
Learning Optimal Tax Design in Nonatomic Congestion Games

Qiwen Cui

Paul G. Allen School of Computer Science
Engineering
University of Washington
Seattle, WA 98195
qwcui@cs.washington.edu

Maryam Fazel

Department of Electrical
Computer Engineering
University of Washington
Seattle, WA 98195
mfazel@uw.edu

Simon S. Du

Paul G. Allen School of Computer Science
Engineering
University of Washington
Seattle, WA 98195
ssdu@cs.washington.edu

Abstract

In multiplayer games, self-interested behavior among the players can harm the social welfare. Tax mechanisms are a common method to alleviate this issue and induce socially optimal behavior. In this work, we take the initial step of learning the optimal tax that can maximize social welfare with limited feedback in congestion games. We propose a new type of feedback named *equilibrium feedback*, where the tax designer can only observe the Nash equilibrium after deploying a tax plan. Existing algorithms are not applicable due to the exponentially large tax function space, nonexistence of the gradient, and nonconvexity of the objective. To tackle these challenges, we design a computationally efficient algorithm that leverages several novel components: (1) a piece-wise linear tax to approximate the optimal tax; (2) extra linear terms to guarantee a strongly convex potential function; (3) an efficient subroutine to find the exploratory tax that can provide critical information about the game. The algorithm can find an ϵ -optimal tax with $O(\beta F^2/\epsilon)$ sample complexity, where β is the smoothness of the cost function and F is the number of facilities.

1 Introduction

In modern society, large-scale systems often consist of many self-interested players with shared resources, such as transportation and communication networks. Importantly, the objectives of individual players are not always aligned with the system efficiency, and the system designer should take this into consideration. A widely known example is Braess's paradox, where adding more roads to a network can make the network more congested [Braess, 1968]. Price of anarchy is a notion that measures the inefficiency caused by selfish behavior compared with optimal centralized behavior [Koutsoupias and Papadimitriou, 1999]. Characterizing such inefficiency has been an active research area with applications in resource allocation [Marden and Roughgarden, 2014], traffic congestion [Roughgarden and Tardos, 2004], and others. The inefficiency motivates research on how to design mechanisms to improve performance even when the players are still behaving selfishly.

Tax mechanisms are a standard approach to resolving the inefficiency issue, which are widely studied in economics, operations research, and game theory. The goal of tax mechanisms is to incentivize self-interested players to follow socially optimal behavior by applying tax/subsidy. Congestion game is a widely studied class of game theory models characterizing the interactions between players sharing facilities, where the cost of each facility depends on the “congestion” level [Wardrop [1952], Rosenthal [1973]]. As a motivating example, in traffic routing games, each facility corresponds to an edge in a network, and each player chooses a path that connects her source node and target node. The cost of each facility corresponds to the latency of each edge, which depends on the number of players using that edge. Then, the tax can be interpreted as the toll collected by the road owner or the government to improve overall traffic efficiency [Bergendorff et al. [1997]].

Most existing works on congestion game tax design focus on the computation complexity of the optimal tax [Nisan et al., 2007], [Caragiannis et al., 2010]. They assume the tax designer has full knowledge of the underlying game, which is unrealistic in many applications. As Nash equilibrium is the only stable state of the system, we study a partial information feedback setting named “equilibrium feedback”, where the tax designer can only observe information about the Nash equilibrium. The limited feedback information brings new challenges to the tax designer, and strategic exploration is necessary to learn or design the optimal tax. In this work, we aim to take the first step in learning optimal tax design for congestion games, and we study the following problem:

How can we learn the optimal tax design in congestion games with equilibrium feedback?

Below we highlight our contributions.

1.1 Main Contributions and Technical Novelties

1. The first algorithm for learning optimal tax design in congestion games. To the best of our knowledge, this is the first result for learning optimal tax in congestion games with partial information feedback. Our algorithm enjoys $O(F^2\beta/\epsilon)$ sample complexity for learning an ϵ -optimal tax, where F is the number of facilities and β is the smoothness coefficient of the cost function. The sample complexity has no dependence on the number of actions, which could be exponential in F . In addition, we provide an efficient implementation for network congestion games with $\tilde{O}(\text{poly}(V, E, \epsilon))$ computational complexity, where V and E are the numbers of the vertexes and edges in the network. Due to space limitation, we defer the computation analysis and experiments to Appendix C and Appendix E.

2. Piece-wise linear function approximation. We only assume the cost functions are smooth and make no parameterization assumptions as they are too strong to be satisfied in real-world applications. To tackle this challenge, we use piece-wise linear functions to approximate the optimal tax function. While only the values of the cost functions can be observed, we show that a carefully designed piece-wise linear function can approximate the unobservable optimal tax function well.

3. Strongly convex potential function. One challenge in tax design is controlling the sensitivity of Nash equilibrium w.r.t. tax perturbation. We always enforce tax functions with subgradient lower bounded by some positive value, which leads to a strongly convex potential function. As a result, the Nash equilibrium will be unique and Lipschitz with respect to tax perturbation. As the potential function for optimal tax is not necessarily strongly convex, we carefully choose the strong-convexity coefficient to balance the induced bias.

4. Exploratory tax design. Given the equilibrium feedback, the tax designer can only indirectly query the cost function by applying tax. Consequently, exploration in tax design becomes much more difficult than that in standard bandit problems where the player can directly query the value of an action [Lattimore and Szepesvári [2020]]. We design an exploratory tax that pushes the equilibrium to the “boundary”, where an additional tax perturbation will change the equilibrium and reveal information about at least one unknown facility.

In this work, we focus on the well-known nonatomic congestion games. We hope our algorithm and analysis provide new insight on the intriguing structure of nonatomic congestion games. In addition, the tax design algorithm might find applications in real-world problems such as toll design in traffic networks. Due to space limitation, most proofs are deferred to the appendix.

Notations. $[m] = \{1, 2, \dots, m\}$. For a set of real numbers \mathcal{K} and a real number $x : \min\{\mathcal{K}\} \leq x \leq \max\{\mathcal{K}\}$, we define $[x]_{\mathcal{K}}^+ := \min_{y \in \mathcal{K}: y \geq x} y$ and $[x]_{\mathcal{K}}^- := \max_{y \in \mathcal{K}: y \leq x} y$. The clip operation $\text{clip}(a, l, r) := \min\{\max\{a, l\}, r\}$ clips a into the interval $[l, r]$. We use $O(\cdot)$ to hide absolute constants and $\tilde{O}(\cdot)$ to hide polylog terms as well. A function $f : \mathcal{X} \mapsto \mathbb{R}$ is α -strongly convex if $f(y) \geq f(x) + \nabla f(x)^\top (y - x) + \frac{\alpha}{2} \|y - x\|_2^2, \forall x, y \in \mathcal{X}$. f is β -smooth if $\|\nabla f(x) - \nabla f(y)\|_2 \leq \beta \|x - y\|_2, \forall x, y \in \mathcal{X}$.

2 Related Work

Learning in congestion games. We refer the readers to the textbook [Nisan et al., 2007] for a general introduction to congestion games, the price of anarchy and tax mechanisms. Nonatomic congestion games were first studied in [Pigou, 1912] and formalized by [Wardrop, 1952]. Atomic congestion games were introduced by [Rosenthal, 1973] and the connection with potential games is developed by [Monderer and Shapley, 1996]. In contrast to general-sum games without structures, (approximate) Nash equilibrium can be computed efficiently in congestion games due to the existence of the potential function. Recently, various algorithms are developed to learn the Nash equilibrium in congestion games with different feedback oracles [Krichene et al., 2015, Chen and Lu, 2016, Cui et al., 2022, Jiang et al., 2022, Panageas et al., 2023, Dong et al., 2023, Dadi et al., 2024]. These algorithms are derived from the perspective of the players in the system, while our algorithm is essentially different in that it is utilized by the system designer to induce better equilibrium.

Optimal tax design in congestion games. For nonatomic congestion games, optimal tax design has a closed-form solution known as the marginal cost mechanism [Nisan et al., 2007]. For atomic congestion games, the marginal cost mechanism can no longer improve the efficiency [Paccagnan et al., 2021]. Instead, other mechanisms are proposed for optimal local/global and congestion dependent/independent tax in atomic congestion games [Caragiannis et al., 2010, Bilò and Vinci, 2019, Paccagnan et al., 2021, Paccagnan and Gairing, 2021, Harks et al., 2015]. Notably, all of these mechanisms assume full knowledge of the game while we consider learning with partial information feedback.

Stackelberg games. Stackelberg game [Von Stackelberg, 2010] models the interactions between leaders and followers such that leaders take actions first and the followers make decisions after observing leaders' actions. Tax design can be formulated as a Stackelberg game where the designer is the leader and the game players are the followers. Equipped with a best response oracle to predict followers' actions, [Letchford et al., 2009], [Blum et al., 2014], [Peng et al., 2019] propose algorithms for learning Stackelberg equilibrium. Recently, [Bai et al., 2021], [Zhong et al., 2021], [Zhao et al., 2023] generalize these results to learning Stackelberg equilibrium with bandit feedback, under finite actions or linear function approximation assumptions. For tax design, the search space is an exponentially large function space with complicated dependence on the objective. Consequently, existing results for Stackelberg games become vacuous when specialized to our problem.

Mathematical programming under equilibrium constraint. Tax design can be formulated as minimizing social cost with respect to tax under the constraint that players are following the equilibrium. This is known as mathematical programs with equilibrium constraints (MPEC). MPEC is a bilevel optimization problem and is NP-hard in general [Luo et al., 1996]. Existing approaches use specific inner loop algorithms to approach the equilibrium so that the gradient can be propagated to the outer loop [Li et al., 2020, Liu et al., 2022, Li et al., 2022, Maheshwari et al., 2023, Li et al., 2023, Grontas et al., 2024], relying on a unique and differentiable equilibrium [Colson et al., 2007]. However, such an approach requires many strong assumptions, such as the tax designer can control the algorithm of the agents, convex objective function and parameterized tax function. In contrast, our results make none of these assumptions.

3 Preliminaries

Nonatomic congestion games. A weighted nonatomic congestion game (congestion game) is described by the tuple $(\mathcal{F}, \mathcal{A}_{[m]}, w_{[m]}, c_{\mathcal{F}})$, where \mathcal{F} is the set of facilities with cardinality F , m is the number of commodities, \mathcal{A}_i is the action set for commodity $i \in [m]$, $w_i \in [0, 1]$ is the weight for

commodity $i \in [m]$ such that $\sum_{i \in [m]} w_i = 1$, and $c_f : [0, 1] \mapsto [0, 1]$ is the cost function for facility $f \in \mathcal{F}$. Each commodity consists of infinite number of infinitesimal players with a total load to be w_i . Each individual player is self-interested and has a negligible effect on the game.

In congestion games, action $a \in \mathcal{A}_i$, $i \in [m]$ is a subset of \mathcal{F} , i.e. $a \subseteq \mathcal{F}$, which denotes the facilities utilized by action a . For commodity $i \in [m]$, we use strategy $x_i = (x_{i,a})_{a \in \mathcal{A}_i} \in [0, w_i]^{|\mathcal{A}_i|}$ with constraint $\sum_{a \in \mathcal{A}_i} x_{i,a} = w_i$ to denote how the load is distributed over all the actions. The joint strategy for the game is represented by $x = (x_1, x_2, \dots, x_m) \in [0, 1]^A$, where $A = \sum_{i \in [m]} |\mathcal{A}_i|$. We use \mathcal{X} to denote the set of all feasible strategies.

A decentralized perspective of strategy x_i for commodity i is that each self-interested infinitesimal player follows a randomized strategy that chooses $a \in \mathcal{A}_i$ with probability proportional to $x_{i,a}$. With the law of large number, the load on action a would be $x_{i,a}$.

Cost function. For a strategy x , the cost of a facility is $c_f(l_f(x))$, where $l_f(x) = \sum_{i \in [m]} \sum_{a \in \mathcal{A}_i : f \in a} x_{i,a}$ is the load on facility f . The cost of an action a is the sum of the facility cost that a utilizes: $c_a(x) := \sum_{f \in a} c_f(l_f(x))$.

We make the following assumption on the cost function. Monotonicity is a standard congestion game assumption, which is also observed in many real-world applications as more players sharing one facility, each player will have less gain or more cost [Nisan et al., 2007]. Smoothness is a standard technical assumption for analysis.

