convergence rate of $O(1/\sqrt{T})$ for linear utilities, which seems plausible, as the value function is not differentiable assuming linear utilities, and hence we are unlikely to achieve a better convergence rate. On the other hand, for Cobb-Douglas utilities, both the value and the objective function are differentiable; in fact, they are both twice continuously differentiable, making them both Lipschitz-smooth. These factors combined seem to provide a much faster convergence rate than $O(1/\sqrt{T})$.

Fisher markets with Leontief utilities, in which the objective function is not differentiable, are the hardest markets of the three for our algorithms to solve. Indeed, our theory does not even predict convergence. Still, convergence is not entirely surprising, as Cheung, Cole, and Devanur [12] have shown that buyer demand throughout tâtonnement is bounded for Leontief utilities, which means that the objective function of Equation (3) is locally Lipschitz: i.e., any subgradient computed by the algorithm will be bounded. Overall, our theory suggests that differentiability of the value function is not essential to guarantee convergence of first-order methods in convex-concave games, while our

⁸The entropic tâtonnement process ensures that prices remain bounded away from zero, so that excess demand is likewise bounded and Assumption 3.1 is satisfied.

⁹Our code can be found at <https://github.com/denizalp/min-max-fisher.git>

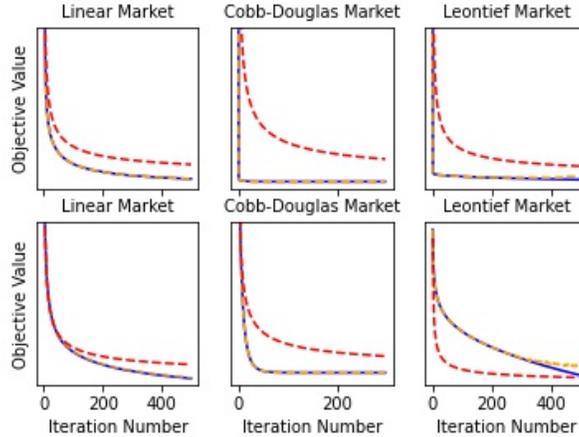


Figure 1: The first row describes the average trajectory of the value of the objective function for a randomly initialized market on each iteration of both Algorithm 3 (in blue) and Algorithm 4 (in orange) when the starting prices are high, while the second row describes the average trajectory of the objective function when starting prices are low for linear, Cobb-Douglas, and Leontief Fisher markets respectively. The dashed red line represents a convergence rate of $O(1/\sqrt{T})$, which corresponds to an iteration complexity of $O(1/\epsilon^2)$.

experiments seem to suggest that differentiability of the objective function is more important than differentiability of the value function in regards to the convergence rate.

In order to investigate whether the outputs of Algorithm 1, which uses an exact max-oracle (i.e., 0-max-oracle) are more precise than those of Algorithm 2, we solved 500 randomly initialized markets with both algorithms. We then ran a James' first-order test [2, 36] on the mean output of both algorithms to see if their difference was statistically significant. Our calculations produced p -values of 0.69, 0, and 1.06×10^{-18} , for Fisher markets with linear, Cobb-Douglas, and Leontief utilities, respectively. At a significance level of 0.05, these results are *not* statistically significant for linear utilities only. This result can be attributed to the fact that the value function is not differentiable in the linear case, which makes the nested gradient descent/ascent algorithm less precise.

5 Conclusion

In this paper, we proposed two first-order methods to solve a class of constrained convex-concave min-max optimization problems, which we call min-max Stackelberg games. As such games do not afford Nash equilibria in general, we focused on the computation of their Stackelberg equilibria. To solve for Stackelberg equilibria, we introduced a novel subdifferential envelope theorem, which formed the heart of two subgradient methods with polynomial-time iteration complexity that converge to Stackelberg equilibria. Finally, we applied our theory to the computation of competitive equilibria in Fisher markets. This application yielded a new variant of the classic tâtonnement process, where, in addition to the auctioneer iteratively adjusting prices, the buyers iteratively compute their demands. A further question of interest both for Fisher market dynamics and convex-concave min-max Stackelberg games more generally is whether gradient-descent-ascent (GDA) converges in the dependent strategy set setting as it does in the independent strategy setting [45]. GDA dynamics for Fisher markets correspond to myopic best-response dynamics (see, for example, [52]). We would expect such a result to be of computational as well as economic interest.

