

Adaptive Incentive Design

Lillian J. Ratliff and Tanner Fiez

{ratliff1,fiezt}@uw.edu

University of Washington*

Abstract

We apply control theoretic and optimization techniques to adaptively design incentives. In particular, we consider the problem of a *planner* with an objective that depends on data from strategic decision makers. The planner does not know the process by which the strategic agents make decisions. Under the assumption that the agents are utility maximizers, we model their interactions as a non-cooperative game and utilize the Nash equilibrium concept as well as myopic update rules to model the selection of their decision. By parameterizing the agents' utility functions and the incentives offered, we develop an algorithm that the planner can employ to learn the agents' decision-making processes while simultaneously designing incentives to change their response to a more desirable response from the planner's perspective. We provide convergence results for this algorithm both in the noise-free and noisy cases and present illustrative examples.

1. Introduction

Due in large part to the increasing adoption of digital technologies, many applications that once treated users as passive entities must now consider users as active participants. In many application domains, a planner or coordinator, such as a platform provider (e.g., transportation network companies), is tasked with optimizing the performance of a system that people are actively interacting with, often in real-time. For instance, the planner may want to drive the system performance to a more desirable behavior. While perhaps on competing ends of the spectrum, both revenue maximization and social welfare maximization fall under this umbrella.

A significant challenge in optimizing such an objective is the fact that human preferences are unknown *a priori* and perhaps their solicited responses, on which the system depends, may not be reported truthfully (i.e. in accordance with their true preferences) due to issues related privacy or trust.

We consider a class of incentive design problems in which a planner does not know the underlying preferences, or decision-making process, of the agents that it is trying to coordinate. In the economics literature these types of problems are known as problems of *asymmetric information*—meaning that the involved parties do not possess the same information sets and, as is often the case, one party posses some information to which the other party is not privy.

The particular type of information asymmetry which we consider, i.e. where the preferences of the agents are unknown to the planner, results in a problem of *adverse selection*. The classic example of adverse selection is the *market for lemons* [1] in which the seller of a used car knows more about the car than the buyer. There are a number of components that are hidden from the buyer such as the maintenance upkeep history, engine health, etc.

*. This work is supported by NSF Award CNS-1656873. T. Fiez was also supported in part by an NDSEG Fellowship.

Hence, the buyer could end up with a *lemon* instead of a *cherry*—i.e. a broken down piece of junk versus a sweet ride. Such problems have long been studied by economists.

The incentive design problem has also been explored by the control community, usually in the context of (reverse) Stackelberg games (see, e.g., [2–4]). More recently, dynamic incentive design in the context of applications such as the power grid [5] or network congestion games [6]. We take a slightly different view by employing techniques from learning and control to develop an adaptive method of designing incentives in a setting where repeated decisions are made by multiple, competing agents whose preferences are unknown to the designer, yet they are subjected to the incentives.

We assume that agents, including the planner, are cost minimizers¹. The decision space of the agents are assumed continuous. We model each agent’s cost as a parametric function that is dependent on the choices of other agents and is modified by an incentive chosen by the planner. The planner not knowing the underlying preferences of the agents is tantamount to it not knowing the value of the parameters of the agents’ cost functions. Such parameters can be thought of as the *type* of the agent.

We formulate an adaptive incentive design problem in which the planner iteratively learns the agents’ preferences and optimizes the incentives offered to the agents so as to drive them to a more desirable set of choices. We derive an algorithm to solve this problem and provide theoretical results on convergence for both the case when the agents play according to a Nash equilibrium as well as the case when the agents play myopically—e.g. the agents play according to a myopic update rule common in the theory of learning in games [7]. Specifically, we formulate an algorithm for iteratively estimating preferences and designing incentives. By adopting tools from adaptive control and online learning, we show that the algorithm converges under reasonable assumptions.

The results have strong ties to both the adaptive control literature [8–10] and the online learning literature [11–13]. The former gives us tools to do tracking of both the observed output (agents’ strategies) and the control input (incentive mechanism). It also allows us to go one step further and prove parameter convergence under some additional assumptions—*persistence of excitation*—on the problem formulation and, in particular, the utility learning and incentive design algorithm. The latter provides tools that allow us to generalize the algorithm and get faster convergence of the observed actions of the agents to a more desirable or even socially optimal outcome.

The remainder of the paper is organized as follows. We first introduce the problem of interest in Section 2. In Sections 3 and 4, we mathematically formulate the utility learning and incentive design problems and provide an algorithm for adaptive incentive design. We present convergence results for the Nash and myopic-play cases in Section 5 after which we draw heavily on adaptive control techniques to provide convergence results when the planner receives noisy observations. We provide illustrative numerical examples in Section 7 and conclude in Section 8.

1. While in the remainder, we formulate the entire problem given all agents are cost minimizers, the utility maximization formulation is completely analogous.

2. Problem Formulation

We consider a problem in which there is a coordinator or planner with an objective, $f_P(x, v)$, that it desires to optimize by selecting $v \in V \subset \mathbb{R}^n$; however, this objective is a function of $x = (x_1, \dots, x_n) \in X \subset \mathbb{R}^n$ which is the response of n non-cooperative strategic agents each providing response $x_i \in X_i \subset \mathbb{R}$. Regarding the dimension of each x_i , we assume without loss of generality that they are scalars. All the theoretical results and insights apply to the more general setting where each agent's choice is of arbitrary finite dimension.

The goal of the planner is to design a mechanism to *coordinate* the agents by incentivizing them choose an x that ultimately leads to minimization of $f_P(x, v)$. Yet, the coordinator does not know the decision making process by which these agents arrive at their collective response x . As a consequence there is *asymmetric information* between the agents and the planner.

Let us suppose that the each agent has some type θ and a process \mathcal{M}_i^θ that determines their *choice* $x_i(\theta_i)$. This process is dependent on the other agents and any mechanism designed by the planner. The classical approach in the economics literature is to solve this problem of so-called *adverse selection* [14] by designing mechanisms that induce agents' to take actions in a way that corresponds with their true decision-making process \mathcal{M}_i^θ . In this approach, it is assumed that the coordinator has a *prior* on the type space of the agents—e.g., a probability distribution on θ . The coordinator then designs a mechanism (usually static) based on this assumed prior that encourages agents to act in accordance with their true preferences.

We take an alternative view in which we adopt control theoretic and optimization techniques to adaptively learn the agents' types while designing incentives to coordinate the agents around a more desirable (from the point of view of the planner) choice x . Such a framework departs from one-shot decisions that assume all prior information is known at the start of the engagement and opens up opportunities for mechanisms that are dynamic and can *learn* over time.

We thus take the view that, in order to optimize its objective, the planner must learn the decision-making process and simultaneously design a mechanism that induces the agents to respond in such a way such that the planner's objective is optimized.

The planner first optimizes its objective function to find the *desired* response x^d and the *desired* v^d . That is, it determines the optimizers of its cost as if x and v are its decision variables. Of course, it may be the case that the set of optimizers of f_P contains more than one pair (x^d, v^d) ; in this case, the coordinator must choose amongst the set of optimizers. In order to realize (x^d, v^d) , the planner must incentivize the agents to play x^d by synthesizing mappings $\gamma_i \in \Gamma \subset C^2(X, \mathbb{R})$ for each $i \in \mathcal{I} = \{1, \dots, n\}$ such that $\gamma_i(x^d) = v_i^d \in V_i \subset \mathbb{R}$ and $x^d = (x_1^d, \dots, x_n^d)$ is the collective response of the agents under their true processes $\{\mathcal{M}_i^\theta\}_{i \in \mathcal{I}}$.

We will consider two scenarios: (i) Agents play according to a *Nash equilibrium* strategy; (ii) Agents play according to a *myopic* update rule—e.g. approximate gradient play or fictitious play [7].