Assumption 1. *We assume the cost function satisfies:*

1. *Monotonicity: $c_f(\cdot)$ is non-decreasing for all $f \in \mathcal{F}$,*
2. *Smoothness: $c_f(\cdot)$ is β -smooth for all $f \in \mathcal{F}$.*

Nash equilibrium. Nash equilibrium in nonatomic congestion games, also known as the Wardrop equilibrium [Wardrop, 1952], is the strategy that no player has the incentive to deviate from its strategy as formalized in Definition 1. In other words, Nash equilibrium is a stable state for a system with selfish players.

Definition 1. *A Nash equilibrium strategy x is a joint strategy such that each player is choosing the best action: for any commodity $i \in [m]$ and actions $a, a' \in \mathcal{A}_i$, we have*

$$c_a(x) \leq c_{a'}(x), \text{ if } x_{i,a} > 0.$$

Similarly, an ϵ -approximate Nash equilibrium x satisfies that

$$\forall i \in [m], a, a' \in \mathcal{A}_i, c_a(x) \leq c_{a'}(x) + \epsilon, \text{ if } x_{i,a} > 0.$$

For a strategy x and commodity i , actions $a \in \mathcal{A}_i$ such that $x_{i,a} > 0$ are named as the “in-support” actions and the others are “off-support” actions. For a Nash equilibrium, in-support actions must all have the same cost and off-support actions are no better than in-support actions. It is well known that Nash equilibrium always exists in congestion games [Beckmann et al., 1956].

Potential Function. An important concept in congestion games is the potential function:

$$\Phi(x) := \sum_f \int_0^{l_f(x)} c_f(u) du.$$

If Assumption 1 is satisfied, then $\Phi(x)$ is a convex function and Nash equilibrium is equivalent to the minimizer of the potential function [Beckmann et al., 1956].

Network congestion games. Network congestion games are congestion games with multicommodity network structure, which are also known as the selfish routing games [Roughgarden, 2005]. A multicommodity network is described by a directed graph $(\mathcal{V}, \mathcal{E})$ where \mathcal{V} is the vertex set and \mathcal{E} is the edge (facility) set. In addition, each commodity $i \in [m]$ corresponds to a pair of source and target vertex (s_i, t_i) , and actions are all feasible paths connecting s_i and t_i . Each edge is associated with a nondecreasing cost (latency) function.

4 Tax Design for Congestion Games

In this section, we introduce tax design in congestion games. Before we get into the details, we will first introduce some notions to simplify the problem.

4.1 Polytope Description for Congestion Games

For a strategy $x_i \in \mathbb{R}^{|\mathcal{A}_i|}$, the dimension $|\mathcal{A}_i|$ can be as large as 2^F . Instead, it would be convenient to consider the facility load $y_i \in \mathbb{R}^F$ such that $y_{i,f} = \sum_{a \in \mathcal{A}_i: f \in a} x_{i,a}$. In addition, we define $y = \sum_{i \in [m]} y_i \in \mathbb{R}^F$ to be the total facility load. We use $\phi_{(i)}(\cdot)$ to denote the reparameterization mapping:

$$\phi(x) = y, \phi_i(x_i) = y_i, \forall i \in [m],$$

and we set $\mathcal{Y} = \{y \in \mathbb{R}^F : \exists x \in \mathcal{X}, y = \phi(x)\}$ to be the set of all feasible loads. Note that ϕ is not necessarily a bijection, i.e., there could exist multiple strategies sharing the same load. We use $\phi^{-1}(y) := \{x \in \mathcal{X} : \phi(x) = y\}$ to denote the set of strategies that are mapped to load y . The potential function can be defined after the repartitioning as well:

$$\Phi^{\text{repa}}(y) := \sum_f \int_0^{y_f} c_f(u) du = \Phi(x), \forall x \in \Phi^{-1}(y).$$

Importantly, $\Phi^{\text{repa}}(y)$ does not depend on the choice of strategy $x \in \phi^{-1}(y)$. For the repartitioned potential function, we have the following lemma showing that it is almost equivalent to the original potential function. When it is clear from the context, we will simplify Φ^{repa} as Φ .

Lemma 1. Φ^{repa} is convex under Assumption I. If $y^* = \operatorname{argmin}_y \Phi^{\text{repa}}(y)$, then for any $x \in \phi^{-1}(y^*)$, x is a Nash equilibrium.

For any Nash equilibrium strategy x , we call $y = \phi(x)$ the Nash equilibrium load (Nash load).

4.2 Optimal Tax for Congestion Games

Nash equilibrium is a stable state for a system with self-interested players, as no player has the incentive to deviate unilaterally. However, Nash equilibrium does not efficiently utilize the facilities, which is measured by the social cost:

$$\Psi(y) := \sum_f y_f c_f(y_f).$$

Price of anarchy is a concept that measures the efficiency of selfish agents in a system, defined as the ratio between the worst-case social cost for equilibria and the optimal social cost:

$$\text{PoA} = \frac{\max_{y \text{ is a Nash equilibrium load}} \Psi(y)}{\min_{y \in \mathcal{Y}} \Psi(y)}$$

For example, in nonatomic congestion games with polynomial cost functions, the price of anarchy grows as $\Theta(d/\ln d)$ where d is the degree of the polynomials [Nisan et al., 2007].

To reduce the price of anarchy, one standard approach is to enforce a tax on each facility to change the behavior of the self-interested players. Formally, a taxed congestion game is described by $(\mathcal{F}, \mathcal{A}_{[m]}, w_{[m]}, c_{\mathcal{F}}, \tau_{\mathcal{F}})$ with an additional tax function $\tau_f : [0, 1] \rightarrow \mathbb{R}$ on facility $f \in \mathcal{F}$. The cost of facility f with load u under tax becomes $c_f(u) + \tau_f(u)$. Correspondingly, we define the potential function with tax τ as

$$\Phi(y; \tau) := \sum_f \int_0^{y_f} [c_f(u) + \tau_f(u)] du,$$

and the Nash load would satisfy $y^* \in \operatorname{argmin}_y \Phi(y; \tau)$.

The optimal tax is defined as the tax that can induce optimal social behavior for self-interested players. We want to note that tax is not included in social cost following the convention in tax design.

Protocol 1 Online Tax Design for Congestion Games

Initialize: Facility set \mathcal{F} .
for $t = 1, 2, \dots, T$ **do**
 Designer chooses tax τ^t .
 Designer observes Nash load $y^t = \operatorname{argmin}_{y \in \mathcal{Y}} \Phi(y; \tau)$ and Nash cost $c^t = [c_f(y_f^t)]_{f \in \mathcal{F}}$ for $f \in \mathcal{F}$.
end for

Definition 2. A tax τ is an optimal tax if all Nash equilibria under tax τ can minimize the social cost:

$$\operatorname{argmin}_{y \in \mathcal{Y}} \Phi(y; \tau) \subseteq \operatorname{argmin}_{y \in \mathcal{Y}} \Psi(y).$$

In addition, a tax τ is an ϵ -optimal tax if we have

$$\Psi(y) \leq \min_{y' \in \mathcal{Y}} \Psi(y') + \epsilon, \forall y \in \operatorname{argmin}_{y'' \in \mathcal{Y}} \Phi(y''; \tau).$$

The marginal cost tax is defined as

$$\tau^* : \tau_f^*(u) = uc'_f(u), \forall f \in \mathcal{F}.$$

As $\Phi(y; \tau^*) = \Psi(y)$, the Nash equilibrium under tax τ^* will minimize the social cost and τ^* is an optimal tax [Nisan et al., 2007]. We will make the following assumption so that the cost combined with tax $c + \tau^*$ is still non-decreasing. In many real world problems, $c'_f(u)$ is non-decreasing due to the law of diminishing marginal utility, which guarantees Assumption 2.

Assumption 2. Marginal cost tax $\tau_f^*(u) = uc'_f(u)$ is non-decreasing for all $f \in \mathcal{F}$.

4.3 Tax Design for Congestion Games

In this paper, we consider the case where the system designer (e.g. government) wants to enforce an (approximate) optimal tax to induce optimal social behavior and maximize social welfare. However, the cost function is unknown so the optimal tax function cannot be computed directly via the marginal cost mechanism. On the other hand, the designer can enforce several taxes and observe the feedback. As Nash equilibrium is the stable state of the system, we assume the designer can observe the equilibrium feedback.

Formally, tax design is formulated as an online learning problem as shown in Protocol 1. At round t , the designer can choose a tax τ^t and observe the corresponding Nash equilibrium load $y^t \in \mathbb{R}^F$ and Nash equilibrium cost $c^t \in \mathbb{R}^F$. The sample complexity of a tax design algorithm is the number of rounds for designing an ϵ -optimal tax.

A naive approach is the designer first enumerates all of the ϵ -approximations of τ^* and chooses the tax with minimal social cost. However, such an approach would require $O((1/\epsilon)^{F\beta/\epsilon})$ samples as the complexity of using piece-wise linear function to approximate τ_f^* (a β -smooth function) with ϵ error is $O((1/\epsilon)^{\beta/\epsilon})$, resulting in exponential dependence on the parameters β , $1/\epsilon$ and F .

Another approach is applying algorithms for mathematical programming under equilibrium constraints. Specifically, we can formulate tax design as solving

$$\min_{\tau} \Psi(y(\tau)), \text{ s.t. } y(\tau) = \operatorname{argmin}_{y \in \mathcal{Y}} \Phi(y; \tau).$$

However, $y(\tau)$ can be non-differentiable or even discontinuous w.r.t. τ , and $\Psi(y(\tau))$ can be non-convex w.r.t. τ (Lemma 5). As a result, previous results do not apply to our problem as they apply gradient-based methods and make convexity assumptions [Li et al., 2020, Liu et al., 2022].

5 Learning Optimal Tax in Nonatomic Congestion Games

In this section, we describe our algorithm that can learn an ϵ -optimal tax with $O(F^2\beta/\epsilon)$ samples. First, we introduce piece-wise linear functions as a nonparametric way to approximate the marginal cost tax τ^* [Takezawa, 2005].

Definition 3. (Piece-wise Linear Function) We use a dictionary* $d = \{(x_1, y_1), \dots, (x_n, y_n)\}$ for $x_i \neq x_j, \forall i \neq j$ (w.l.o.g. we let $x_1 < x_2 < \dots < x_n$) to represent a piece-wise linear function $d(\cdot)$ on $[x_1, x_n]$ such that

$$d(x) = \frac{x - x_{i+1}}{x_i - x_{i+1}} y_i + \frac{x_i - x}{x_i - x_{i+1}} y_{i+1}, \forall x \in [x_i, x_{i+1}].$$

In addition, we use \cup to represent the update method for dictionary. I.e., $d \cup (x, y)$ is the piece-wise linear function interpolating one more point (x, y) if $(x, d(x))$ is not already in d , otherwise it will update $d(x)$ to y .

We will maintain the piece-wise function on a grid $\mathcal{L} = \{0, \Delta, 2\Delta, \dots, K\Delta = 1\}$ with $K = \lceil \frac{2\beta}{\epsilon} \rceil$ and $\Delta = 1/K$. The time complexity for computing $d(x)$ is $O(\log K)$ for any $x \in [0, 1]$.