Acknowledgments and Disclosure of Funding

We would like to thank Dustin Morrill for his feedback on an early draft of this paper. This work was partially supported by NSF Grant CMMI-1761546.

References

- [1] S. N. Afriat. “Theory of Maxima and the Method of Lagrange”. In: *SIAM Journal on Applied Mathematics* 20.3 (1971), pp. 343–357. ISSN: 00361399. URL: <http://www.jstor.org/stable/2099955>.
- [2] James Algina, T. C. Oshima, and Wen-Ying Lin. “Type I Error Rates for Welch’s Test and James’s Second-Order Test under Nonnormality and Inequality of Variance When There are Two Groups”. In: *Journal of Educational and Behavioral Statistics* 19.3 (1994), pp. 275–291. ISSN: 10769986, 19351054. URL: <http://www.jstor.org/stable/1165297>.
- [3] Mohammad Alkousa et al. *Accelerated methods for composite non-bilinear saddle point problem*. 2020. arXiv: [1906.03620 \[math.OC\]](https://arxiv.org/abs/1906.03620).
- [4] Kenneth Arrow and Gerard Debreu. “Existence of an equilibrium for a competitive economy”. In: *Econometrica: Journal of the Econometric Society* (1954), pp. 265–290.
- [5] Kenneth J. Arrow and Leonid Hurwicz. “On the Stability of the Competitive Equilibrium, I”. In: *Econometrica* 26.4 (1958), pp. 522–552. ISSN: 00129682, 14680262. URL: <http://www.jstor.org/stable/1907515>.
- [6] Aharon Ben-Tal et al. “Oracle-based robust optimization via online learning”. In: *Operations Research* 63.3 (2015), pp. 628–638.
- [7] Claude Berge. *Topological Spaces: including a treatment of multi-valued functions, vector spaces, and convexity*. Courier Corporation, 1997.
- [8] Stephen Boyd, Stephen P Boyd, and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [9] Stephen Boyd and Lieven Vandenberghe. *Subgradients*. Apr. 2018.
- [10] William C Brainard, Herbert E Scarf, et al. *How to compute equilibrium prices in 1891*. Citeseer, 2000.
- [11] Qi Cai et al. *On the Global Convergence of Imitation Learning: A Case for Linear Quadratic Regulator*. 2019. arXiv: [1901.03674 \[cs.LG\]](https://arxiv.org/abs/1901.03674).
- [12] Yun Kuen Cheung, Richard Cole, and Nikhil Devanur. “Tatonnement beyond Gross Substitutes? Gradient Descent to the Rescue”. In: *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing*. STOC ’13. Palo Alto, California, USA: Association for Computing Machinery, 2013, pp. 191–200. ISBN: 9781450320290. DOI: [10.1145/2488608.2488633](https://doi.org/10.1145/2488608.2488633), URL: <https://doi.org/10.1145/2488608.2488633>.
- [13] Richard Cole and Yixin Tao. *Balancing the Robustness and Convergence of Tatonnement*. 2019. arXiv: [1908.00844 \[cs.GT\]](https://arxiv.org/abs/1908.00844).
- [14] Bo Dai et al. “Kernel Exponential Family Estimation via Doubly Dual Embedding”. In: *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*. Ed. by Kamalika Chaudhuri and Masashi Sugiyama. Vol. 89. Proceedings of Machine Learning Research. PMLR, Apr. 2019, pp. 2321–2330. URL: <http://proceedings.mlr.press/v89/dai19a.html>.
- [15] Bo Dai et al. “SBEED: Convergent Reinforcement Learning with Nonlinear Function Approximation”. In: *Proceedings of the 35th International Conference on Machine Learning*. Ed. by Jennifer Dy and Andreas Krause. Vol. 80. Proceedings of Machine Learning Research. PMLR, July 2018, pp. 1125–1134. URL: <http://proceedings.mlr.press/v80/dai18c.html>.
- [16] Yu-Hong Dai and Liwei Zhang. *Optimality Conditions for Constrained Minimax Optimization*. 2020. arXiv: [2004.09730 \[math.OC\]](https://arxiv.org/abs/2004.09730).
- [17] John M. Danskin. “The Theory of Max-Min, with Applications”. In: *SIAM Journal on Applied Mathematics* 14.4 (1966), pp. 641–664. ISSN: 00361399. URL: <http://www.jstor.org/stable/2946123>.
- [18] Constantinos Daskalakis, Stratis Skoulakis, and Manolis Zampetakis. *The Complexity of Constrained Min-Max Optimization*. 2020. arXiv: [2009.09623 \[cs.CC\]](https://arxiv.org/abs/2009.09623).
- [19] Trevor Davis, Kevin Waugh, and Michael Bowling. *Solving Large Extensive-Form Games with Strategy Constraints*. 2019. arXiv: [1809.07893 \[cs.GT\]](https://arxiv.org/abs/1809.07893).
- [20] N. R. Devanur et al. “Market equilibrium via a primal-dual-type algorithm”. In: *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings*. 2002, pp. 389–395. DOI: [10.1109/SFCS.2002.1181963](https://doi.org/10.1109/SFCS.2002.1181963).