In the first scenario, if the agents are assumed to play according to a Nash equilibrium strategy, then x^d must be a Nash equilibrium in the game induced by $\gamma = (\gamma_1, \dots, \gamma_n)$. In particular, using the notation $x_{-i} = \{x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n\}$, let agent i have *nominal*

cost $f_i \in C^2(X, \mathbb{R})$ and *incentivized* cost

$$f_i^{\gamma_i}(x_i, x_{-i}) = f_i(x_i, x_{-i}) + \gamma_i(x_i, x_{-i}). \quad (1)$$

The desired response x^d is a Nash equilibrium of the incentivized game $(f_1^{\gamma_1}, \dots, f_n^{\gamma_n})$ if

$$x_i^d \in \arg \min_{x_i \in X_i} f_i^{\gamma_i}(x_i, x_{-i}^d). \quad (2)$$

Hence, x_i^d is a *best response* to x_{-i}^d for each $i \in \mathcal{I}$. Formally, we define a Nash equilibrium as follows.

Definition 1 (Nash Equilibrium of the Incentivized Game). *A point $x \in X$ is a Nash equilibrium of the incentivized game $(f_1^{\gamma_1}, \dots, f_n^{\gamma_n})$ if*

$$f_i^{\gamma_i}(x) \leq f_i^{\gamma_i}(x'_i, x_{-i}), \quad \forall x'_i \in X_i. \quad (3)$$

If, for each $i \in \mathcal{I}$, the inequality in (3) holds only for a neighborhood $W_i \subset X_i$ of x_i , then x is a *local Nash equilibrium*.

We make use of a sub-class of Nash equilibria called *differential Nash equilibria*, as they can be characterized locally and thus, amenable to computation. Let the *differential game form* [15, Definition 2] $\omega : X \rightarrow \mathbb{R}^n$ be defined by $\omega(x) = (D_1 f_1(x), \dots, D_n f_n(x))$.

Definition 2 ([15, Definition 4]). *A strategy $x = (x_1, \dots, x_n) \in X$ is a differential Nash equilibrium of $(f_1^{\gamma_1}, \dots, f_n^{\gamma_n})$ if $\omega(x) = 0$ and $D_{ii}^2 f_i(x)$ is positive definite for each $i \in \mathcal{I}$.*

Differential Nash equilibria are known to be generic amongst local Nash equilibria [16], structurally stable and attracting under *tâtonnement* [15].

In the second scenario, we assume the agents play according to a *myopic update rule* [7] defined as follows. Given the incentive γ_i , agent i 's response is determined by the mapping

$$g_i^{\gamma_i}(x) = g_i(x) + \gamma_i(x) \quad (4)$$

In addition, function $g_i^{\gamma_i} \in C^2(X^{k+1}, \mathbb{R})$ maps the history, from time 0 up to time k , of the agents' previous collective response to the current response where X^{k+1} is the product space $X \times \dots \times X$ with $k+1$ copies of the space X .

We aim to design an algorithm in which the planner performs a *utility learning* step and an *incentive design* step such that as the planner iterates through the algorithm, agents' collective observed response converges to the desired response x^d and the value of the incentive mapping evaluated at x^d converges to the desired value v^d . In essence, we aim to ensure *asymptotic* or *approximate* incentive compatibility. In the sections that follow, we describe the utility learning and the incentive design steps of the algorithm and then, present the algorithm itself.

3. Utility Learning Formulation

We first formulate a general utility learning problem, then we give examples in the the Nash–play and myopic–play cases.

3.1 Utility Learning Under Nash–Play

We assume that the planner knows the parametric structure of the agents' nominal cost functions and receives observations of the agents' choices over time. That is, for each $i \in \mathcal{I}$, we assume that the nominal cost function of agent i has the form of a generalized linear model

$$f_i(x) = \langle \Phi(x), \theta_i^* \rangle = \sum_{j=1}^m \phi_j(x) \theta_{i,j}^* \quad (5)$$

where $\Phi(x)$ is a vector of basis functions given by $\Phi(x) = [\phi_1(x) \cdots \phi_m(x)]^T$, assumed to be known to the planner, and $\theta_i^* \in \mathbb{R}^m$ is a parameter vector, $\theta_i^* = [\theta_{i,1}^* \cdots \theta_{i,m}^*]^T$, assumed unknown to the planner.

While our theory is developed for this case, we show through simulations in Section 7 that the planner can be agnostic to the agents' decision-making processes and still drive them to the desired outcome.

Let the set of basis functions for the agents' cost functions be denoted by \mathcal{F}_ϕ . We assume that elements of \mathcal{F}_ϕ are $C^2(X, \mathbb{R})$ and *Lipschitz continuous*. Thus the derivative of any function in \mathcal{F}_ϕ is uniformly bounded.

The admissible set of parameters for agent i , denoted by Θ_i , is assumed to be a compact subset of \mathbb{R}^m and to contain the true parameter vector θ_i^* . We will use the notation $f_i(x; \theta)$ when we need to make the dependence on the parameter θ explicit.

Note that we are limiting the problem of asymmetric information to one of *adverse selection* [14] since it is the parameters of the cost functions that are unknown to the coordinator.

Similarly, we assume that the admissible incentive mappings have a generalized linear model of the form

$$\gamma_i(x) = \langle \Psi(x), \alpha_i \rangle = \sum_{j=1}^s \psi_j(x) \alpha_{i,j} \quad (6)$$

where $\Psi(x) = [\psi_1(x) \cdots \psi_s(x)]^T$ is a vector of basis functions, belonging to a finite collection \mathcal{F}_ψ , and assumed to be $C^2(X, \mathbb{R})$ and *Lipschitz continuous*, and $\alpha_i = (\alpha_{i,1} \cdots \alpha_{i,s}) \in \mathbb{R}^s$ are parameters.

Remark 1. *This framework can be generalized to use different subsets of the basis functions for different players, simply by constraining some of the parameters θ_i or α_i to be zero. We choose to present the theory with a common number of basis functions across players in an effort to minimize the amount of notation that needs to be tracked by the reader.*

At each iteration k , the planner receives the collective response from the agents, i.e. $x^k = (x_1^k, \dots, x_n^k)$, and has the incentive parameters $\alpha^k = (\alpha_1^k, \dots, \alpha_n^k)$ that were issued.

We denote the set of observations up to time $k+1$ by $\{x^t\}_{t=0}^{k+1}$ —where x^0 is the observed Nash equilibrium of the nominal game (without incentives)—and the set of incentive parameters $\{\alpha^t\}_{t=0}^k$. Each of the observations is assumed to be a Nash equilibrium.

For the incentivized game $(f_1^{\gamma_i}, \dots, f_n^{\gamma_i})$, a Nash equilibrium x necessarily satisfies the first- and second-order conditions $D_i f_i^{\gamma_i}(x) = 0$ and $D_{ii}^2 f_i^{\gamma_i}(x) \geq 0$ for each $i \in \mathcal{I}$ (see [15, Proposition 1]).

Under this model, we assume that the agents are playing a local Nash—that is, each x^k is a local Nash equilibrium so that

$$0 = D_i f_i^{\gamma_i}(x^k) = \langle D_i \Phi(x^k), \theta_i^* \rangle + \langle D_i \Psi(x^k), \alpha_i^{k-1} \rangle, \quad (7)$$

for $k \in \mathbb{N}_+$ and $0 = \langle D_i \Phi(x^0), \theta_i^* \rangle$, where $D_i \Phi(x^k) = [D_i \phi_1(x^k) \ \dots \ D_i \phi_m(x^k)]^T$ with $D_i \phi_j$ denoting the derivative of ϕ_j with respect to x_i and where we define $D_i \Psi(x^k)$ similarly. By an abuse of notation, we treat derivatives as vectors instead of co-vectors.