5.1 Main Algorithm

Algorithm 1 Tax Design for Congestion Game

```

1: Initialize: Facility set  $\mathcal{F}$ , number of rounds  $T$ , tolerance  $\epsilon$ , smoothness  $\beta$ , perturbation  $\delta = \epsilon\Delta^2/8$ .
2: Set initial tax  $\tau^1 : \tau_f^1 = \{(0, 0), (1, \beta + \epsilon)\}$  for all  $f \in \mathcal{F}$ . Set  $\mathcal{K}_f^1$  to be  $\{0\}$  for all  $f \in \mathcal{F}$ .
3: for  $t = 1, 2, \dots, T$  do
4:   Observe Nash load  $y^t \in \mathbb{R}^F$  and Nash cost  $c^t \in \mathbb{R}^F$  under tax  $\tau^t$ .
5:   Set  $\bar{\mathcal{F}}$  to be the unknown facility set (Definition 4).
6:   Set  $l_f = \tau_f^t([y_f^t]_{\mathcal{K}_f^t}^-) + \epsilon(y_f^t - [y_f^t]_{\mathcal{K}_f^t}^-)$  and  $r_f = \tau_f^t([y_f^t]_{\mathcal{K}_f^t \cup \{1\}}^+) + \epsilon(y_f^t - [y_f^t]_{\mathcal{K}_f^t \cup \{1\}}^+)$  for
      each  $f \in \bar{\mathcal{F}}$ .
7:   Run Algorithm 2 with input  $y^t, c^t, \tau^t = [\tau_f^t(y_f^t)]_f, \bar{\mathcal{F}}$  and  $[l_f, r_f]_{f \in \bar{\mathcal{F}}}$ .
8:   if Algorithm 2 return False then
9:     return  $\tau^t$ 
10:  else
11:    Algorithm 2 return  $\tilde{\tau} \in \mathbb{R}^F, \tilde{f} \in \bar{\mathcal{F}}, \text{sign} \in \{-1, 1\}$ .
12:    Apply tax  $\dot{\tau}^t : \dot{\tau}_{\tilde{f}}^t = \tau_{\tilde{f}}^t \cup (y_{\tilde{f}}^t, \tilde{\tau}_{\tilde{f}}^t) + \text{sign} \cdot \delta$  and  $\dot{\tau}_f^t = \tau_f^t \cup (y_f^t, \tilde{\tau}_f^t)$  for  $f \neq \tilde{f}$ .
13:    Observe  $\dot{y}^t, \dot{c}^t \in \mathbb{R}^F$  as the Nash load and the Nash cost of each facility.
14:    Update  $\tau_{t+1}$  and  $\mathcal{K}_{t+1}$  according to (1).
15:  end if
16: end for

```

In this section, we introduce our main algorithm. At each round t , we will maintain a known index set $\mathcal{K}_f^t \subseteq \mathcal{L}$ where the marginal cost tax can be accurately estimated (Lemma 7), and use a piece-wise linear function to approximate the tax function by interpolating the values at the known indexes. The piece-wise linear function takes the form $\tau_f^t = \{(x_i^t, y_i^t)\}_i$ and the known index set \mathcal{K}_f^t satisfies $\{x_i^t\}_i = \mathcal{K}_f^t \cup \{1\}$ and $\mathcal{K}_f^t \subseteq \mathcal{L}$. Here 1 is a special case as it is not in the known index set initially but it is needed as the boundary for the piece-wise linear function τ_f^t . Initially, the tax is set to be $\tau_f^1(u) = \{(0, 0), (1, \beta + \epsilon)\}$ and the auxiliary tax is $\hat{\tau}_f^1 = \{(0, 0), (1, \beta)\}$ for $f \in \mathcal{F}$ (Line 2). Here we set $\hat{\tau}_f^1(1) = \beta$ as β is always an upper bound on $\tau_f^*(1)$. The auxiliary tax $\hat{\tau}_f^t$ is a non-decreasing piece-wise linear approximation of τ_f^* and we always set tax $\tau_f^t(u) = \hat{\tau}_f^t(u) + \epsilon u$ to ensure that the subgradient of the tax enforced is lower bounded by ϵ .

At round t , after observing Nash equilibrium load $y^t \in \mathbb{R}^F$ and Nash equilibrium cost $c^t \in \mathbb{R}^F$, the facilities are split into two sets: known facilities and unknown facilities.

Definition 4. For each round t , facility f is known if the Nash load $y_f^t \in [0, 1]$ satisfies $[y_f^t]_{\mathcal{L}}^- \in \mathcal{K}_f^t$ and $[y_f^t]_{\mathcal{L}}^+ \in \mathcal{K}_f^t$. Otherwise, facility f is unknown for round t .

For a known facility f , the Nash load is either in the known index set or sandwiched by two consecutive known indexes. As a result, the tax estimate for the Nash load $\tau_f^t(y_f^t)$ will be close to the

*In this dictionary, key is x_i and value is y_i . For readers unfamiliar with the dictionary data structure, it can be regarded as a set with a special update operation.

true optimal tax $\tau_f^*(y_f^t)$ with error 2ϵ (Lemma 8). We will apply Algorithm 2 to find the exploratory tax to gather information about unknown facilities (Line 7).

Proposition 1. *If Algorithm 2 return False at round t , then tax τ^t is an $6\epsilon F$ -optimal tax. If Algorithm 2 output $\tilde{\tau}^t, \tilde{f}^t, \text{sign}^t$ at round t , then we have*

$$0 < |y_{\tilde{f}^t}^t - \dot{y}_{\tilde{f}^t}^t| \leq \Delta.$$

If Algorithm 2 output $\tilde{\tau}^t, \tilde{f}^t, \text{sign}^t$ at round t , we update the tax and the known index set by the following rule. For $u \in \{[y_{\tilde{f}^t}^t]_{\mathcal{L}}^+, [y_{\tilde{f}^t}^t]_{\mathcal{L}}^-\} \setminus \mathcal{K}_{\tilde{f}^t}^t$ (this set is not empty as \tilde{f}^t is an unknown facility), we set

$$\hat{\tau}_{\tilde{f}^t}^{t+1} = \hat{\tau}_{\tilde{f}^t}^t \cup \left(u, \text{clip}\left(u \cdot \frac{c_{\tilde{f}^t}^t - \dot{c}_{\tilde{f}^t}^t}{y_{\tilde{f}^t}^t - \dot{y}_{\tilde{f}^t}^t}, \hat{\tau}_{\tilde{f}^t}^t([y_{\tilde{f}^t}^t]_{\mathcal{K}_{\tilde{f}^t}^t}^-), \hat{\tau}_{\tilde{f}^t}^t([y_{\tilde{f}^t}^t]_{\mathcal{K}_{\tilde{f}^t}^t}^+ \cup \{1\}) \right) \right), \quad (1)$$

$$\mathcal{K}_{\tilde{f}^t}^{t+1} = \mathcal{K}_{\tilde{f}^t}^t \cup \{u\}. \quad (2)$$

and $\hat{\tau}_f^{t+1} = \hat{\tau}_f^t, \mathcal{K}_f^{t+1} = \mathcal{K}_f^t$ for $f \neq \tilde{f}^t$. Then we set $\tau_f^{t+1}(u) = \hat{\tau}_f^{t+1}(u) + \epsilon u$ for all $f \in \mathcal{F}$ and $u \in [0, 1]$.

In words, we clip the two-point estimate $u \cdot \frac{c_{\tilde{f}^t}^t - \dot{c}_{\tilde{f}^t}^t}{y_{\tilde{f}^t}^t - \dot{y}_{\tilde{f}^t}^t}$ on the left and right known index of $\hat{\tau}_{\tilde{f}^t}^t(u)$ so that $\hat{\tau}_{\tilde{f}^t}^{t+1}(u)$ is still a non-decreasing piece-wise linear approximation of the marginal cost tax τ_f^* . τ_f^{t+1} is added with an extra linear term to guarantee a strongly convex potential function (Lemma 2).

As $0 < |y_f^t - \dot{y}_f^t| \leq \Delta$, the two point estimate of the gradient $\frac{c_f^t - \dot{c}_f^t}{y_f^t - \dot{y}_f^t}$ is accurate enough for $c'_f(u)$ such that $|\tau_{\tilde{f}^t}^{t+1}(u) - \tau_f(u)| \leq \epsilon$ (Lemma 6).

As $|\mathcal{K}_{\tilde{f}^t}^t|$ increases by 1 at round t and there are F such sets with size bounded by $O(\beta/\epsilon)$, Algorithm 2 will output False within at most $O(F\beta/\epsilon)$ rounds, which implies τ^t is an ϵF -optimal tax (Proposition 1). With proper rescaling, the sample complexity for learning ϵ -optimal tax is $O(F^2\beta/\epsilon)$.

Theorem 1. *Under Assumption 1 and Assumption 2, Algorithm 1 will output a $6\epsilon F$ tax within $T \leq 2F\beta/\epsilon$ rounds. In addition, each round has at most two tax realizations.*

Remark 1. *To uniformly approximate a β -smooth function, we have to know its value at $O(\beta/\epsilon)$ points [Takezawa, 2005]. For an ϵ -optimal tax, we need to estimate τ_f^* with ϵ/F accuracy as the error accumulates with all the facilities. As a result, we conjecture that $O(F^2\beta/\epsilon)$ sample complexity is tight and we leave the lower bound to future work.*

Remark 2. *Our algorithm can be easily adapted to the case where we have feedback other than only the equilibrium feedback. Specifically, when the tax designer obtain a non-equilibrium feedback, she can still update the optimal tax estimate if the feedback provides new information. It is possible for our algorithm to find the optimal tax even if no equilibrium feedback is provided. In addition, as long as the equilibrium can be reached after applying a tax, the algorithm can always find the optimal tax.*

5.2 Subroutine for Finding Exploratory Tax

In this section, we describe Algorithm 2 which can find an exploratory tax that satisfies Proposition 1. The idea is we can observe another similar but different Nash equilibrium load by perturbing the tax. However, there are two challenges:

1. Perturbing the tax might change the Nash equilibrium load drastically.
2. Perturbing the tax might not change the Nash equilibrium load at all.

To resolve the first issue, we always apply taxes that have (sub)gradient lower bounded by $\epsilon > 0$. The feasible range $[l_f, r_f]$ for updating tax τ_f^t with (y_f^t, \cdot) guarantees that the updated tax still maintains the subgradient lower bound. By Lemma 2 the potential function is always ϵ -strongly convex. As a result, the Nash load for any feasible tax is unique and Lipschitz w.r.t. tax perturbation. To resolve

Algorithm 2 Test Tax Design

```

1: Initialize: Nash flow  $y \in \mathbb{R}^F$ , tax  $\tau \in \mathbb{R}^F$ , cost  $c \in \mathbb{R}^F$ , unknown facility set  $\bar{\mathcal{F}}$ , unknown
   facility range  $[l_f, r_f]$  for  $f \in \bar{\mathcal{F}}$ .
2: Set strategy  $x \in \phi^{-1}(y)$ . Compute commodity load  $y_i = \phi_i(x_i) \in \mathbb{R}^F$  for  $i \in [m]$ .
3: Set  $I = \text{False}$ .
4: if Exists  $f \in \bar{\mathcal{F}}$  and  $i \in [m]$  such that  $0 < y_i(f) < w_i$  then
5:   return  $\tau, f, 1$ .
6: end if
7: for Commodity  $i \in [m]$  do
8:   Let  $\bar{\mathcal{F}}_i = \{f \in \bar{\mathcal{F}} : \sum_{a:f \in a} x_{i,a} = w_i\}$  and  $\bar{\mathcal{F}}'_i = \{f \in \bar{\mathcal{F}} : \sum_{a:f \in a} x_{i,a} = 0\}$ .
9:   Set  $\bar{\tau} : \bar{\tau}_{\bar{\mathcal{F}}_i} = r_{\bar{\mathcal{F}}_i}, \bar{\tau}_{\bar{\mathcal{F}}'_i} = l_{\bar{\mathcal{F}}'_i}, \bar{\tau}_{\mathcal{F} \setminus (\bar{\mathcal{F}}_i \cup \bar{\mathcal{F}}'_i)} = \tau_{\mathcal{F} \setminus (\bar{\mathcal{F}}_i \cup \bar{\mathcal{F}}'_i)}$ .
10:  if  $\text{Gap}_i(x, c + \bar{\tau}) < 0$  then
11:    Set  $I = \text{True}$ .
12:    break
13:  end if
14: end for
15: if  $I = \text{False}$  then
16:   return False.
17: end if
18: Set  $\tau' = \tau$ 
19: for  $f \in \bar{\mathcal{F}}_i$  do
20:   Set  $\tilde{\tau}^u : \tilde{\tau}_f^u = u, \tilde{\tau}_{\mathcal{F} \setminus \{f\}}^u = \tau'_{\mathcal{F} \setminus \{f\}}$ .
21:   Set  $u = \text{argmax}\{u : \text{Gap}_j(x, c + \tilde{\tau}^u) \geq 0, \forall j\}$ .
22:   if  $u \leq r_f$  then
23:     return  $\tilde{\tau}^u, f, 1$ .
24:   end if
25:   Set  $\tau' = \tilde{\tau}^u$ .
26: end for
27: for  $f \in \bar{\mathcal{F}}'_i$  do
28:   Set  $\tilde{\tau}^u : \tilde{\tau}_f^u = u, \tilde{\tau}_{\mathcal{F} \setminus \{f\}}^u = \tau'_{\mathcal{F} \setminus \{f\}}$ .
29:   Set  $u = \text{argmin}\{u : \text{Gap}_j(x, c + \tilde{\tau}^u) \geq 0, \forall j\}$ .
30:   if  $u \geq l_f$  then
31:     return  $\tilde{\tau}^u, f, -1$ .
32:   end if
33:   Set  $\tau' = \tilde{\tau}^u$ .
34: end for

```

the second issue, we find the tax that makes the current Nash equilibrium on the “boundary”. I.e., an additional perturbation will make the Nash equilibrium change. Intuitively, this is similar to removing the slackness in a constrained optimization problem. By Lemma 3, we can observe a different Nash load on f if we make the additional perturbation.