- [21] Steven Diamond and Stephen Boyd. “CVXPY: A Python-embedded modeling language for convex optimization”. In: *Journal of Machine Learning Research* 17.83 (2016), pp. 1–5.
- [22] Paul Duetting et al. “Optimal auctions through deep learning”. In: *International Conference on Machine Learning*. PMLR, 2019, pp. 1706–1715.
- [23] Harrison Edwards and Amos Storkey. *Censoring Representations with an Adversary*. 2016. arXiv: [1511.05897 \[cs.LG\]](https://arxiv.org/abs/1511.05897).
- [24] Edmund Eisenberg and David Gale. “Consensus of subjective probabilities: The pari-mutuel method”. In: *The Annals of Mathematical Statistics* 30.1 (1959), pp. 165–168.
- [25] Filippo Fabiani et al. “Local Stackelberg equilibrium seeking in generalized aggregative games”. In: *IEEE Transactions on Automatic Control* (2021).
- [26] Francisco Facchinei and Christian Kanzow. “Generalized Nash equilibrium problems”. In: *4or* 5.3 (2007), pp. 173–210.
- [27] Francisco Facchinei and Christian Kanzow. “Generalized Nash equilibrium problems”. In: *Annals of Operations Research* 175.1 (2010), pp. 177–211.
- [28] Gabriele Farina, Christian Kroer, and Tuomas Sandholm. *Regret Circuits: Composability of Regret Minimizers*. 2019. arXiv: [1811.02540 \[cs.LG\]](https://arxiv.org/abs/1811.02540).
- [29] Yuan Gao and Christian Kroer. “First-Order Methods for Large-Scale Market Equilibrium Computation”. In: *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*. Ed. by Hugo Larochelle et al. 2020. URL: <https://proceedings.neurips.cc/paper/2020/hash/f75526659f31040afeb61cb7133e4e6d-Abstract.html>.
- [30] Gauthier Gidel et al. *A Variational Inequality Perspective on Generative Adversarial Networks*. 2020. arXiv: [1802.10551 \[cs.LG\]](https://arxiv.org/abs/1802.10551).
- [31] Denizalp Goktas, Enrique Areyan Viqueira, and Amy Greenwald. “A Consumer-Theoretic Characterization of Fisher Market Equilibria”. In: *Proceedings of the Conference on Web and Internet Economics, (WINE)*. 2021.
- [32] E. Yazdandoost Hamedani et al. *Iteration Complexity of Randomized Primal-Dual Methods for Convex-Concave Saddle Point Problems*. 2018. arXiv: [1806.04118 \[math.OA\]](https://arxiv.org/abs/1806.04118).