As noted earlier, without loss of generality, we take $x_i \in \mathbb{R}$. This makes the notation significantly simpler and the presentation of results much more clear and clean. All details for the general setting are provided in [17].

In addition, for each $i \in \mathcal{I}$, we have

$$0 \leq D_{ii}^2 f_i^{\gamma_i}(x^k) = \langle D_{ii}^2 \Phi(x^k), \theta_i^* \rangle + \langle D_{ii}^2 \Psi(x^k), \alpha_i^{k-1} \rangle \quad (8)$$

for $k \in \mathbb{N}_+$ and $0 \leq \langle D_{ii}^2 \Phi(x^0), \theta_i^* \rangle$ where $D_{ii}^2 \Phi$ and $D_{ii}^2 \Psi$ are the second derivative of Φ and Ψ , respectively, with respect to x_i .

Let the admissible set of θ_i 's at iteration k be denoted by Θ_i^k . They are defined using the second-order conditions from the assumption that the observations at times $t \in \{1, \dots, k\}$ are local Nash equilibria and are given by

$$\begin{aligned} \Theta_i^k = \{ \theta_i \in \Theta_i \mid & \langle D_{ii}^2 \Phi(x^t), \theta_i \rangle + \langle D_{ii}^2 \Psi(x^t), \alpha_i^{t-1} \rangle \geq 0, \\ & t \in \{1, \dots, k\} \} \subseteq \Theta_i. \end{aligned} \quad (9)$$

These sets are nested, i.e. $\Theta_i^k \subseteq \Theta_i^{k-1} \subseteq \dots \subseteq \Theta_i^0 \subseteq \Theta_i$, since at each iteration an additional constraint is added to the previous set. These sets are also convex since they are defined by semi-definite constraints [18]. Moreover, $\theta_i^* \in \Theta_i^k$ for all k since, by assumption, each observation x^k is a local Nash equilibrium.

Since the planner sets the incentives, given the response x^k , they can compute the quantity $-\langle D_i \Psi_i(x^k), \alpha_i^{k-1} \rangle$, which is equal to

$$-\langle D_i \Psi_i(x^k), \alpha_i^{k-1} \rangle = \langle D_i \Phi_i(x^k), \theta_i^* \rangle$$

by the first order Nash condition (7). Thus, if we let $y_i^{k+1} = -\langle D_i \Psi_i(x^{k+1}), \alpha_i^k \rangle$ and $\xi_i^k = D_i \Phi_i(x^{k+1})$, we have

$$y_i^{k+1} = \langle \xi_i^k, \theta_i^* \rangle.$$

Then, the coordinator has observations $\{y_i^t\}_{t=1}^{k+1}$ and *regression vectors* $\{\xi_i^t\}_{t=0}^k$. We use the notation $\xi^k = (\xi_1^k, \dots, \xi_n^k)$ for the regression vectors of all the agents at iteration k .

3.2 Utility Learning Under Myopic-Play

As in the Nash-play case, we assume the planner knows the parametric structure of the myopic update rule. That is to say, the nominal update function g_i is parameterized by θ_i over basis functions $\{\phi_1, \dots, \phi_m\} \subset \mathcal{F}_\phi$ and the incentive mapping γ_i^k at iteration k is parameterized by α_i^k over basis functions $\{\psi_1, \dots, \psi_s\} \subset \mathcal{F}_\psi$. We assume the planner observes the initial response x^0 and we denote the past responses up to iteration k by $x^{(0,k)} = (x^0, \dots, x^k)$. The general architecture for the myopic update rule is given by

$$x_i^{k+1} = g_i^{\gamma_i}(x_i^{(0,k)}, x_{-i}^{(0,k)}) \quad (10)$$

$$= \langle \Phi_i(x_i^{(0,k)}, x_{-i}^{(0,k)}), \theta_i^* \rangle + \langle \Psi_i(x_i^{(0,k)}, x_{-i}^{(0,k)}), \alpha_i^k \rangle. \quad (11)$$

Note that the update rule does not need to depend on the whole sequence of past response. It could depend just on the past response x^k or a subset, say $x^{(j,l)}$ for $0 \leq j \leq l \leq k$.

As before, we denote the set of admissible parameters for player i by Θ_i which we assume to be a compact subset of \mathbb{R}^m . In contrast to the Nash–play case, our admissible set of parameters is no long time varying so that $\Theta_i^k = \Theta_i$ for all k .

Keeping consistent with the notation of the previous sections, we let $\xi_i^k = \Phi_i(x^k)$ and $y_i^{k+1} = x_i^{k+1} - \langle \Psi_i(x^k), \alpha_i^k \rangle$ so that the myopic update rule can be re-written as $y_i^{k+1} = \langle \xi_i^k, \theta_i^* \rangle$. Analogous to the previous case, the coordinator has observations $\{y_i^t\}_{t=1}^{k+1}$ and regression vectors $\{\xi_i^t\}_{t=0}^k$. Again, we use the notation $\xi^k = (\xi_1^k, \dots, \xi_n^k)$ for the regression vectors of all the agents at iteration k .

Note that the form of the myopic update rule is general enough to accommodate a number of game-theoretic learning algorithms including approximate fictitious play and gradient play [7].

3.3 Unified Framework for Utility Learning

We can describe both the Nash–play and myopic–play cases in a unified framework as follows. At iteration k , the planner receives a response x^k which lives in the set $\mathcal{X}_k(x^{(0,k-1)}, \theta^*, \alpha^{k-1})$ which is either the set of local Nash equilibria of the incentivized game or the unique response determined by the incentivized myopic update rule at iteration k . The planner uses the past responses $\{x^t\}_{t=0}^k$ and incentive parameters $\{\alpha^t\}_{t=0}^k$ to generate the set of observations $\{y^t\}_{t=1}^{k+1}$ and regression vectors $\{\xi^t\}_{t=0}^k$.

The utility learning problem is formulated as an online optimization problem in which parameter updates are calculated following the gradient of a loss function. For each $i \in \mathcal{I}$, consider the loss function given by

$$\ell(\theta_i^k) = \frac{1}{2} \|y_i^{k+1} - \langle \xi_i^k, \theta_i^k \rangle\|_2^2 \quad (12)$$

that evaluates the error between the predicted observation and the true observation at time k for each player.

In order to minimize this loss, we introduce a well-known generalization of the projection operator. Denote by $\partial f(x)$ the set of *subgradients* of $f : X \rightarrow \mathbb{R}$ at x . A convex continuous function $\beta : \Theta \rightarrow \mathbb{R}$ is a *distance generating* function with modulus $\nu > 0$ with respect to a reference norm $\|\cdot\|$, if the set $\Theta^\circ = \{\theta \in \Theta | \partial\beta(\theta) \neq \emptyset\}$ is convex and restricted to Θ° , β is continuously differentiable and *strongly convex* with parameter ν , that is

$$\langle \theta' - \theta, \nabla\beta(\theta') - \nabla\beta(\theta) \rangle \geq \nu \|\theta' - \theta\|^2, \quad \forall \theta', \theta \in \Theta^\circ.$$

The function $V : \Theta^\circ \times \Theta \rightarrow \mathbb{R}_+$, defined by

$$V(\theta_1, \theta_2) = \beta(\theta_2) - (\beta(\theta_1) + \nabla\beta(\theta_1)^T(\theta_2 - \theta_1))$$

is the Bregman divergence [19] associated with β . By definition, $V(\theta_1, \cdot)$ is non-negative and strongly convex with modulus ν . Given a subset $\Theta^k \subset \Theta$ and a point $\theta \in \Theta^k$, the mapping $P_{\Theta^k, \theta} : \mathbb{R}^m \rightarrow \Theta^k$ defined by

$$P_{\Theta^k, \theta}(g) = \arg \min_{\theta' \in \Theta^k} \{ \langle g, \theta' - \theta \rangle + V(\theta, \theta') \}. \quad (13)$$

is the prox-mapping induced by V on Θ^k . This mapping is well-defined, the minimizer is unique by strong convexity of $V(\theta, \cdot)$, and is a contraction at iteration k , [20, Proposition 5.b].