Lemma 2. *If the subgradient of the cost function c_f is lower bounded by $\epsilon > 0$ for all $f \in \mathcal{F}$, then the potential function $\Phi^{\text{repa}}(y)$ is ϵ -strongly convex. However, $\Phi(x)$ is not necessarily strongly convex.*

Lemma 3. *If two taxes τ and $\tilde{\tau}$ only differ in facility f and the Nash loads y and \tilde{y} are different, then $y_f \neq \tilde{y}_f$.*

Definition 5. *The gap for a strategy $x \in \mathcal{X}$ with cost $c \in \mathbb{R}^F$ is defined as*

$$\text{Gap}_i(x, c) = \min_{a:x_{i,a}=0} \sum_{f:f \in a} c_f - \max_{a:x_{i,a} \neq 0} \sum_{f:f \in a} c_f. \quad (3)$$

In the algorithm, we use $\text{Gap}_i(x, c)$ to measure the cost gap between in-support actions and off-support actions for commodity i and strategy x . If x is a Nash equilibrium and c is the Nash cost, then all of the in-support actions have the same minimal cost and $\text{Gap}_i(x, c) \geq 0$ holds. Informally, “boundary” tax τ means that $\text{Gap}_i(x, c + \tau) = 0$ for a Nash equilibrium x and perturbing τ results in $\text{Gap}_i(x, c + \tau) < 0$, so the Nash equilibrium under the perturbed tax will be different from x .

Now we discuss how Algorithm 2 finds the “boundary” tax in detail. The input to the algorithm is the Nash flow y , the Nash cost c , the Nash tax τ , the unknown facility set $\bar{\mathcal{F}}$ and the feasible tax range $[l_f, r_f]$ for each unknown facility $f \in \bar{\mathcal{F}}$. We emphasize that here the Nash cost/tax are the values of the cost/tax function on the Nash load and they are vectors in \mathbb{R}^F instead of functions. $[l_f, r_f]$ is the feasible range for the perturbed tax value at facility f . By the definition of l_f and r_f , current tax τ_f^t updated with (y_f^t, u) , $u \in [l_f, r_f]$ is still a tax with subgradient lower bounded by ϵ .

For the first step, the algorithm will compute strategy $x \in \phi^{-1}(y)$ as the Nash equilibrium strategy (Line 2). If there exists an unknown facility f and commodity i such that not all load of commodity i is using f or not using f (Line 4), then perturbing the tax at facility f will make x not longer a Nash equilibrium as in-support actions have different costs.

Otherwise, for each unknown facility f and commodity i , either all of the load is using f or all of the load is not using f . As a result, in-support actions always have the same cost after perturbing the tax. For the next step, we verify if there exists a tax within the feasible ranges for unknown facilities that makes x not a Nash equilibrium. However, there does not exist a universal worst-case tax that can verify if x is always a Nash equilibrium or not.

Fortunately, the worst-case tax has a closed form for each commodity separately: the taxes for facilities used by all of the Nash load would be the upper bound r_f and the taxes for facilities used by none of the Nash load would be the lower bound l_f , thus maximizing the cost for in-support actions and minimizing the cost for off-support actions (Line 9). For each commodity i , we apply the corresponding worst-case tax and check if the in-support actions are still the optimal actions (Line 10). If for all commodities, the in-support actions are optimal under the worst-case tax, then for any tax within the feasible range, y is the Nash load and the algorithm will output False (Line 16). As τ^* is approximately within the range, y^t approximately minimizes the social cost (Lemma 10).

Otherwise, the algorithm finds commodity i such that x is not the Nash equilibrium under the worst-case tax (Line 11). For the last step, we gradually transform the initial tax to this worst-case tax and stop when x is not the Nash equilibrium for some commodity. Specifically, the algorithm iteratively changes the tax in the unknown facility set $\bar{\mathcal{F}} = \bar{\mathcal{F}}_i \cup \bar{\mathcal{F}}'_i$ (Line 19 and Line 27) to the worst-case tax.

For facility $f \in \bar{\mathcal{F}}_i$, the algorithm finds the boundary tax for facility f that satisfies

$$u = \underset{u}{\operatorname{argmax}} \{u : \operatorname{Gap}_j(x, c + \tilde{\tau}^u) \geq 0, \forall j \in [m]\}.$$

If $u \leq r_f$, we will output $\tilde{\tau}^u, f, 1$. By the definition of u , if we further increase u , one of the gaps will become negative and x is no longer the Nash equilibrium. Otherwise, all feasible taxes for f have a nonnegative gap for all commodities, which means x is still the Nash equilibrium, and we continue for the next facility. After enumerating all the facilities in $\bar{\mathcal{F}}_i$, we enumerate $\bar{\mathcal{F}}'_i$ in the same way. Eventually, the tax is transformed into the worst-case tax with negative gap for commodity i , so this process will end and output $\tilde{\tau}^u, f, \text{sign}$ such that $\tilde{\tau}^u$ is the tax that makes the Nash equilibrium on the boundary, f is the facility to perturb and sign is the direction to perturb the tax at f .

6 Conclusion

We proposed the first algorithm with polynomial sample complexity for learning optimal tax in nonatomic congestion games. The algorithm leverages several novel designs to exploit the special structure of congestion games, which can also be implemented efficiently. Below we list a few potential future research directions:

1. Relaxing the Nash equilibrium assumption to players following no-regret dynamics or quantal response equilibrium.
2. Design algorithms that do not require prior knowledge of the smoothness coefficient.
3. Generalize the algorithm to atomic congestion games.

7 Acknowledgement

SSD acknowledges the support of NSF IIS 2110170, NSF DMS 2134106, NSF CCF 2212261, NSF IIS 2143493, NSF CCF 2019844, and NSF IIS 2229881. MF acknowledges the support of NSF awards CCF 2007036, TRIPDS II DMS 2023166, CCF 2212261, and CCF-AF 2312775.

References

Yu Bai, Chi Jin, Huan Wang, and Caiming Xiong. Sample-efficient learning of stackelberg equilibria in general-sum games. *Advances in Neural Information Processing Systems*, 34:25799–25811, 2021.

Martin Beckmann, Charles B McGuire, and Christopher B Winsten. Studies in the economics of transportation. Technical report, 1956.

Pia Bergendorff, Donald W Hearn, and Motakuri V Ramana. *Congestion toll pricing of traffic networks*. Springer, 1997.

Vittorio Bilò and Cosimo Vinci. Dynamic taxes for polynomial congestion games. *ACM Transactions on Economics and Computation (TEAC)*, 7(3):1–36, 2019.

Avrim Blum, Nika Haghtalab, and Ariel D Procaccia. Learning optimal commitment to overcome insecurity. *Advances in Neural Information Processing Systems*, 27, 2014.

Dietrich Braess. Über ein paradoxon aus der verkehrsplanung. *Unternehmensforschung*, 12:258–268, 1968.

Ioannis Caragiannis, Christos Kaklamanis, and Panagiotis Kanellopoulos. Taxes for linear atomic congestion games. *ACM Transactions on Algorithms (TALG)*, 7(1):1–31, 2010.

Po-An Chen and Chi-Jen Lu. Generalized mirror descents in congestion games. *Artificial Intelligence*, 241:217–243, 2016.

Michael B Cohen, Yin Tat Lee, and Zhao Song. Solving linear programs in the current matrix multiplication time. *Journal of the ACM (JACM)*, 68(1):1–39, 2021.

Benoît Colson, Patrice Marcotte, and Gilles Savard. An overview of bilevel optimization. *Annals of operations research*, 153:235–256, 2007.

Qiwen Cui, Zhihan Xiong, Maryam Fazel, and Simon S Du. Learning in congestion games with bandit feedback. *Advances in Neural Information Processing Systems*, 35:11009–11022, 2022.

Leello Dadi, Ioannis Panageas, Stratis Skoulakis, Luca Viano, and Volkan Cevher. Polynomial convergence of bandit no-regret dynamics in congestion games. *arXiv preprint arXiv:2401.09628*, 2024.

Jing Dong, Jingyu Wu, Siwei Wang, Baoxiang Wang, and Wei Chen. Taming the exponential action set: Sublinear regret and fast convergence to nash equilibrium in online congestion games. *arXiv preprint arXiv:2306.13673*, 2023.

Panagiotis D Grontas, Giuseppe Belgioioso, Carlo Cenedese, Marta Fochesato, John Lygeros, and Florian Dörfler. Big hype: Best intervention in games via distributed hypergradient descent. *IEEE Transactions on Automatic Control*, 2024.

Tobias Harks, Ingo Kleinert, Max Klimm, and Rolf H Möhring. Computing network tolls with support constraints. *Networks*, 65(3):262–285, 2015.

Haozhe Jiang, Qiwen Cui, Zhihan Xiong, Maryam Fazel, and Simon S Du. Offline congestion games: How feedback type affects data coverage requirement. *arXiv preprint arXiv:2210.13396*, 2022.

Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *Annual symposium on theoretical aspects of computer science*, pages 404–413. Springer, 1999.

Walid Krichene, Benjamin Drighès, and Alexandre M Bayen. Online learning of nash equilibria in congestion games. *SIAM Journal on Control and Optimization*, 53(2):1056–1081, 2015.

Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.

Joshua Letchford, Vincent Conitzer, and Kamesh Munagala. Learning and approximating the optimal strategy to commit to. In *Algorithmic Game Theory: Second International Symposium, SAGT 2009, Paphos, Cyprus, October 18–20, 2009. Proceedings* 2, pages 250–262. Springer, 2009.

Jiayang Li, Jing Yu, Yu Nie, and Zhaoran Wang. End-to-end learning and intervention in games. *Advances in Neural Information Processing Systems*, 33:16653–16665, 2020.

Jiayang Li, Jing Yu, Qianni Wang, Boyi Liu, Zhaoran Wang, and Yu Marco Nie. Differentiable bilevel programming for stackelberg congestion games. *arXiv preprint arXiv:2209.07618*, 2022.

Jiayang Li, Jing Yu, Boyi Liu, Yu Nie, and Zhaoran Wang. Achieving hierarchy-free approximation for bilevel programs with equilibrium constraints. In *International Conference on Machine Learning*, pages 20312–20335. PMLR, 2023.

Boyi Liu, Jiayang Li, Zhuoran Yang, Hoi-To Wai, Mingyi Hong, Yu Nie, and Zhaoran Wang. Inducing equilibria via incentives: Simultaneous design-and-play ensures global convergence. *Advances in Neural Information Processing Systems*, 35:29001–29013, 2022.

Zhi-Quan Luo, Jong-Shi Pang, and Daniel Ralph. *Mathematical programs with equilibrium constraints*. Cambridge University Press, 1996.

Chinmay Maheshwari, S Shankar Sasty, Lillian Ratliff, and Eric Mazumdar. Convergent first-order methods for bi-level optimization and stackelberg games. *arXiv preprint arXiv:2302.01421*, 2023.

Jason R Marden and Tim Roughgarden. Generalized efficiency bounds in distributed resource allocation. *IEEE Transactions on Automatic Control*, 59(3):571–584, 2014.

Dov Monderer and Lloyd S Shapley. Potential games. *Games and economic behavior*, 14(1):124–143, 1996.

Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.

Dario Paccagnan and Martin Gairing. In congestion games, taxes achieve optimal approximation. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 743–744, 2021.

Dario Paccagnan, Rahul Chandan, Bryce L Ferguson, and Jason R Marden. Optimal taxes in atomic congestion games. *ACM Transactions on Economics and Computation (TEAC)*, 9(3):1–33, 2021.

Ioannis Panageas, Stratis Skoulakis, Luca Viano, Xiao Wang, and Volkan Cevher. Semi bandit dynamics in congestion games: Convergence to nash equilibrium and no-regret guarantees. In *International Conference on Machine Learning*, pages 26904–26930. PMLR, 2023.

Binghui Peng, Weiran Shen, Pingzhong Tang, and Song Zuo. Learning optimal strategies to commit to. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 2149–2156, 2019.

Arthur Cecil Pigou. *Wealth and welfare*. Macmillan and Company, limited, 1912.

Robert W Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.

Tim Roughgarden. *Selfish routing and the price of anarchy*. MIT press, 2005.

Tim Roughgarden and Éva Tardos. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and economic behavior*, 47(2):389–403, 2004.

Kunio Takezawa. *Introduction to nonparametric regression*. John Wiley & Sons, 2005.

Heinrich Von Stackelberg. *Market structure and equilibrium*. Springer Science & Business Media, 2010.