- [33] Erfan Yazdandoost Hamedani and Necdet Serhat Aybat. “A primal-dual algorithm for general convex-concave saddle point problems”. In: *arXiv preprint arXiv:1803.01401* 2 (2018).
- [34] Charles R. Harris et al. “Array programming with NumPy”. In: *Nature* 585 (2020), pp. 357–362. DOI: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2).
- [35] R Henrion. “On constraint qualifications”. In: *Journal of optimization theory and applications* 72.1 (1992), pp. 187–197.
- [36] Freddy Hernandez et al. *stests: Package with useful statistical tests*. R package version 0.1.0. 2021. URL: <https://fhernanb.github.io/stests>.
- [37] J. D. Hunter. “Matplotlib: A 2D graphics environment”. In: *Computing in Science & Engineering* 9.3 (2007), pp. 90–95. DOI: [10.1109/MCSE.2007.55](https://doi.org/10.1109/MCSE.2007.55).
- [38] Adam Ibrahim et al. “Lower bounds and conditioning of differentiable games”. In: *arXiv preprint arXiv:1906.07300* (2019).
- [39] Tatsuuro Ichiishi. “4 - Noncooperative Behavior and Equilibrium”. In: *Game Theory for Economic Analysis*. Ed. by Tatsuuro Ichiishi. Economic Theory, Econometrics, and Mathematical Economics. San Diego: Academic Press, 1983, pp. 55–76. ISBN: 978-0-12-370180-0. DOI: <https://doi.org/10.1016/B978-0-12-370180-0.50009-3>, URL: <https://www.sciencedirect.com/science/article/pii/B9780123701800500093>.
- [40] Kamal Jain, Vijay V Vazirani, and Yinyu Ye. “Market equilibria for homothetic, quasi-concave utilities and economies of scale in production”. In: *SODA*. Vol. 5. 2005, pp. 63–71.
- [41] Chi Jin, Praneeth Netrapalli, and Michael I. Jordan. *What is Local Optimality in Nonconvex-Nonconcave Minimax Optimization?* 2020. arXiv: [1902.00618 \[cs.LG\]](https://arxiv.org/abs/1902.00618).
- [42] Anatoli Juditsky, Arkadi Nemirovski, et al. “First order methods for nonsmooth convex large-scale optimization, ii: utilizing problems structure”. In: *Optimization for Machine Learning* 30.9 (2011), pp. 149–183.
- [43] Hamed Kebriaei and Luigi Iannelli. “Discrete-time robust hierarchical linear-quadratic dynamic games”. In: *IEEE Transactions on Automatic Control* 63.3 (2017), pp. 902–909.