Given the loss function $\ell : \Theta \rightarrow \mathbb{R}$, a positive, non-increasing sequence of learning rates (η_k) , and a distance generating function β , the parameter estimate of each θ_i is updated at iteration k as follows

$$\theta_i^{k+1} = P_{\Theta_i^k, \theta_i^k}^{k+1} \left(\eta_k \nabla \ell(\theta_i^k) \right). \quad (14)$$

Note that if the distance generating function is $\beta(\theta) = \frac{1}{2} \|\theta\|_2^2$, then the associated Bregman divergence is the Euclidean distance, $V(\theta, \theta') = \frac{1}{2} \|\theta - \theta'\|_2^2$, and the corresponding prox-mapping is the Euclidean projection on the set Θ_i , which we denote by $P_{\Theta_i, \theta}(g) = \Pi_{\Theta_i}(\theta - g)$, so that

$$\theta_i^{k+1} = \Pi_{\Theta_i^k} \left(\theta_i^k - \eta_k \nabla \ell(\theta_i^k) \right). \quad (15)$$

4. Incentive Design Formulation

In the previous section, we described the parameter update step that will be used in our utility learning and incentive design problem. We now describe how the incentive parameters α^k for each iteration are selected. In particular, at iteration k , after updating parameter estimates for each agent, the data the planner has includes the past observations $\{y^t\}_{t=1}^{k+1}$, incentive parameters $\{\alpha^t\}_{t=0}^k$, and has an estimate of each θ_i^{k+1} for $i \in \mathcal{I}$. The planner then uses the past data along with the parameter estimates to find an $\alpha^{k+1} = (\alpha_1^{k+1}, \dots, \alpha_n^{k+1})$ such that the incentive mapping for each player evaluates to v_i^d at x_i^d and $x^d \in \mathcal{X}_{k+1}(x^{(0, k+1)}, \theta^*, \alpha^{k+1})$. This is to say that if the agents are rational and play Nash, then x^d is a local Nash equilibrium of the game $(f_1^{\gamma_1}(x; \theta_1^{k+1}), \dots, f_n^{\gamma_n}(x; \theta_n^{k+1}))$ where $f_i^{\gamma_i}(x; \theta_i^{k+1})$ denotes the incentivized cost of player i parameterized by θ_i^{k+1} . On the other hand, if the agents are myopic, then, for each $i \in \mathcal{I}$, $x_i^d = \langle \Phi_i(x^k), \theta_i^{k+1} \rangle + \langle \Psi_i(x^k), \alpha_i^{k+1} \rangle$.

In the following two subsections, for each of these cases, we describe how α^{k+1} is selected.

4.1 Incentive Design: Nash–Play

Given that γ_i^{k+1} is parameterized by α_i^{k+1} , the goal is to find α_i^{k+1} for each $i \in \mathcal{I}$ such that x^d is a local Nash equilibrium of the game

$$(\langle \Phi_1(x), \theta_1^{k+1} \rangle + \langle \Psi_1(x), \alpha_1^{k+1} \rangle, \dots, \langle \Phi_n(x), \theta_n^{k+1} \rangle + \langle \Psi_n(x), \alpha_n^{k+1} \rangle)$$

and such that $\langle \Psi_i(x^d), \alpha_i^{k+1} \rangle = v_i^d$ for each $i \in \mathcal{I}$.

Assumption 1. *For every $\{\theta_i\}_{i=1}^n$ where $\theta_i \in \Theta_i$, there exist $\alpha_i \in \mathbb{R}^s$ for each $i \in \mathcal{I}$ such that x^d is the induced differential Nash equilibrium in the game $(f_1^{\gamma_1}(x; \theta_1), \dots, f_n^{\gamma_n}(x; \theta_n))$ and $\gamma_i(x^d) = v_i^d$ where $\gamma_i(x) = \langle \Psi_i(x), \alpha_i \rangle$.*

We remark that the above assumption is not restrictive in the following sense. Finding α_i that induces the desired Nash equilibrium and results in γ_i evaluating to the desired incentive value amounts to finding α_i^{k+1} such that the first- and second-order sufficient conditions for a local Nash equilibrium are satisfied given our estimate of the agents' cost

functions. That is, for each $i \in \mathcal{I}$, we need to find α_i^{k+1} satisfying

$$\begin{cases} 0_{2 \times 1} = \zeta_i^{k+1} + \Lambda_i \alpha_i^{k+1} \\ 0 < \langle D_{ii}^2 \Phi_i(x^d), \theta_i^k \rangle + \langle D_{ii}^2 \Psi_i(x^d), \alpha_i^{k+1} \rangle \end{cases} \quad (16)$$

where

$$\zeta_i^{k+1} = \begin{bmatrix} \langle D_i \Phi_i(x^d), \theta_i^{k+1} \rangle \\ -v_i^d \end{bmatrix}$$

and

$$\Lambda_i = \begin{bmatrix} D_i \Psi_i(x^d)^T \\ \Psi_i(x^d)^T \end{bmatrix} \in \mathbb{R}^{2 \times s}$$

If Λ_i is full rank, i.e. has rank 2, then there exists a α_i^{k+1} that solves the first equation in (16). If the number of basis functions s satisfies $s > 2$, then the rank condition is not unreasonable and in fact, there are multiple solutions. In essence, by selecting s to be *large enough*, the planner is allowing for enough degrees of freedom to ensure there exists a set of parameters α that induce the desired result. Moreover, the problem of finding α_i^{k+1} reduces to a convex feasibility problem.

The convex feasibility problem defined by (16) can be formulated as a constrained least-squares optimization problem. Indeed, for each $i \in \mathcal{I}$,

$$(P1) \begin{cases} \min_{\alpha_i^{k+1}} \|\zeta_i^{k+1} + \Lambda_i \alpha_i^{k+1}\|_2^2 \\ \text{s.t.} \quad \langle D_{ii}^2 \Phi_i(x^d), \theta_i^{k+1} \rangle + \langle D_{ii}^2 \Psi_i(x^d), \alpha_i^{k+1} \rangle \geq \varepsilon \end{cases}$$

for some $\varepsilon > 0$. By Assumption 1, for each $i \in \mathcal{I}$, there is an α^{k+1} such that the cost is exactly minimized.

The choice of ε determines how well-conditioned the second-order derivatives of agents' costs with respect to their own choice variables is. In addition, we note that if there are a large number of incentive basis functions, it may be reasonable to incorporate a cost for sparsity—e.g., $\lambda_i \|\alpha_i^{k+1}\|_1$; however, the optimal solution in this case is not guaranteed to satisfy (16).