John Glen Wardrop. Road paper. some theoretical aspects of road traffic research. *Proceedings of the institution of civil engineers*, 1(3):325–362, 1952.

Geng Zhao, Banghua Zhu, Jiantao Jiao, and Michael Jordan. Online learning in stackelberg games with an omniscient follower. In *International Conference on Machine Learning*, pages 42304–42316. PMLR, 2023.

Han Zhong, Zhuoran Yang, Zhaoran Wang, and Michael I Jordan. Can reinforcement learning find stackelberg-nash equilibria in general-sum markov games with myopic followers? *arXiv preprint arXiv:2112.13521*, 2021.

A Basics about Congestion Games

Lemma 4. *If strategy x is an ϵ -NE in a congestion game, then x is an ϵ -minimizer of the corresponding potential function $\Phi(\cdot)$.*

Proof. Let $x^* = \operatorname{argmin}_{x \in \mathcal{X}} \Phi(x)$ and $y = \phi(x)$. First, we show that

$$\begin{aligned} \nabla_{i,a} \Phi(x) &= \nabla_{i,a} \sum_f \int_0^{y_f} c_f(u) du \\ &= \nabla_{i,a} \sum_{f \in a} \int_0^{y_f} c_f(u) du \\ &= \sum_{f \in a} c_f(y_f) \nabla_{i,a} y_f \\ &= \sum_{f \in a} c_f(y_f) \nabla_{i,a} \sum_{i',a':f \in a'} x_{i',a'} \\ &= \sum_{f \in a} c_f(y_f). \end{aligned}$$

Then we have

$$\begin{aligned} \Phi(x) - \Phi(x^*) &\leq \langle x - x^*, \nabla \Phi(x) \rangle && \text{(Convexity)} \\ &\leq \sum_{i \in [m]} \langle x_i - x_i^*, \nabla_i \Phi(x) \rangle \\ &\leq \sum_{i \in [m]} \left[\sum_{a \in \mathcal{A}_i} x_{i,a} \sum_{f \in a} c_f(y_f) - \min_{a \in \mathcal{A}_i} w_i \sum_{f \in a} c_f(y_f) \right] \\ &\leq \sum_{i \in [m]} \sum_{a \in \mathcal{A}_i} x_{i,a} \epsilon \\ &= \sum_{i \in [m]} w_i \epsilon \\ &= \epsilon. \end{aligned}$$

□

Lemma 1. Φ^{repa} is convex under Assumption L. If $y^* = \operatorname{argmin}_y \Phi^{\text{repa}}(y)$, then for any $x \in \phi^{-1}(y^*)$, x is a Nash equilibrium.

Proof. For $y^1, y^2 \in \mathcal{Y}$, we have

$$\Phi^{\text{repa}}(y^1) + \Phi^{\text{repa}}(y^2) - 2\Phi^{\text{repa}}\left(\frac{y^1 + y^2}{2}\right) = \sum_f \left[\int_0^{y_f^1} c_f(u) du + \int_0^{y_f^2} c_f(u) du - 2 \int_0^{\frac{y_f^1 + y_f^2}{2}} c_f(u) du \right].$$

Now we show that $\int_0^{y_f^1} c_f(u) du + \int_0^{y_f^2} c_f(u) du - 2 \int_0^{\frac{y_f^1 + y_f^2}{2}} c_f(u) du$ is nonnegative for all $f \in \mathcal{F}$. W.l.o.g., we assume $y_f^1 \leq y_f^2$ and we have

$$\begin{aligned} \int_0^{y_f^1} c_f(u) du + \int_0^{y_f^2} c_f(u) du - 2 \int_0^{\frac{y_f^1 + y_f^2}{2}} c_f(u) du &= \int_{\frac{y_f^1 + y_f^2}{2}}^{y_f^2} c_f(u) du - \int_{y_f^1}^{\frac{y_f^1 + y_f^2}{2}} c_f(u) du \\ &= \int_{y_f^1}^{\frac{y_f^1 + y_f^2}{2}} \left[c_f(u + \frac{y_f^2 - y_f^1}{2}) - c_f(u) \right] du \\ &\geq 0, \end{aligned}$$

where the last step is from Assumption 1 (monotonicity). As a result, Φ^{repa} is convex.

Let $y^* = \operatorname{argmin}_y \Phi^{\text{repa}}(y)$ and $x \in \phi^{-1}(y^*)$. If there exists $x' \in \mathcal{X}$ such that $\Phi(x') < \Phi(x)$, then we have $\Phi^{\text{repa}}(\phi(x')) < \Phi^{\text{repa}}(y^*)$, which contradicts the definition of y^* . As a result, x is the minimizer of $\Phi(\cdot)$, which means x is a Nash equilibrium. \square

Lemma 5. *The Nash load under tax τ : $y(\tau) = \operatorname{argmin}_{y \in \mathcal{Y}} \Phi(y; \tau)$ is not continuous w.r.t. τ . In addition, the social welfare $\Psi(y(\tau))$ is not convex w.r.t. τ .*

Proof. For the first part, we construct a congestion game with two facilities f_1, f_2 , one commodity with action set $\{f_1, f_2\}$, and constant cost $c_1 = 1, c_2 = 1 - \epsilon$ with $\epsilon > 0$. Then for tax $\tau = 0$, we have $y(\tau) = [0, 1]$. For constant tax $\tau_1 = 0, \tau_2 = 2\epsilon$, we have $y(\tau) = [1, 0]$. As ϵ can be arbitrarily small, $y(\tau)$ is not continuous w.r.t. τ .

For the second part, we construct a congestion game with two facilities f_1, f_2 , one commodity with action set $\{f_1, f_2\}$, and cost function $c_1 = 1, c_2(u) = \sqrt{u}$ for $u \in [0, 1]$. We apply constant tax $\tau : \tau_1 = t, \tau_2 = 0$ for $t \in [-1, 0]$. The Nash equilibrium under tax τ is $y(\tau) = [1 - (1+t)^2, (1+t)^2]$. Then the social cost is $\Psi(y(\tau)) = 1 - t(1+t)^2$, which is not convex on $[-1, 0]$. \square

B Missing Proofs in Section 5

Lemma 2. *If the subgradient of the cost function c_f is lower bounded by $\epsilon > 0$ for all $f \in \mathcal{F}$, then the potential function $\Phi^{\text{repa}}(y)$ is ϵ -strongly convex. However, $\Phi(x)$ is not necessarily strongly convex.*

Proof. First, by the definition of the potential function Φ , it is easy to show that $\nabla \Phi(y) = [c_f(y_f)]_{f \in \mathcal{F}}$. For $y^1, y^2 \in \mathcal{Y}$, we have

$$(\nabla \Phi(y^1) - \nabla \Phi(y^2))^\top (y^1 - y^2) = \sum_{f \in \mathcal{F}} (c_f(y_f^1) - c_f(y_f^2))(y_f^1 - y_f^2) \geq \sum_{f \in \mathcal{F}} \epsilon (y_f^1 - y_f^2)^2 = \epsilon \|y^1 - y^2\|_2^2,$$

which implies $\Phi(\cdot)$ is a ϵ -strongly convex function.

For the second argument, we only need to construct a congestion game such that there exists two strategy $x^1, x^2 \in \mathcal{X}$ such that for $\phi(tx^1 + (1-t)x^2)$ is a constant for $t \in [0, 1]$, which implies the potential function $\Phi(tx^1 + (1-t)x^2) = \Phi^{\text{repa}}(\phi(tx^1 + (1-t)x^2))$ is a constant w.r.t. t . However, a strongly convex function cannot be a constant on a line, which implies Φ is not strongly convex.

We construct a congestion game with three facilities f_1, f_2, f_3 , three actions $a_1 = \{f_1\}, a_2 = \{f_2\}, a_3 = \{f_3\}$ and three commodities with action set $\{a_1, a_2\}, \{a_2, a_3\}, \{a_3, a_1\}$. Strategy $x^1 : x_1^1 = [1, 0, 0], x_2^1 = [0, 1, 0], x_3^1 = [0, 0, 1]$ and $x^2 : x_1^2 = [0, 1, 0], x_2^2 = [0, 0, 1], x_3^2 = [1, 0, 0]$. Then $tx^1 + (1-t)x^2$ is a feasible strategy and we have $\phi(tx^1 + (1-t)x^2) = [1, 1, 1]$ for all $t \in [0, 1]$. \square

Lemma 3. *If two taxes τ and $\dot{\tau}$ only differ in facility f and the Nash loads y and \dot{y} are different, then $y_f \neq \dot{y}_f$.*

Proof. For simplicity, we consider the equivalent tax-free case that we have two costs c, \dot{c} with subgradient lower bounded by ϵ and they only differ in facility f . The potential functions are

$$\Phi(Y) = \sum_f \int_0^{Y_f} c_f(u) du, \dot{\Phi}(Y) = \sum_f \int_0^{Y_f} \dot{c}_f(u) du.$$

By Lemma 2, Φ and $\dot{\Phi}$ are strongly convex and thus the Nash equilibrium load y and \dot{y} are unique. Suppose $y_f = \dot{y}_f$. Consider any $Y \in \mathcal{Y}$ such that $Y_f = y_f$, we have

$$\begin{aligned} \dot{\Phi}(Y) - \dot{\Phi}(y) &= \sum_f \int_{y_f}^{Y_f} \dot{c}_f(u) du = \sum_{f \neq f} \int_{y_f}^{Y_f} \dot{c}_f(u) du + \int_{y_f}^{Y_f} \dot{c}_f(u) du \\ &= \sum_{f \neq f} \int_{y_f}^{Y_f} c_f(u) du = \Phi(Y) - \Phi(y) \geq 0. \end{aligned}$$

As a result, we have $\dot{\Phi}(y) \leq \dot{\Phi}(\dot{y})$. By the optimality of \dot{y} , we have $y = \dot{y}$. By contradiction, if $y \neq \dot{y}$, we have $y_f = \dot{y}_f$. \square

Lemma 6. *If $|u_1 - u_2| \leq \Delta$, then for any $|u_3 - u_1| \leq \Delta$, we have*

$$\left| \frac{c_f(u_1) - c_f(u_2)}{u_1 - u_2} - c'_f(u_3) \right| \leq \epsilon.$$

Proof. This is a direct corollary of the β -smoothness. By mean value theorem, we have $\frac{c_f(u_1) - c_f(u_2)}{u_1 - u_2} = c'_f(u)$ for some $u \in [u_1, u_2]$. As $|u - u_3| \leq |u - u_1| + |u_1 - u_3| \leq 2\Delta \leq \frac{\epsilon}{2\beta}$, we have

$$\left| \frac{c_f(u_1) - c_f(u_2)}{u_1 - u_2} - c'_f(u_3) \right| \leq \epsilon.$$

\square

Lemma 7. *For round t and facility f , if $u \in \mathcal{K}_f^t$, then we have $|\tau_f^t(u) - \tau_f^*(u)| \leq 2\epsilon$.*

Proof. By the algorithm design, for each $u \in \mathcal{K}_f^t$, $\tau_f^t(u)$ will not change after u is added to \mathcal{K}_f . We will use induction on t to prove $|\tau_f^t(u) - \tau_f(u)| \leq \epsilon$ for $u \in \mathcal{K}_f^t$. At round $t = 1$, $\mathcal{K}_f^1 = \{0\}$ and $\widehat{\tau}_f^1(0) = \tau_f^*(0) = 0$ holds.

Suppose at round t , we have $\mathcal{K}_{\tilde{f}^t}^{t+1} = \mathcal{K}_{\tilde{f}^t}^t \cup \{u\}$ with $u \in \{[y_{\tilde{f}^t}^t]_{\mathcal{L}}^+, [y_{\tilde{f}^t}^t]_{\mathcal{L}}^-\} \setminus \mathcal{K}_{\tilde{f}^t}^t$, and $\mathcal{K}_f^{t+1} = \mathcal{K}_f^t$ for $f \neq \tilde{f}^t$. By the induction hypothesis, we only need to prove $|\widehat{\tau}_{\tilde{f}^t}^{t+1}(u) - \tau_{\tilde{f}^t}^*(u)| \leq 2\epsilon$. Recall that

$$\widehat{\tau}_{\tilde{f}^t}^{t+1}(u) = \text{clip}\left(u \cdot \frac{c_{\tilde{f}^t}^t - \dot{c}_{\tilde{f}^t}^t}{y_{\tilde{f}^t}^t - \dot{y}_{\tilde{f}^t}^t}, \widehat{\tau}_{\tilde{f}^t}^t([y_{\tilde{f}^t}^t]_{\mathcal{K}_{\tilde{f}^t}^t}^-), \widehat{\tau}_{\tilde{f}^t}^t([y_{\tilde{f}^t}^t]_{\mathcal{K}_{\tilde{f}^t}^t}^+ \cup \{1\})\right).$$

Then we have the following three cases. For simplicity we replace \tilde{f}^t with f .