- [44] HW Kuhn and AW Tucker. *Proceedings of 2nd berkeley symposium*. 1951.
- [45] Tianyi Lin, Chi Jin, and Michael Jordan. “On gradient descent ascent for nonconvex-concave minimax problems”. In: *International Conference on Machine Learning*. PMLR. 2020, pp. 6083–6093.
- [46] Tianyi Lin, Chi Jin, and Michael I Jordan. “Near-optimal algorithms for minimax optimization”. In: *Conference on Learning Theory*. PMLR. 2020, pp. 2738–2779.
- [47] Songtao Lu, Ioannis Tsaknakis, and Mingyi Hong. “Block alternating optimization for non-convex min-max problems: algorithms and applications in signal processing and communications”. In: *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE. 2019, pp. 4754–4758.
- [48] David Madras et al. *Learning Adversarially Fair and Transferable Representations*. 2018. arXiv: [1802.06309 \[cs.LG\]](https://arxiv.org/abs/1802.06309).
- [49] Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. OUP Catalogue 9780195102680. Oxford University Press, 1995. ISBN: ARRAY(0x4cf9c5c0). URL: <https://ideas.repec.org/b/oxp/obooks/9780195102680.html>
- [50] Paul Milgrom and Ilya Segal. “Envelope theorems for arbitrary choice sets”. In: *Econometrica* 70.2 (2002), pp. 583–601.
- [51] Aryan Mokhtari, Asuman Ozdaglar, and Sarath Pattathil. *Convergence Rate of $\mathcal{O}(1/k)$ for Optimistic Gradient and Extra-gradient Methods in Smooth Convex-Concave Saddle Point Problems*. 2020. arXiv: [1906.01115 \[math.OA\]](https://arxiv.org/abs/1906.01115).
- [52] Dov Monderer and Lloyd S Shapley. “Potential games”. In: *Games and economic behavior* 14.1 (1996), pp. 124–143.
- [53] Arkadi Nemirovski. “Prox-method with rate of convergence $\mathcal{O}(1/t)$ for variational inequalities with Lipschitz continuous monotone operators and smooth convex-concave saddle point problems”. In: *SIAM Journal on Optimization* 15.1 (2004), pp. 229–251.
- [54] Yurii Nesterov. “Dual extrapolation and its applications to solving variational inequalities and related problems”. In: *Mathematical Programming* 109.2 (2007), pp. 319–344.
- [55] J v Neumann. “Zur theorie der gesellschaftsspiele”. In: *Mathematische annalen* 100.1 (1928), pp. 295–320.
- [56] Nam Ho-Nguyen and Fatma Kilinc-Karzan. “Online first-order framework for robust convex optimization”. In: *Operations Research* 66.6 (2018), pp. 1670–1692.
- [57] Noam Nissan and Tim Roughgarden. *Algorithmic Game Theory*. Cambridge University Press, 2007. DOI: [10.1017/CBO9780511800481](https://doi.org/10.1017/CBO9780511800481)
- [58] Maher Nouiehed et al. “Solving a class of non-convex min-max games using iterative first order methods”. In: *arXiv preprint arXiv:1902.08297* (2019).
- [59] Dmitrii M Ostrovskii, Andrew Lowy, and Meisam Razaviyayn. “Efficient search of first-order nash equilibria in nonconvex-concave smooth min-max problems”. In: *arXiv preprint arXiv:2002.07919* (2020).
- [60] Yuyuan Ouyang and Yangyang Xu. *Lower complexity bounds of first-order methods for convex-concave bilinear saddle-point problems*. 2018. arXiv: [1808.02901 \[math.OA\]](https://arxiv.org/abs/1808.02901).
- [61] Krzysztof Postek and Shimrit Shtern. *First-order algorithms for robust optimization problems via convex-concave saddle-point Lagrangian reformulation*. 2021. arXiv: [2101.02669 \[math.OA\]](https://arxiv.org/abs/2101.02669).
- [62] Murray H Protter, B Charles Jr, et al. *A first course in real analysis*. Springer Science & Business Media, 2012.
- [63] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. Vienna, Austria, 2013. URL: <http://www.R-project.org/>.
- [64] Hassan Rafique et al. *Non-Convex Min-Max Optimization: Provable Algorithms and Applications in Machine Learning*. 2019. arXiv: [1810.02060 \[math.OA\]](https://arxiv.org/abs/1810.02060).
- [65] R Tyrrell Rockafellar and Roger J-B Wets. *Variational analysis*. Vol. 317. Springer Science & Business Media, 2009.
- [66] Maziar Sanjabi et al. *On the Convergence and Robustness of Training GANs with Regularized Optimal Transport*. 2018. arXiv: [1802.08249 \[cs.LG\]](https://arxiv.org/abs/1802.08249).