It is desirable for the induced local Nash equilibrium to be a *stable, non-degenerate differential Nash equilibrium* so that it is attracting in a neighborhood under the gradient flow [15]. To enforce this, the planner must add additional constraints to the feasibility problem defined by (16). In particular, second-order conditions on player cost functions must be satisfied, i.e. that the derivative of the differential game form ω is positive-definite [15, Theorem 2]. This reduces to ensuring $D^2\Phi(x, \theta^{k+1}) + D^2\Psi(x, \alpha^{k+1}) > 0$ where

$$D^2\Phi(x, \theta^{k+1}) = \begin{bmatrix} \langle D_{11}^2 \Phi_1(x), \theta_1^{k+1} \rangle & \cdots & \langle D_{n1}^2 \Phi_1(x), \theta_1^{k+1} \rangle \\ \vdots & \ddots & \vdots \\ \langle D_{1n}^2 \Phi_n(x), \theta_n^{k+1} \rangle & \cdots & \langle D_{nn}^2 \Phi_n(x), \theta_n^{k+1} \rangle \end{bmatrix},$$

and $D^2\Psi(x, \alpha^{k+1})$ is defined analogously. Notice that this constraint is a semi-definite constraint [18] and thus, the problem of finding α^{k+1} that induces x^d to be a stable, non-degenerate differential Nash equilibrium can be formulated as a constrained least-squares

optimization problem. Indeed,

$$(P2) \begin{cases} \min_{\alpha^{k+1}} & \sum_i \|\zeta_i^{k+1} + \Lambda_i \alpha_i^{k+1}\|_2^2 \\ \text{s.t.} & D^2 \Phi(x^d, \theta^{k+1}) + D^2 \Psi(x^d, \alpha^{k+1}) \geq \varepsilon \end{cases}$$

for some $\varepsilon > 0$. The optimization problem (P2) can be written as a semi-definite program.

Again, a regularization term can be incorporated in order to find sparse parameters α_i^{k+1} . However, by introducing regularization, the condition $\zeta_i^{k+1} + \Lambda_i \alpha_i^{k+1} = 0$ will in general no longer be satisfied by the solution.

Ensuring the desired Nash equilibrium is a stable, non-degenerate differential Nash equilibrium means that, first and foremost, the desired Nash equilibrium is *isolated* [15, Theorem 2]. Thus, there is no nearby Nash equilibria. Furthermore, non-degenerate differential Nash equilibria are generic [16, Theorem 1] and structurally stable [15, Theorem 3] so that they are robust to small modeling errors and environmental noise. Stability ensures that if at each iteration players play according to a myopic approximate best response strategy (gradient play), then they will converge to the desired Nash equilibrium [15, Proposition 2]. Hence, if a stable equilibrium is desired by the planner, we can consider a modified version of Assumption 1.

Assumption 1' (Modified—Stable Differential Nash). *For every $\{\theta_i\}_{i=1}^n$ where $\theta_i \in \Theta_i$, there exist $\alpha_i \in \mathbb{R}^s$ for each $i \in \mathcal{I}$ such that x^d is the induced stable, non-degenerate differential Nash equilibrium of $(f_1^{\gamma_1}(x; \theta_1), \dots, f_n^{\gamma_n}(x; \theta_n))$ and $\gamma_i(x^d) = v_i^d$ where $\gamma_i(x) = \Psi_i(x)^T \alpha_i$.*

Given $\theta^k = (\theta_1^k, \dots, \theta_n^k)$, let $\mathcal{A}(\theta^k, x, v)$ be the set of $\alpha^k = (\alpha_1^k, \dots, \alpha_n^k)$ such that x is a differential Nash equilibrium of $(f_1^{\gamma_1}(x), \dots, f_n^{\gamma_n}(x))$ and $\gamma_i(x) = v_i$ where $\gamma_i(x) = \langle \Psi_i(x), \alpha_i^k \rangle$. Similarly, let $\mathcal{A}_s(\theta^k, x, v)$ be the set of α^k that induce x to be a stable, non-degenerate differential Nash equilibrium where $\gamma_i(x) = v_i$. By Assumptions 1 and 1', $\mathcal{A}(\theta^k, x, v)$ and $\mathcal{A}_s(\theta^k, x, v)$, respectively, are non-empty. Further, it is straightforward to find an α^k belonging to $\mathcal{A}(\theta^k, x, v)$ (resp., $\mathcal{A}_s(\theta^k, x, v)$) by solving the convex problem stated in (P1) (resp., (P2)).

4.2 Incentive Design: Myopic—Play

Given θ_i^{k+1} for each $i \in \mathcal{I}$, the planner seeks an incentive mapping $\gamma^{k+1} = (\gamma_1^{k+1}, \dots, \gamma_n^{k+1})$ that induces the desired response x^d and such that $\gamma_i^{k+1}(x^d) = v_i^d$ for each $i \in \mathcal{I}$. As before, given that γ_i has been parameterized, this amounts to finding α_i^{k+1} such that

$$x_i^d = \langle \Phi(x^k), \theta_i^{k+1} \rangle + \langle \Psi(x^k), \alpha_i^{k+1} \rangle \quad (17)$$

and such that $\langle \Psi(x^d), \alpha_i^{k+1} \rangle = v_i^d$ for each $i \in \mathcal{I}$.

Assumption 2. *For every $\{\theta_i\}_{i=1}^n$ where $\theta_i \in \Theta_i$, there exist $\alpha_i \in \mathbb{R}^{s_i}$ for each $i \in \mathcal{I}$ such that x^d is the estimated collective response—that is, (17) is satisfied for each $i \in \mathcal{I}$ —and such that $\langle \Psi(x^d), \alpha_i^{k+1} \rangle = v_i^d$.*

As in the Nash—play case, finding the incentive parameters at each iteration that induce the desired response amounts to solving a set of linear equations for each player. That is,

Algorithm 1: Adaptive Incentive Design

```

1:  $k \leftarrow 0$ 
2: If Myopic–play: receive  $x^0$ 
3: Choose  $\theta_i^0 \in \Theta_i$ ,  $\alpha_i^0 \in \mathbb{R}^s$  for each  $i \in \mathcal{I}$ 
4: Issue incentive mapping with parameters  $\alpha_i^0$ 
5: do
6:   issue incentives with parameters  $\alpha^k$ 
7:   Receive  $x^{k+1} \in \mathcal{X}_k(x^{(0,k)}, \theta^*, \alpha^k)$ .
8:   Compute  $y^{k+1}$ ,  $\xi^k$  and incur loss  $\ell(\theta_i^k)$ 
9:    $\theta_i^{k+1} = P_{\theta_i^k}^{k+1}(\eta_k \nabla \ell(\theta_i^k))$  for each  $i \in \mathcal{I}$ 
10:  for  $i \in \mathcal{I}$ :
11:    if Nash–play:
12:       $\alpha_i^{k+1} \in \mathcal{A}(\theta^{k+1}, x^d, v^d)$  [or  $\mathcal{A}_s(\theta^{k+1}, x^d, v^d)$ ]
13:    elif myopic–play:  $\alpha_i^{k+1} \in \mathcal{A}_m(\theta^{k+1}, x^d, v^d)$ 
14:     $k \leftarrow k + 1$ 
15: end do

```

for each $i \in \mathcal{I}$, the coordinator must solve

$$0_{p_i \times 1} = \tilde{\zeta}_i^{k+1} - \tilde{\Lambda}_i^k \alpha_i^{k+1} = \begin{bmatrix} x_i^d - \langle \Phi(x^k), \theta_i^{k+1} \rangle \\ v_i^d \end{bmatrix} - \begin{bmatrix} \Psi_i(x^k)^T \\ \Psi_i(x^d)^T \end{bmatrix} \alpha_i^{k+1} \quad (18)$$

for α_i^{k+1} . Define $\mathcal{A}_m(\theta^{k+1}, x^k, x^d, v^d)$ to be the set of α_i^{k+1} that satisfy (18).

The above set of equations will have a solution if the matrix $\tilde{\Lambda}_i^k$ has rank 2. Choosing the set of basis functions $\{\psi_j\}_{j=1}^s$ such that $s > 2$ makes this rank condition not unreasonable. One unfortunate difference between the Nash–play case and the present case of myopic–play is that in the former the coordinator could check the rank condition *a priori* given that it does not depend on the observations. On the other hand, $\tilde{\Lambda}_i^k$ depends on the observation at each iteration and thus, can only be verify online.