(1) $\widehat{\tau}_f^{t+1}(u) = u \cdot \frac{c_f^t - \dot{c}_f^t}{y_f^t - \dot{y}_f^t}$. By Lemma 11 and Lemma 6, we have

$$|\widehat{\tau}_f^{t+1}(u) - \tau_f^*(u)| = \left| u \frac{c_f^t - \dot{c}_f^t}{y_f^t - \dot{y}_f^t} - u c'_f(u) \right| \leq \epsilon.$$

(2) $\widehat{\tau}_f^{t+1}(u) = \widehat{\tau}_f^t([y_f^t]_{\mathcal{K}_f^t}^-)$ and $u \cdot \frac{c_f^t - \dot{c}_f^t}{y_f^t - \dot{y}_f^t} \leq \widehat{\tau}_f^t([y_f^t]_{\mathcal{K}_f^t}^-)$. Then we have

$$\widehat{\tau}_f^{t+1}(u) = \widehat{\tau}_f^t([y_f^t]_{\mathcal{K}_f^t}^-) \leq \tau_f^*([y_f^t]_{\mathcal{K}_f^t}^-) + \epsilon \leq \tau_f^*(u) + \epsilon,$$

where the first inequality is from the induction hypothesis as $[y_f^t]_{\mathcal{K}_f^t}^- \in \mathcal{K}_f^t$ and the second inequality is from Assumption 2. In addition, we have

$$\widehat{\tau}_f^{t+1}(u) \geq u \cdot \frac{c_f^t - \dot{c}_f^t}{y_f^t - \dot{y}_f^t} \geq u c'_f(u) - \epsilon = \tau_f^*(u) - \epsilon.$$

(3) $\widehat{\tau}_f^{t+1}(u) = \widehat{\tau}_f^t([y_f^t]_{\mathcal{K}_f^t}^+ \cup \{1\})$ and $u \cdot \frac{c_f^t - \dot{c}_f^t}{y_f^t - \dot{y}_f^t} \geq \widehat{\tau}_f^t([y_f^t]_{\mathcal{K}_f^t}^+ \cup \{1\})$. Then we have

$$\widehat{\tau}_f^{t+1}(u) \leq u \frac{c_f^t - \dot{c}_f^t}{y_f^t - \dot{y}_f^t} \leq u c'_f(u) + \epsilon \leq \tau_f^*(u) + \epsilon.$$

If $[y_f^t]_{\mathcal{K}_f^t}^+ \cup \{1\} \in \mathcal{K}_f^t$, then we have

$$\widehat{\tau}_f^t(u) = \widehat{\tau}_f^t([y_f^t]_{\mathcal{K}_f^t}^+ \cup \{1\}) \geq \tau_f^*([y_f^t]_{\mathcal{K}_f^t}^+ \cup \{1\}) - \epsilon \geq \tau_f^*(u) - \epsilon.$$

If $[y_f^t]_{\mathcal{K}_f^t \cup \{1\}}^+ = 1$ and $1 \notin \mathcal{K}_f^t$, we still have

$$\widehat{\tau}_f^t(u) = \beta \geq uc'_f(u) = \tau_f^*(u).$$

For each of these three cases, the induction holds.

As $\tau_f^t(u) = \widehat{\tau}_f^t(u) + \epsilon u$ for all $f \in \mathcal{F}$ and $u \in [0, 1]$, we have $|\tau_f^t(u) - \tau_f^*(u)| \leq 2\epsilon$. \square

Lemma 8. For round t , if facility f is known, then we have $|\tau_f^t(y_f^t) - \tau_f^*(y_f^t)| \leq 3\epsilon$.

Proof. If $y_f^t \in \mathcal{K}_f^t$, we can directly apply Lemma 7. Otherwise, we set $u_1 = [y_f^t]_{\mathcal{L}}^-$ and $u_2 = [y_f^t]_{\mathcal{L}}^+$. Then we have $u_1 < y_f^t < u_2$ and $u_1, u_2 \in \mathcal{K}_f^t$. There exists $\lambda_1 \in [0, 1], \lambda_1 + \lambda_2 = 1$ such that $y_f^t = \lambda_1 u_1 + \lambda_2 u_2$. By Lemma 6, we have $|\tau_f^t(u_i) - \tau_f^*(u_i)| \leq 2\epsilon$ for $i \in \{1, 2\}$. Then we have

$$\begin{aligned} & |\tau_f^t(y_f^t) - \tau_f^*(y_f^t)| \\ &= |\lambda_1 \tau_f^t(u_1) + \lambda_2 \tau_f^t(u_2) - (\lambda_1 u_1 + \lambda_2 u_2) c'_f(u)| \\ &\leq |\lambda_1 \tau_f^t(u_1) - \lambda_1 u_1 c'_f(u)| + |\lambda_2 \tau_f^t(u_2) - \lambda_2 u_2 c'_f(u)| \\ &\leq \lambda_1 |\tau_f^t(u_1) - \tau_f^*(u_1)| + \lambda_1 u_1 |c'_f(u_1) - c'_f(u)| + \lambda_2 |\tau_f^t(u_2) - \tau_f^*(u_2)| + \lambda_2 u_2 |c'_f(u_2) - c'_f(u)| \\ &\leq 2\lambda_1 \epsilon + \lambda_1 \epsilon + 2\lambda_2 \epsilon + \lambda_2 \epsilon \quad (\text{Lemma 6, } \beta\text{-smoothness and } |u - u_i| \leq \epsilon/\beta.) \\ &\leq 3\epsilon. \end{aligned}$$

\square

Lemma 9. If Algorithm 2 return False at round t , then for any $\tilde{\tau} \in \mathbb{R}^{\mathcal{F}}$ such that $\tilde{\tau}_f = \tau_f^t(y_f)$ for $f \in \mathcal{F} \setminus \bar{\mathcal{F}}^t$ and $\tilde{\tau}_f \in [l_f^t, r_f^t]$ for $f \in \bar{\mathcal{F}}^t$, we have $\text{Gap}_i(x^t, c^t + \tilde{\tau}) \geq 0$ for all $i \in [m]$. In addition, x^t is a Nash equilibrium for tax $\tilde{\tau}$.

Proof. For simplicity, we will omit t when there is no confusion. Algorithm 2 return False if and only if for all $i \in [m]$ and tax $\bar{\tau}_i : \bar{\tau}_{\bar{F}_i} = r_{\bar{F}_i}, \bar{\tau}_{\bar{F}'_i} = l_{\bar{F}'_i}, \bar{\tau}_{\mathcal{F} \setminus (\bar{F}_i \cup \bar{F}'_i)} = \tau_{\mathcal{F} \setminus (\bar{F}_i \cup \bar{F}'_i)}$, we have

$$\text{Gap}_i(x, c + \bar{\tau}) = \min_{a: x_{i,a} = 0} \sum_{f: f \in a} (c_f + \bar{\tau}_f) - \max_{a: x_{i,a} \neq 0} \sum_{f: f \in a} (c_f + \bar{\tau}_f) \geq 0.$$

By the definition of \bar{F}_i , for any $f \in \bar{F}_i$ and $a : x_{i,a} \neq 0$, we have $f \in a$. Similarly, for any $f \in \bar{F}'_i$ and $a : x_{i,a} \neq 0$, we have $f \notin a$. Thus for any $a : x_{i,a} \neq 0$, we have

$$\sum_{f: f \in a} (c_f + \bar{\tau}_f) - \sum_{f: f \in a} (c_f + \tilde{\tau}_f) = \sum_{f \in \bar{F}_i} (r_f - \tilde{\tau}_f) \geq 0.$$

For any $a : x_{i,a} = 0$, we have

$$\sum_{f: f \in a} (c_f + \bar{\tau}_f) - \sum_{f: f \in a} (c_f + \tilde{\tau}_f) = \sum_{f \in \bar{F}'_i \cap a} (l_f - \tilde{\tau}_f) \leq 0.$$

As a result, we have

$$\begin{aligned} \text{Gap}_i(x^t, c^t + \tilde{\tau}) &= \min_{a: x_{i,a} = 0} \sum_{f: f \in a} (c_f + \tilde{\tau}_f) - \max_{a: x_{i,a} \neq 0} \sum_{f: f \in a} (c_f + \tilde{\tau}_f) \\ &\geq \min_{a: x_{i,a} = 0} \sum_{f: f \in a} (c_f + \bar{\tau}_f) - \max_{a: x_{i,a} \neq 0} \sum_{f: f \in a} (c_f + \bar{\tau}_f) \geq 0. \end{aligned}$$

To prove that x^t is Nash equilibrium for tax $\tilde{\tau}$, we only need to show that for in-support actions $a : x_{i,a}^t \neq 0$, the action costs $\sum_{f: f \in a} (c_f^t + \tilde{\tau}_f)$ are the same. This can be derived by

$$\sum_{f: f \in a} (c_f^t + \tau_f^t) - \sum_{f: f \in a} (c_f^t + \tilde{\tau}_f) = \sum_{f \in \bar{F}_i} (\tau_f^t - \tilde{\tau}_f), \forall a : x_{i,a}^t \neq 0,$$

which is independent of a . As x^t is Nash equilibrium for tax τ^t , $\sum_{f: f \in a} (c_f^t + \tau_f^t)$ is also independent of a . \square

Lemma 10. If Algorithm 2 return False at round t , then tax τ^t is an $6F\epsilon$ -optimal tax.

Proof. For known facility f , by Lemma 8, we have $|\tau_f^*(y_f^t) - \tau_f^t(y_f^t)| \leq 3\epsilon$. By Lemma 7, for any $u \in \mathcal{K}_f^t$, we have $|\tau_f^*(u) - \tau_f^t(u)| \leq 2\epsilon$. Thus for unknown facility f , we have

$$l_f^t = \tau_f^t([y_f^t]_{\mathcal{K}_f^t}^-) + \epsilon \cdot (y_f^t - [y_f^t]_{\mathcal{K}_f^t}^-) \leq \tau_f^*([y_f^t]_{\mathcal{K}_f^t}^-) + 2\epsilon + \epsilon = \tau_f^*([y_f^t]_{\mathcal{K}_f^t}^-) + 3\epsilon,$$

$$r_f^t = \tau_f^*([y_f^t]_{\mathcal{K}_f^t \cup \{1\}}^+) + \epsilon \cdot (y_f^t - [y_f^t]_{\mathcal{K}_f^t \cup \{1\}}^+) \geq \tau_f^*([y_f^t]_{\mathcal{K}_f^t \cup \{1\}}^+) - 2\epsilon - \epsilon = \tau_f^*([y_f^t]_{\mathcal{K}_f^t \cup \{1\}}^+) - 3\epsilon,$$

As τ_f^* is nondecreasing (Assumption 2), we have

$$l_f^t - 3\epsilon \leq \tau_f^*([y_f^t]_{\mathcal{K}_f^t}^-) \leq \tau_f^*(y_f^t) \leq \tau_f^*([y_f^t]_{\mathcal{K}_f^t \cup \{1\}}^+) \leq r_f^t + 3\epsilon.$$

Thus there exists tax $\tilde{\tau}^t$ such that $\tilde{\tau}_f^t(y_f^t)$ satisfies the condition of Lemma 9 and $|\tau_f^*(y_f^t) - \tilde{\tau}_f^t(y_f^t)| \leq 3\epsilon$ for all $f \in \mathcal{F}$. x^t is the Nash equilibrium for tax $\tilde{\tau}^t$, we have

$$\forall i \in [m], a, a' \in \mathcal{A}_i, \sum_{f \in a} c_f(y_f^t) + \tilde{\tau}_f^t(y_f^t) \leq \sum_{f \in a'} c_f(y_f^t) + \tilde{\tau}_f^t(y_f^t), \text{ if } x_{i,a}^t > 0.$$

Thus we have

$$\forall i \in [m], a, a' \in \mathcal{A}_i, \sum_{f \in a} c_f(y_f^t) + \tau_f^*(y_f^t) \leq \sum_{f \in a'} c_f(y_f^t) + \tau_f^*(y_f^t) + 6F\epsilon, \text{ if } x_{i,a}^t > 0.$$

By Lemma 4, $\Psi(y_f^t) - \min_{y \in \mathcal{Y}} \Psi(y) \leq 6F\epsilon$. \square

Lemma 11. If Algorithm 2 output $\tilde{\tau}^t, \tilde{f}^t, \text{sign}^t$ at round t , then we have

$$0 < |y_{\tilde{f}^t}^t - \dot{y}_{\tilde{f}^t}^t| \leq \Delta.$$

Proof. First, we prove $|y_{\tilde{f}^t}^t - \dot{y}_{\tilde{f}^t}^t| > 0$. We consider the following two cases.