- [67] Maziar Sanjabi et al. “On the Convergence and Robustness of Training GANs with Regularized Optimal Transport”. In: *Proceedings of the 32nd International Conference on Neural Information Processing Systems*. NIPS’18. Montréal, Canada: Curran Associates Inc., 2018, pp. 7091–7101.
- [68] Prasanna Sattigeri et al. *Fairness GAN*. 2018. arXiv: [1805.09910](https://arxiv.org/abs/1805.09910) [stat.ML].
- [69] Gilberto A.S. Segundo, Renato A. Krohling, and Rodrigo C. Cosme. “A differential evolution approach for solving constrained min–max optimization problems”. In: *Expert Systems with Applications* 39.18 (2012), pp. 13440–13450. ISSN: 0957-4174. DOI: <https://doi.org/10.1016/j.eswa.2012.05.059>. URL: <https://www.sciencedirect.com/science/article/pii/S0957417412007750>.
- [70] K. Shimizu and E. Aiyoshi. “A new computational method for Stackelberg and min-max problems by use of a penalty method”. In: *IEEE Transactions on Automatic Control* 26.2 (1981), pp. 460–466. DOI: [10.1109/TAC.1981.1102607](https://doi.org/10.1109/TAC.1981.1102607).
- [71] K. Shimizu and E. Aiyoshi. “Necessary conditions for min-max problems and algorithms by a relaxation procedure”. In: *IEEE Transactions on Automatic Control* 25.1 (1980), pp. 62–66. DOI: [10.1109/TAC.1980.1102226](https://doi.org/10.1109/TAC.1980.1102226).
- [72] Aman Sinha et al. *Certifying Some Distributional Robustness with Principled Adversarial Training*. 2020. arXiv: [1710.10571](https://arxiv.org/abs/1710.10571) [stat.ML].
- [73] The pandas development team. *pandas-dev/pandas: Pandas*. Version latest. Feb. 2020. DOI: [10.5281/zenodo.3509134](https://doi.org/10.5281/zenodo.3509134). URL: <https://doi.org/10.5281/zenodo.3509134>.
- [74] Kiran Koshy Thekumparampil et al. *Efficient Algorithms for Smooth Minimax Optimization*. 2019. arXiv: [1907.01543](https://arxiv.org/abs/1907.01543) [math.OA].
- [75] Paul Tseng. “On accelerated proximal gradient methods for convex-concave optimization”. In: *submitted to SIAM Journal on Optimization* 1 (2008).
- [76] Paul Tseng. “On linear convergence of iterative methods for the variational inequality problem”. In: *Journal of Computational and Applied Mathematics* 60.1 (1995). Proceedings of the International Meeting on Linear/Nonlinear Iterative Methods and Verification of Solution, pp. 237–252. ISSN: 0377-0427. DOI: [https://doi.org/10.1016/0377-0427\(94\)00094-H](https://doi.org/10.1016/0377-0427(94)00094-H). URL: <https://www.sciencedirect.com/science/article/pii/S037704279400094H>.
- [77] Guido Van Rossum and Fred L Drake Jr. *Python tutorial*. Centrum voor Wiskunde en Informatica Amsterdam, The Netherlands, 1995.
- [78] Heinrich Von Stackelberg. *Marktform und gleichgewicht*. J. springer, 1934.
- [79] Abraham Wald. “Statistical Decision Functions Which Minimize the Maximum Risk”. In: *Annals of Mathematics* 46.2 (1945), pp. 265–280. ISSN: 0003486X. URL: <http://www.jstor.org/stable/1969022>.
- [80] Léon Walras. *Elements of Pure Economics; or, The Theory of Social Wealth*. Translated by William Jaffé. Vol. 2. Orion Editions, 1969, pp. 457–458.
- [81] Hadley Wickham et al. “Welcome to the tidyverse”. In: *Journal of Open Source Software* 4.43 (2019), p. 1686. DOI: [10.21105/joss.01686](https://doi.org/10.21105/joss.01686).
- [82] Depeng Xu et al. *FairGAN: Fairness-aware Generative Adversarial Networks*. 2018. arXiv: [1805.11202](https://arxiv.org/abs/1805.11202) [cs.LG].
- [83] Laura Scrimali Yurii Nesterov. “Solving strongly monotone variational and quasi-variational inequalities”. In: *Discrete & Continuous Dynamical Systems* 31.4 (2011), pp. 1383–1396.
- [84] Junyu Zhang, Mingyi Hong, and Shuzhong Zhang. *On Lower Iteration Complexity Bounds for the Saddle Point Problems*. 2020. arXiv: [1912.07481](https://arxiv.org/abs/1912.07481) [math.OA].
- [85] Renbo Zhao. *A Primal Dual Smoothing Framework for Max-Structured Nonconvex Optimization*. 2020. arXiv: [2003.04375](https://arxiv.org/abs/2003.04375) [math.OA].
- [86] Renbo Zhao. “Optimal algorithms for stochastic three-composite convex-concave saddle point problems”. In: *arXiv preprint arXiv:1903.01687* (2019).