As before, the problem can be cast as a least–squares optimization problem with cost $\|\tilde{\zeta}_i^{k+1} - \tilde{\Lambda}_i^k \alpha_i^{k+1}\|_2^2$.

5. Convergence in the Noise Free Case

In Algorithm 1, the steps of the utility learning and incentive design algorithm are formalized. We now discuss the convergence results for the proposed algorithm.

Definition 3. *If, for each $i \in \mathcal{I}$, there exists a constant $0 < c_{i,s} < \infty$ such that $\xi_i^k(\xi_i^k)^T \leq c_{i,s} I$ for all k , then we say the algorithm is stable.*

Lipschitz continuity of the functions in \mathcal{F}_ϕ implies stability.

Definition 4. *If for each $i \in \mathcal{I}$, there exists a constant $0 < c_{i,p} < \infty$ such that $c_{i,p} I \leq \xi_i^k(\xi_i^k)^T$ for all k , we will say the algorithm is persistently exciting.*

Let $c_s = \max_{i \in \mathcal{I}} c_{i,s}$ and $c_p = \min_{i \in \mathcal{I}} c_{i,p}$. The following lemma is a straightforward extension of [12, Lemma 2.1] and we leave the proof to Appendix A.

Lemma 1. For every $\theta^* \in \Theta^{k+1}$, $\theta^k \in (\Theta^k)^\circ$, and $g \in \mathbb{R}^m$, we have

$$V(P_{\theta^k}^{k+1}(g), \theta^*) \leq V(\theta^k, \theta^*) + \langle g, \theta^* - \theta^k \rangle + \frac{1}{2\nu} \|g\|_*^2$$

To compactify the notation, let $V_k(\theta_i) \equiv V(\theta_i^k, \theta_i)$ and $\Delta\theta_i^k = \theta_i^* - \theta_i^k$.

Theorem 1. Suppose Algorithm 1 with prox-mapping defined by β (modulus ν) is persistently exciting, stable and that the step-size η is chosen such that $\eta - \frac{\eta^2}{2\nu} c_s > \varepsilon$ for some ε such that $0 < \varepsilon < \frac{1}{2c_p}$ with $c_p = \min_{i \in \mathcal{I}} c_{i,p}$ and where $0 < c_s < \infty$ is such that $\|\xi_i^k\|_*^2 \leq c_s$. We have the following:

(a) For each $i \in \mathcal{I}$, $V_k(\theta_i^*)$ converges and

$$\lim_{k \rightarrow \infty} \|(\xi_i^k)^T(\theta_i^* - \theta_i^k)\|^2 = 0. \quad (19)$$

(b) If $\beta(\theta_i) = \frac{1}{2}\|\theta_i\|_2^2$, then for each $i \in \mathcal{I}$, θ_i^k converges exponentially fast to θ_i^* .

Proof. We prove part (a) first. Since elements of \mathcal{F}_ϕ are Lipschitz, $\xi_i^k(\xi_i^k)^T \leq \hat{c}_s I$ so that we may find a $c_s > 0$ such that for all k , $\|\xi_i^k\|_*^2 \leq c_s$. Lemma 1 implies that

$$V_{k+1}(\theta_i^*) \leq V_k(\theta_i^*) - \eta \langle \Delta\theta_i^k, \nabla \ell(\theta_i^k) \rangle + \frac{\eta^2}{2\nu} \|\nabla \ell(\theta_i^k)\|_*^2. \quad (20)$$

Hence,

$$\begin{aligned} V_{k+1}(\theta_i^*) &\leq V_k(\theta_i^*) - \eta \|(\xi_i^k)^T(\Delta\theta_i^k)\|^2 + \frac{1}{2\nu} \eta^2 \|\xi_i^k\|_*^2 \|(\xi_i^k)^T(\Delta\theta_i^k)\|^2 \\ &\leq V_k(\theta_i^*) - \left(\eta - \frac{1}{2\nu} \eta^2 \|\xi_i^k\|_*^2\right) \|(\xi_i^k)^T(\Delta\theta_i^k)\|^2 \\ &\leq V_k(\theta_i^*) - \left(\eta - \frac{1}{2\nu} \eta^2 c_s\right) \|(\xi_i^k)^T(\Delta\theta_i^k)\|^2 \\ &\leq V_k(\theta_i^*) - \varepsilon \|(\xi_i^k)^T(\Delta\theta_i^k)\|^2 \end{aligned} \quad (21)$$

Thus, $\|(\xi_i^k)^T(\Delta\theta_i^k)\|^2 \leq \frac{1}{\varepsilon}(V_k(\theta_i^*) - V_{k+1}(\theta_i^*))$.

Summing k from 0 to K , we have

$$\sum_{k=0}^K \|(\xi_i^k)^T(\theta_i^* - \theta_i^k)\|^2 \leq \varepsilon^{-1} (V_0(\theta_i^*) - V_{K+1}(\theta_i^*)) \leq \varepsilon^{-1} V_0(\theta_i^*)$$

so that $\lim_{K \rightarrow \infty} \sum_{k=0}^K \|(\xi_i^k)^T(\theta_i^* - \theta_i^k)\|^2 \leq \varepsilon^{-1} V_0(\theta_i^*) < \infty$. This, in turn, implies that $\lim_{k \rightarrow \infty} \|\xi_i^k(\theta_i^* - \theta_i^k)\|^2 = 0$. From (21) and the fact that $V_k(\theta_i^*)$ is always positive, we see that $V_k(\theta_i^*)$ is a decreasing sequence and hence, it converges. The analysis holds for each $i \in \mathcal{I}$.

Now, we show part (b). Suppose that $\beta(\theta_i) = \frac{1}{2}\|\theta_i\|_2^2$. Then, starting with the inequality in (21), we have that

$$V_{k+1}(\theta_i^*) \leq V_k(\theta_i^*) - \varepsilon (\Delta\theta_i^k)^T \xi_i^k (\xi_i^k)^T (\Delta\theta_i^k) \leq V_k(\theta_i^*) - \varepsilon c_p (\Delta\theta_i^k)^T (\Delta\theta_i^k) \leq V_k(\theta_i^*) (1 - 2c_p \varepsilon)$$

since $\xi_i^k(\xi_i^k)^T \leq c_s I$ (i.e. stability), $\eta - \frac{1}{2}\eta^2 c_s > \varepsilon$ by construction, and $c_p I \leq \xi_i^k(\xi_i^k)^T$ (i.e. persistence of excitation). Since $0 < \varepsilon < 1/(2c_p)$, we have that $1 - 2c_p \varepsilon < e^{-2c_p \varepsilon}$ so that $V_{k+1}(\theta_i^*) < e^{-2c_p \varepsilon} V_k(\theta_i^*)$. This implies that $V_K(\theta_i^*) < e^{-2c_p K \varepsilon} V_0(\theta_i^*)$. Therefore we have that $\theta_i^k \rightarrow \theta_i^*$ exponentially fast. The same argument holds for each $i \in \mathcal{I}$. \square

For general prox-mappings, Theorem 1 lets us conclude that the observations converge to zero and that the prox-function $V_k(\theta_i)$ converges. Knowing the parameter values—a consequence of choosing $\beta(\theta_i) = \frac{1}{2}\|\theta_i\|_2^2$ —allows for the opportunity to gain qualitative insights into how agents' preferences affect the outcome of their strategic interaction. On the other hand, more general distance generating functions selected to reflect the geometry of Θ have the potential to improve convergence rates [12].