(1) Algorithm 2 return at Line 5. As we have $0 < y_i(\tilde{f}^t) < w_i$, there exists $a, a' \in \mathcal{A}_i$ such that $x_{i,a} > 0, x_{i,a'} > 0$ and $\tilde{f}^t \in a, \tilde{f}^t \notin a'$. Suppose $y_{\tilde{f}^t}^t = \dot{y}_{\tilde{f}^t}^t$. Then by Lemma 3, we have $y^t = \dot{y}^t$ as τ^t and $\dot{\tau}^t$ only differ in facility \tilde{f}^t . As a result, x^t is Nash equilibrium for tax $\dot{\tau}^t$. However, x^t is the Nash equilibrium for tax τ_f^t implies

$$\sum_{f \in a} c_f(y_f^t) + \tau_f^t(y_f^t) = \sum_{f \in a'} c_f(y_f^t) + \tau_f^t(y_f^t).$$

As τ^t and $\dot{\tau}^t$ only differ in facility \tilde{f}^t and $\tilde{f}^t \in a, \tilde{f}^t \notin a'$, we have

$$\sum_{f \in a} c_f(y_f^t) + \dot{\tau}_f^t(y_f^t) \neq \sum_{f \in a'} c_f(y_f^t) + \dot{\tau}_f^t(y_f^t),$$

which means x^t is not the Nash equilibrium for tax $\dot{\tau}$. By contradiction, we have $y_{\tilde{f}^t}^t = \dot{y}_{\tilde{f}^t}^t$.

(2) Algorithm 2 return $\tilde{\tau}^u, \tilde{f}$, sign at Line 23 or Line 31. As there exists $j \in [m]$ such that $\text{Gap}_j(x, c + \tilde{\tau}^u + \text{sign} \cdot \epsilon) < 0$, x is not a Nash equilibrium under tax τ^t . Let $\ddot{\tau}^t : \ddot{\tau}_f^t = \tau_f^t \cup (y_f^t, \tilde{\tau}_f^u)$ for $f \in \mathcal{F}$. Then $\dot{\tau}^t$ and $\ddot{\tau}^t$ only differs in \tilde{f} and x is the Nash equilibrium under tax $\ddot{\tau}^t$. By applying Lemma 3 with $\dot{\tau}^t$ and $\ddot{\tau}^t$, we have $y_{\tilde{f}^t}^t = \dot{y}_{\tilde{f}^t}^t$.

Second, we prove $|y_{\tilde{f}^t}^t - \dot{y}_{\tilde{f}^t}^t| \leq \Delta$. (1) Algorithm 2 return at Line 5. Suppose we have $|y_{\tilde{f}^t}^t - \dot{y}_{\tilde{f}^t}^t| > \Delta$. By the tax design, the (sub)gradient of the tax $(\tau_f^t)'(u) \geq \epsilon$ for $u \in [0, 1]$. As a result, $\Phi(y, c + \tau^t)$ is ϵ -strongly convex by Lemma 2. As $y^t = \arg\min_{y \in \mathcal{Y}} \Phi(y; c + \tau^t)$, we have

$$\Phi(\dot{y}^t; c + \tau^t) - \Phi(y^t; c + \tau^t) > \epsilon \Delta^2 / 2.$$

However, we have $|\Phi(y; c + \tau^t) - \Phi(y; c + \dot{\tau}^t)| \leq \delta$ for all $y \in \mathcal{Y}$. Thus we have

$$\Phi(\dot{y}^t; c + \dot{\tau}^t) - \delta \leq \Phi(y^t; c + \dot{\tau}^t) - \delta/2 \leq \Phi(y^t; c + \tau^t) \leq \Phi(\dot{y}^t; c + \tau^t) \leq \Phi(\dot{y}^t; c + \dot{\tau}^t) + \delta.$$

Comparing to the inequality above, we have $2\delta > \epsilon\Delta^2/2$, which is incorrect by the definition of δ .

By contradiction, we have $|y_f^t - \dot{y}_f^t| \leq \Delta$.

(2) Algorithm 2 return $\tilde{\tau}^u, \tilde{f}$, sign at Line 23 or Line 31. Let $\ddot{\tau}^t : \ddot{\tau}_f^t = \tau_f^t \cup (y_f^t, \tilde{\tau}_f^u)$ for $f \in \mathcal{F}$. Then x is the Nash equilibrium under tax $\ddot{\tau}^t$. Let $\ddot{\tau} : \ddot{\tau}_f^t = \tau_f^t \cup (y_f^t, \tilde{\tau}_f)$ for all $f \in \mathcal{F}$. By the definition of $\ddot{\tau}^t$ and the feasible range $\tilde{\tau}_f \in [l_f, r_f]$, the subgradient of $\ddot{\tau}_f^t$ is lower bounded by ϵ . As a result, $\Phi(\cdot; c + \ddot{\tau}^t)$ is ϵ -strongly convex on $[0, 1]$. We can prove $|y_f^t - \dot{y}_f^t| \leq \Delta$ by following the analysis for case (1) and replacing τ^t with $\ddot{\tau}^t$.

□

Proposition 1. *If Algorithm 2 return False at round t , then tax τ^t is an $6\epsilon F$ -optimal tax. If Algorithm 2 output $\tilde{\tau}^t, \tilde{f}^t, \text{sign}^t$ at round t , then we have*

$$0 < |y_{\tilde{f}^t}^t - \dot{y}_{\tilde{f}^t}^t| \leq \Delta.$$

Proof. This is directly from Lemma 10 and Lemma 11. □

Lemma 12. *Algorithm 1 return False in at most KF rounds.*

Proof. By Lemma 11 and the update rule (1), if Algorithm 2 return $\tilde{\tau}^t, f, \text{sign}$ at round t , then we will have one more known point, i.e., $\sum_{f \in \mathcal{F}} \mathcal{K}_f^{t+1} = \sum_{f \in \mathcal{F}} \mathcal{K}_f^t + 1$. As $\mathcal{K}_f^t \subseteq \mathcal{L}$ for all $f \in \mathcal{F}$ and $|\mathcal{L}| = K + 1$, we proved the lemma. □

Theorem 1. *Under Assumption 1 and Assumption 2, Algorithm 1 will output a $6\epsilon F$ tax within $T \leq 2F\beta/\epsilon$ rounds. In addition, each round has at most two tax realizations.*

Proof. The proof is directly from Proposition 1 and Lemma 12. □

C Computation Complexity

In this section, we discuss the computation complexity of Algorithm 1 and Algorithm 2. We will show that these two algorithms can be implemented with $\tilde{O}(\text{poly}(A, F, m))$ complexity for each round. For network congestion games, the computation complexity can be sharpened to $\tilde{O}(\text{poly}(V, E, m))$, avoiding the dependence on A that can be exponential in V and E .

C.1 General Congestion Games

For Algorithm 1, we compute/update the value of the cost/tax function for each facility. As we use the dictionary data structure, computing value and updating value only have $O(\log K) = O(\log \beta/\epsilon)$ complexity. As a result, the complexity of one round in Algorithm 1 is $\tilde{O}(F)$.

For Algorithm 2, $x \in \phi^{-1}(y)$ is a Caratheodory decomposition problem and can be formulated as a linear program with A variables, $F + m$ equation constraints and A inequality constraints (Proposition 2), which can be solved in polynomial time [Cohen et al., 2021].

The bottleneck is in computing $u = \arg\max_u \{u : \text{Gap}_j(x, c + \tilde{\tau}^u) \geq 0, \forall j \in [m]\}$ for $\tilde{\tau}^u : \tilde{\tau}_f^u = u, \tilde{\tau}_{\mathcal{F} \setminus \{f\}}^u = \tau_{\mathcal{F} \setminus \{f\}}'$, $f \in \bar{F}_i$. For simplicity, we use the notation: $\tilde{c}^u = c + \tilde{\tau}^u$ as the cost with tax $\tilde{\tau}^u$. By Definition 5 and the definition of action cost, we have

$$\text{Gap}_j(x, c + \tilde{\tau}^u) = \min_{a: x_{j,a}=0} \tilde{c}_a^u - \max_{a: x_{j,a} \neq 0} \tilde{c}_a^u. \quad (4)$$

For action cost \tilde{c}_a^u , if $f \in a$, it is a linear function w.r.t. u in the form of $u + C$ for some constant C . Otherwise, it is a constant w.r.t. u . As a result, we can determine the function \tilde{c}_a^u with $O(F)$ computation as we only need to compute $\tilde{c}_a^{\tau_f'}$ to decide the constant. Then we can compute $\text{Gap}_j(x, c + \tilde{\tau}^u)$ in closed form and compute $u_j = \text{argmax}_u \{u : \text{Gap}_j(x, c + \tilde{\tau}^u) \geq 0\}$ with $O(AF)$ complexity. Finally, $u = \min_{j \in [m]} u_j$ can be computed with $\mathcal{O}(mAF)$ complexity. Similarly, $u = \text{argmin}_u \{u : \text{Gap}_j(x, c + \tilde{\tau}^u) \geq 0, \forall j \in [m]\}$ has $\tilde{O}(mAF)$ computation complexity.

C.2 Network Congestion Games

For network congestion games, Algorithm 2 can be implemented by applying shortest path algorithms on a modified network, thus avoiding the dependence on A . We will apply Dijkstra's algorithm with $\tilde{O}(V + E)$ complexity while other shortest path algorithms can be used as well.

First, the Caratheodory decomposition $x \in \phi^{-1}(y)$ can be done efficiently with $O(VE + E^2)$ steps similar to the decomposition algorithm in [Panageas et al., 2023]. While their algorithm is for the flow polytopes with one commodity, it can be directly generalized to the multi-commodity case. We defer the algorithm and analysis to Appendix D.

For $u = \text{argmax}_u \{u : \text{Gap}_j(x, c + \tilde{\tau}^u) \geq 0, \forall j \in [m]\}$, the computation complexity can be boosted to $\tilde{O}(m(E + V))$. To achieve this, we consider how ④ changes as u increases from τ_f' to r_f . By Algorithm 2, we have $\text{Gap}_j(x, c + \tilde{\tau}^u) \geq 0$ when $u = \tau_f'$ as otherwise the algorithm ends at the previous iteration. In addition, facility f either has none of the Nash load or has all of the Nash load for facility j according to the algorithm design. For the first case, the in-support action costs will not change as u increases. $\text{Gap}_j(x, c + \tilde{\tau}^u) \geq 0$ always holds as the off-support action costs are nondecreasing w.r.t. u .

For the second case (all in-support actions use f), the in-support action costs take the form of $u + C$ and C can be determined by applying shortest path algorithm with edge weight $c + \tilde{\tau}^{\tau_f'}$. For off-support action cost, we observe that

$$\min_{a: x_{j,a}=0} \tilde{c}_a^u = \min \left\{ \min_{a: x_{j,a}=0, f \in a} \tilde{c}_a^u, \min_{a: x_{j,a}=0, f \notin a} \tilde{c}_a^u \right\} = \min \left\{ \min_{a: x_{j,a}=0, f \in a} \tilde{c}_a^u, \min_{a: f \notin a} \tilde{c}_a^{\tau_f'} \right\}, \quad (5)$$

where the second equation is from that the action cost does not depend on u and $x_{j,a} = 0$ if $f \notin a$. The first term in (5) grows linearly w.r.t. u as $f \in a$, so it is always larger than the in-support action cost. The second term in (5) is the shortest path length for commodity j that does not use facility f , which can be computed as the shortest path in the network after removing edge f . As a result, $u_j = \text{argmax}_u \{u : \text{Gap}_j(x, c + \tilde{\tau}^u) \geq 0\}$ can be computed with $O(E + V)$ complexity. Then the complexity for computing $u = \min_{j \in [m]} u_j$ is $\tilde{O}(m(E + V))$.

Similarly, $u_j = \text{argmin}_u \{u : \text{Gap}_j(x, c + \tilde{\tau}^u) \geq 0\}$ can be reduced to solving the shortest path that must use edge f in the network. We consider how ④ changes as u decreases from τ_f to l_f . Initially, the gap is nonnegative. If f has all of the Nash load, then the in-support action cost is a linear function $u + C$ and it decreases at least as fast as the first term. As a result, the gap is always nonnegative. Otherwise, f has none of the Nash load and in-support action costs remain constant.

We notice the following equation:

$$\min_{a: x_{j,a}=0} \tilde{c}_a^u = \min \left\{ \min_{a: x_{j,a}=0, f \in a} \tilde{c}_a^u, \min_{a: x_{j,a}=0, f \notin a} \tilde{c}_a^u \right\} = \min \left\{ \min_{a: f \in a} \tilde{c}_a^u, \min_{a: x_{j,a}, f \notin a} \tilde{c}_a^{\tau_f'} \right\}, \quad (6)$$

where the second equation is from that $f \in a$ implies a is off-support ($x_{j,a} = 0$) and $f \notin a$ implies \tilde{c}_a^u is independent of u . The second term in (6) is a constant and is always greater than the in-support action cost. The first term in (6) is a linear function u and it can be determined by computing the shortest path that always uses f and with edge weights $c + \tilde{\tau}^{\tau_f}$. This subproblem can be solved by applying the shortest path algorithm twice: the first one is to connect the source node and the starting node of f , and the second one is to connect the end node of f and the target node. As a result, the complexity for $u = \max_{j \in [m]} u_j$ is $\tilde{O}(m(E + V))$ as well. Thus the computation complexity for Algorithm 2 in network congestion games is $O(VE + E^2 + mV + mE)$.

D Missing Proofs in Section C

Proposition 2. *Finding $x \in \phi^{-1}(y)$ can be formulated as the following linear program.*

$$\begin{aligned} & \min_{x \in \mathbb{R}^A} 1 \\ \text{s.t. } & y = \sum_{i \in [m]} \sum_{a_i \in \mathcal{A}_i} x_{i,a_i} a_i \\ & w_i = \sum_{a_i \in \mathcal{A}_i} x_{i,a_i}, \forall i \in [m] \\ & x_{i,a_i} \geq 0, \forall i \in [m], a_i \in \mathcal{A}_i \end{aligned}$$

Proof. The second and third constraints guarantees x is a feasible strategy. The first constraint indicates $y = \phi(x)$. As a result, any feasible point of the program is a solution of $\phi^{-1}(y)$. \square

Algorithm 3 Efficient Computation of Flow Decomposition (Modified from [Panageas et al., 2023])

```

1: Input: A load  $y \in \mathcal{Y}$ .
2:  $x_{i,a} = 0$  for all  $i \in [m]$  and  $a \in \mathcal{A}_i$ .
3: while  $\exists f : y_f > 0$  do
4:   Let  $A = \{f : y_f > 0\}$ .
5:   Let  $f_{\min} = \operatorname{argmin}_{f \in A} y_f$  and  $y_{\min} = \min_{f \in A} y_f$ .
6:   Let  $a$  be a  $(s_i, t_i)$  path of network  $G(V, A)$  with  $f_{\min} \in p$ .
7:   Let  $x_{i,a} = y_{\min}$ ,  $y_f = y_f - y_{\min}$  if  $f \in a$ .
8: end while

```

Proposition 3. *Algorithm 3 can output a Caratheodory decomposition of y within E steps.*

Proof. During the algorithm, load y will always be nonnegative: $y_f \geq 0, \forall f \in \mathcal{F}$. For each round, we will have $y_{f_{\min}}$ reduced to 0. As a result, the algorithm will end within at most E rounds.

We only need to prove that path a always exists in Line (6) for each round. First, y always remains a multi-commodity flow as Line (7) will not affect the law of conservation in the network. By flow decomposition theorem, there exists simple paths a_1, a_2, \dots, a_p such that

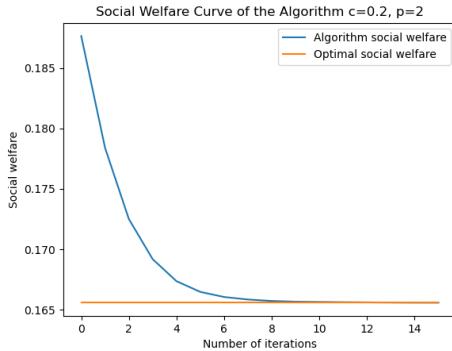
$$y = \sum_{i \in [p]} w_i a_i,$$

where $w_i > 0$ are positive flow weights. As $y_{f_{\min}} > 0$, there exists a_i such that $f_{\min} \in a_i$. Then for any $f \in a_i$, $y_f \geq w_i > 0$. As a result, path a exists for Line (6). \square

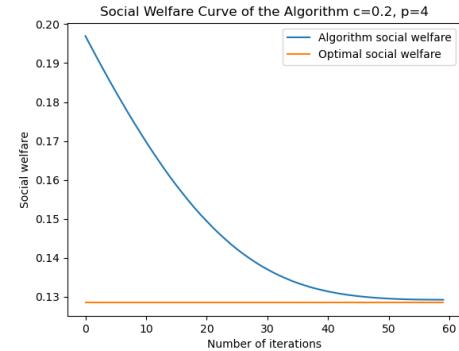
E Experiments

We implemented our algorithm and conducted experiments on a classic example known as the nonlinear variant of Pigou's example [Nisan et al., 2007]. Concretely, nonlinear variant of Pigou's example is a routing game with one source node s and one target node t . There are two edges connecting s and t . One edge has constant cost $c_0(x) = c, \forall x \in [0, 1]$ for some $c \in [0, 1]$, and the other edge has polynomial cost $c_1(x) = x^p$. One important property of such games is the price of anarchy grows without bound as $p \rightarrow \infty$, which urges proper tax to induce socially optimal behavior.

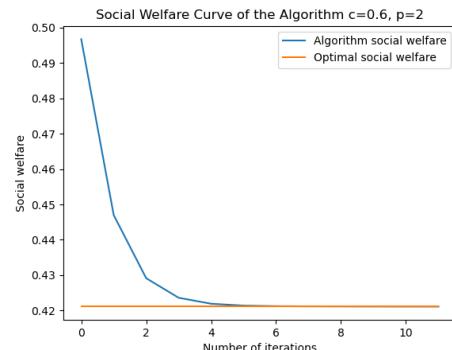
We apply our algorithm to learn the optimal tax with different c_0 and p . As we can see, the social welfare quickly converges to the optimal one. Another important observation is the learned tax function does not uniformly converge to the marginal cost tax, which is reasonable as accurate estimate is only necessary around the Nash equilibrium induced by the tax.



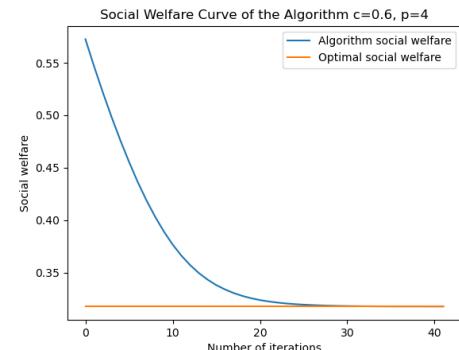
(a) $c=0.2, p=2$



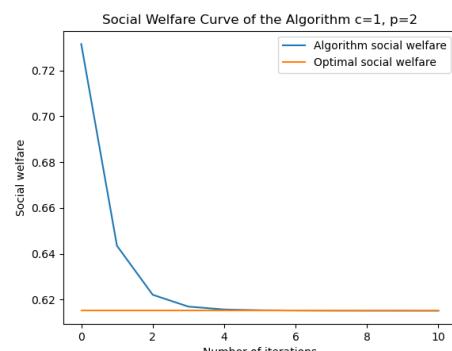
(b) $c=0.2, p=4$



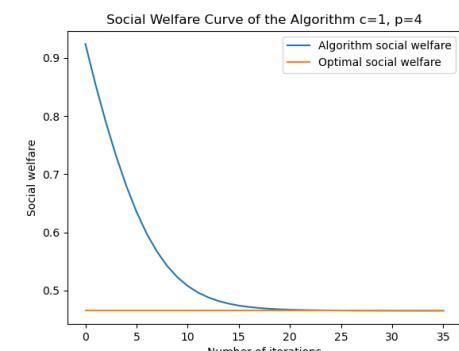
(c) $c=0.6, p=2$



(d) $c=0.6, p=4$

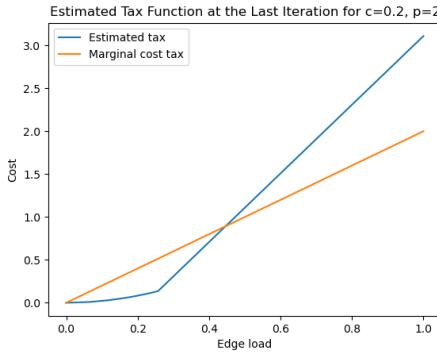


(e) $c=1, p=2$

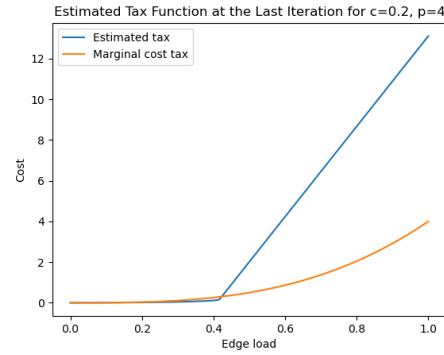


(f) $c=1, p=4$

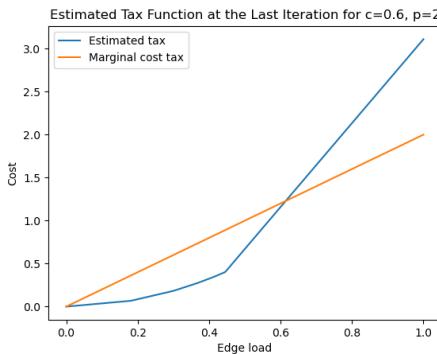
Figure 1: Social Welfare Curves of the Algorithm for various values of c and p . We can observe that the social welfare converges to the optimal one quickly.



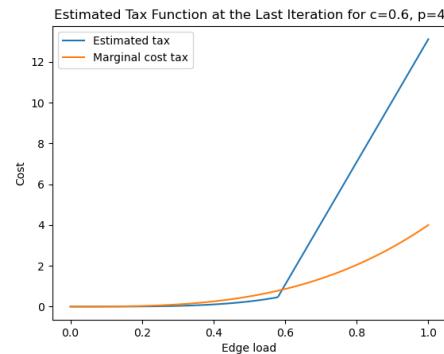
(a) $c=0.2, p=2$



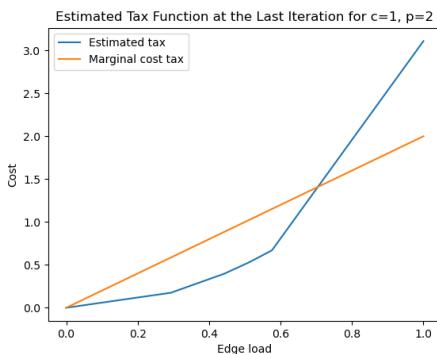
(b) $c=0.2, p=4$



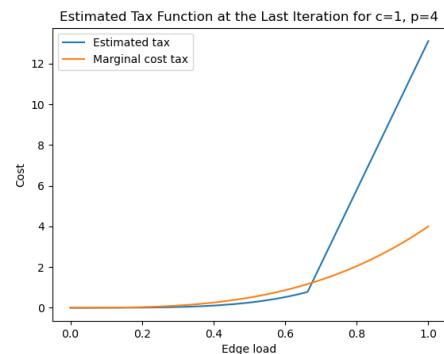
(c) $c=0.6, p=2$



(d) $c=0.6, p=4$



(e) $c=1, p=2$



(f) $c=1, p=4$

Figure 2: Estimated Tax Functions at the Last Iteration for various values of c and p . The estimation is not uniformly accurate but they are accurate at the induced Nash equilibrium.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: The abstract and introduction accurately reflect the paper's contribution.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: We discussed the limitations of our assumptions and algorithms.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[Yes\]](#)

Justification: We provide the proof in the appendix.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: **[Yes]**

Justification: Our experiment is a direct implementation of our algorithm on a well-known task and can be reproduced.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification:

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Our experiment is a direct implementation of our algorithm on a well-known task and can be reproduced.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [NA]

Justification: Our experiments are deterministic.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: Our experiment require minimal computation resources.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [\[Yes\]](#)

Justification:

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [\[Yes\]](#)

Justification: Our research can be potentially applied to real-world tax design.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We did not use public dataset.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.