Corollary 1. *Suppose agents play according to the myopic update rule (11). Under the assumptions of Theorem 1, for each $i \in \mathcal{I}$, $\|x_i^{t+1} - x_i^d\|^2 \rightarrow 0$ and $|v_i^{t+1} - v_i^d|^2 \rightarrow 0$.*

Proof. Since agents play myopically, we have that $x_i^d = \langle \Phi(x^k), \theta_i^k \rangle + \langle \Psi(x^k), \alpha_i^k \rangle$ and the predicted induced response is $x_i^d = \langle \Phi(x^k), \theta_i^k \rangle + \langle \Psi(x^k), \alpha_i^k \rangle$ so that $\|x_i^{k+1} - x_i^d\|^2 = \|(\xi_i^k)^T(\theta_i^* - \theta_i^k)\|^2$. Thus, by Theorem 1-(a), $\lim_{k \rightarrow \infty} \|x_i^{k+1} - x_i^d\|^2 = 0$. Moreover, by Assumption 2, we know that each α_i^{k+1} satisfies $\langle \Psi(x^d), \alpha_i^{k+1} \rangle = v_i^d$. Define $v_i^{k+1} = \langle \Psi(x^{k+1}), \alpha^{k+1} \rangle$. Then,

$$\begin{aligned} |v_i^{k+1} - v_i^d|^2 &= |\langle \Psi(x^{k+1}) - \Psi(x^d), \alpha_i^{k+1} \rangle|^2 \\ &\leq \|\Psi(x^{k+1}) - \Psi(x^d)\|^2 \|\alpha_i^{k+1}\|_*^2. \end{aligned}$$

Since each element of \mathcal{F}_ψ is Lipschitz, $\|\Psi(x^{k+1}) - \Psi(x^d)\| \leq C\|x^{k+1} - x^d\|$ for some constant $C > 0$. Hence, $|v_i^{k+1} - v_i^d|^2 \leq C^2 \|\alpha_i^{k+1}\|_*^2 \|x^{k+1} - x^d\|^2$ and since $\|x^{k+1} - x^d\|^2$ converges to zero and $\|\alpha_i^{k+1}\|_*^2 < \infty$, we get that $|v_i^{k+1} - v_i^d|^2$ converges to zero. \square

In the Nash-play case, we can use the fact that non-degenerate differential Nash equilibria are structurally stable [15, Theorem 3] to determine a bound on how close an equilibrium of the incentivized game, $(f_1^{\gamma_1}(x; \theta_1^*), \dots, f_n^{\gamma_n}(x; \theta_n^*))$ with $\gamma_i(x) = \langle \Psi(x), \alpha_i^k \rangle$ for each $i \in \mathcal{I}$, is to the desired Nash equilibrium x^d . Note that the observed Nash equilibrium x^{k+1} is in the set of Nash equilibria of the incentivized game.

Let $\omega_{\mathcal{G}_\gamma}(\theta, x) = (D_1 f_1^{\gamma_1}(x; \theta_1), \dots, D_n f_n^{\gamma_n}(x; \theta))$ be the *differential game form* [15] of the incentivized game $\mathcal{G}_\gamma = (f_1^{\gamma_1}, \dots, f_n^{\gamma_n})$. By a slight abuse of notation, we will denote $D_1 \omega_{\mathcal{G}_\gamma}(\theta, x)$ and $D_2 \omega_{\mathcal{G}_\gamma}(\theta, x)$ as the local representation of the differential of $\omega_{\mathcal{G}_\gamma}$ with respect to θ and x respectively. If the parameters of the incentive mapping α^k at each iteration k are C^2 with respect to x^k and θ^k , then the differential of ω is well-defined. We remark that we formulated the optimization problem for finding the α 's as a constrained least-squares problem and there are existing results for determining when solutions to such problems are continuously dependent on parameter perturbations [21, 22].

Theorem 2. *Suppose that for each k , $\alpha^k(\theta, x) \in C^2(\mathbb{R}^{nm} \times \mathbb{R}^n, \mathbb{R}^s)$ is chosen such that x^d is a non-degenerate differential Nash equilibrium. For $\|\theta^k - \theta^*\|$ sufficiently small, there is a Nash equilibrium x^* of $\mathcal{G}_\gamma = (f_1^{\gamma_1}(x; \theta_1^*), \dots, f_n^{\gamma_n}(x; \theta_n^*))$ that is near the desired Nash equilibrium, i.e. there exists $\bar{\varepsilon} > 0$, such that for all $\theta^k \in B_{\bar{\varepsilon}}(\theta^*)$,*

$$\|x^* - x^d\| \leq (\sup_{0 \leq \lambda \leq 1} \|Dg((1 - \lambda)\theta^* + \lambda\theta^t)\|) \|\theta^k - \theta^*\|$$

where $Dg(\theta) = -(D_2 \omega_{\mathcal{G}_\gamma})^{-1}(\theta, x^d) \circ D_1 \omega_{\mathcal{G}_\gamma}(\theta, x^d)$. Furthermore, if $\|Dg(\theta)\|$ is uniformly bounded by $M > 0$ on $B_{\bar{\varepsilon}}(\theta^*)$, then $\|x^* - x^d\| \leq M \|\theta^k - \theta^*\|$.

Proof. Consider the differential game form $\omega_{\mathcal{G}_\gamma}(\theta, x)$ which is given by

$$\omega_{\mathcal{G}_\gamma}(\theta, x) = \sum_{i=1}^n \left(\sum_{j=1}^m D_i \phi_j(x) \theta_{i,j} + \sum_{k=1}^s D_i \psi_k(x) \alpha_{i,j} \right) dx_i.$$

Since x^d is a non-degenerate differential Nash equilibrium, $D_2\omega_{\mathcal{G}_\gamma}(\theta^*, x^d)$ is an isomorphism. Thus, by the Implicit Function Theorem [23, Theorem 2.5.7], there exists a neighborhood W_0 of θ^* and a C^1 function $g : W_0 \rightarrow X$ such that for all $\theta \in W_0$, $\omega(\theta, g(\theta)) = 0$. Furthermore, $Dg(\theta) = -(D_2\omega_{\mathcal{G}_\gamma})^{-1}(\theta, x^d) \circ D_1\omega_{\mathcal{G}_\gamma}(\theta, x^d)$. Let $B_{\bar{\varepsilon}}(\theta^*)$ be the largest $\bar{\varepsilon}$ -ball inside of W_0 . Since $B_{\bar{\varepsilon}}(\theta^*)$ is convex, by Proposition [23, Proposition 2.4.7], we have that

$$g(\theta^k) - g(\theta^*) = \left(\int_0^1 Dg((1-\lambda)\theta^* + \lambda\theta^k) \, d\lambda \right) \cdot (\theta^k - \theta^*)$$

Hence, since $\|x^* - x^d\| = \|g(\theta^k) - g(\theta^*)\|$, we have that

$$\|x^* - x^d\| \leq (\sup_{0 \leq \lambda \leq 1} \|Dg((1-\lambda)\theta^* + \lambda\theta^k)\|) \|\theta^k - \theta^*\|$$

Now, if $\|Dg(\theta^*)\|$ is uniformly bounded by $M > 0$ on $B_{\bar{\varepsilon}}(\theta^*)$, then it is straightforward to see from the above inequality that $\|x^* - x^d\| \leq M\|\theta^k - \theta^*\|$. \square

As a consequence of Theorem 1-(b) and Theorem 2, there exists a finite iteration k for which $\|\theta_i^* - \theta_i^k\|_2^2$ is sufficiently small for each $i \in \mathcal{I}$ so that a local Nash equilibrium of the incentivized game at time k is arbitrarily close to the desired Nash equilibrium x^d .

There may be multiple Nash equilibria of the incentivized game; hence, if the agents converge to x^* then the observed local Nash equilibrium is near the desired Nash equilibria. We know that for stable, non-degenerate differential Nash equilibria, agents will converge *locally* if following the gradient flow determined by the differential game form ω [15, Proposition 2].

Corollary 2. *Suppose the assumptions of Theorem 2 hold and that x^* is stable. If agents follow the gradient of their cost, i.e. $-D_i f_i$, then they will converge locally to x^* . Moreover, there exists an $\bar{\varepsilon} > 0$ such that for all $\theta^k \in B_{\bar{\varepsilon}}(\theta^*)$,*

$$\|x^* - x^d\| \leq (\sup_{0 \leq \lambda \leq 1} \|Dg((1-\lambda)\theta^* + \lambda\theta^k)\|) \|\theta^k - \theta^*\|$$

The proof follows directly from Theorem 2 and [15, Proposition 2].

The size of the neighborhood of initial conditions for which agents converge to the desired Nash can be approximated using techniques for computation of region of attraction via a Lyapunov function [10, Chapter 5]. This is in part due to the fact that in the case where α^k is chosen so that x^d is stable, i.e. $d\omega_{\mathcal{G}_\gamma}(\theta^k, x^d) > 0$, we have that $d\omega_{\mathcal{G}_\gamma}(\theta^*, x^*) > 0$ for θ^k near θ^* since the spectrum of $d\omega_{\mathcal{G}_\gamma}$ varies continuously.

Moreover, it is possible to explicitly construct the neighborhood W_0 obtained via the Implicit Function Theorem in Theorem 2 (see, e.g. [24, Theorem 2.9.10] or [17]).

The result of Theorem 1-(b) implies that the incentive value under x^* from Theorem 2 is arbitrarily close to the desired incentive value.

Corollary 3. *Under the assumptions of Theorem 1-(b) and Theorem 2, there exists a finite K such that for all $k \geq K$, $\|x^* - x^d\|_2^2 \leq M\bar{C}e^{-2c_p k \varepsilon}$ where $\bar{C} = n \max_i \{2V_0(\theta_i^*)\}$, $c_p = \min_{i \in \mathcal{I}} c_{i,p}$, and x^* is the (local) Nash equilibrium of the incentivized game satisfying $\|x^* - x^d\| \leq M\|\theta^k - \theta^*\|$ for all $\theta^k \in B_{\bar{\varepsilon}}(\theta^*)$. Furthermore, for each $i \in \mathcal{I}$, $|v_i^* - v_i^d|^2 \leq MC^2\|\alpha_i^k\|_2^2\bar{C}e^{-2c_p k \varepsilon}$ for all $k \geq K$ where C is the Lipschitz bound on Ψ and $v_i^* = \langle \Psi(x^*), \alpha_i^k \rangle$.*

Proof. Choose K such that, for each $i \in \mathcal{I}$, $2V_0(\theta_i^*)e^{-2\varepsilon c_p K} < \bar{\varepsilon}$ so that $\|\theta_i^k - \theta_i\|_2^2 \leq 2V_0(\theta_i^*)e^{-2\varepsilon c_p k}$ for all $k \geq K$. Thus, $\|\theta_i^k - \theta_i^*\|_2^2 \leq Ce^{-2\varepsilon c_p k}$, for all $k \geq K$. By Theorem 2, we have that $\|x^* - x^d\|_2^2 \leq M\bar{C}e^{-2\varepsilon c_p k}$ for all $k \geq K$ where M is the uniform bound on $\|Dg(\theta^*)\|$. We know that $v_i^d = \langle \Psi(x^d), \alpha_i^k \rangle$ and $v_i^* = \langle \Psi(x^*), \alpha_i^k \rangle$. Hence, since each element of \mathcal{F}_ψ is Lipschitz, we have that for all $k \geq K$, that

$$\begin{aligned} |v_i^* - v_i^d|^2 &= |\langle \Psi(x^*) - \Psi(x^d), \alpha_i^k \rangle|^2 \\ &\leq \|\Psi(x^*) - \Psi(x^d)\| \|\alpha_i^k\|_2^2 \\ &\leq C^2 \|\alpha_i^k\|_2^2 \|x^* - x^d\|_2^2 \\ &\leq MC^2 \|\alpha_i^k\|_2^2 \bar{C}e^{-2\varepsilon c_p k}. \end{aligned}$$

□

We have argued that following Algorithm 1 with a particular choice of prox-mapping, the parameter estimates of each θ_i^* converge to the true values and as a consequence we can characterize the bound on how close the observed response and incentive value are to their desired values. Knowing the true parameter values (even if obtained asymptotically) for θ^* allows the planner to make qualitative insights into the rationale behind the observed responses.

6. Convergence in the Presence of Noise

In this section, we will use the unified framework that describes both the case where the agents play according to Nash and where the agents play myopically. However, we consider *noisy* updates given by

$$y_i^{k+1} = \langle \xi_i^k, \theta_i^* \rangle + w_i^{k+1} \quad (22)$$

for each $i \in \mathcal{I}$ where w_i^{k+1} is an independent, identically distributed (i.i.d) real stochastic process defined on a probability space $(\Omega, \mathcal{F}, \mathcal{P})$ adapted to the sequence of increasing sub- σ -algebras $(\mathcal{F}_k, k \in \mathbb{N})$, where \mathcal{F}_k is the σ -algebra generated by the set $\{y_i^t, \alpha_i^t, w_i^t, t \leq k\}$ and such that the following hold²:

$$\mathbb{E}[w_i^{k+1} | \mathcal{F}_k] = 0 \quad \forall k \quad (23a)$$

$$\mathbb{E}[(w_i^{k+1})^2 | \mathcal{F}_k] = \sigma^2 > 0 \text{ a.s. } \forall k \quad (23b)$$

$$\sup_k \mathbb{E}[(w_i^{k+1})^4 | \mathcal{F}_k] < +\infty \text{ a.s.} \quad (23c)$$

Note that \mathcal{F}_k is also the σ -algebra generated by $\{y_i^t, \xi_i^t, t \leq k\}$ since w_i^k can be deduced from y_i^k and ξ_i^{k-1} through the relationship $w_i^k = y_i^k - \langle \xi_i^{k-1}, \theta_i^* \rangle$ [9].

Theorem 3. Suppose that for each $i \in \mathcal{I}$, $\{w_i^k\}$ satisfies (23a) and (23b) and that Algorithm 1, with prox-mapping P_θ associated with β (modulus ν), is persistently exciting and stable. Let the step-size η_k be selected such that $\sum_{k=1}^{\infty} \eta_k^2 < \infty$ and $\eta_k - \frac{\eta_k^2}{2\nu} \tilde{c}_s > 0$ where $0 < \tilde{c}_s < \infty$ is such that $\|\xi_i^k\|_*^2 \leq \tilde{c}_s$. Then, for each $i \in \mathcal{I}$, $V_k(\theta_i^*)$ converges a.s. Furthermore, if the sequence $\{r_k\}$ where $r_k = (\eta_k)^{-1}$ is a non-decreasing, non-negative sequence

2. We use the abbreviation a.s. for *almost surely*.

such that r_k is \mathcal{F}_k measurable and there exists constants $0 < K_1, K_2 < \infty$ and $0 < T < \infty$ such that

$$\frac{1}{k}r_{k-1} \leq K_1 + \frac{K_2}{k} \sum_{t=0}^{k-1} \|y_i^{t+1} - \langle \xi_i^t, \theta_i^t \rangle - w_i^{t+1}\|^2, \quad \forall k \geq T, \quad (24)$$

then

$$\lim_{k \rightarrow \infty} \frac{1}{k} \sum_{t=0}^{k-1} \mathbb{E}[\|w_i^{t+1} - \langle \xi_i^t, \theta_i^t - \theta_i^* \rangle\|^2 | \mathcal{F}_t] = \sigma^2 \quad a.s. \quad (25)$$

In addition, if (23c) holds, then

$$\lim_{k \rightarrow \infty} \frac{1}{k} \sum_{t=0}^{k-1} \|w_i^{t+1} - \langle \xi_i^t, \theta_i^t - \theta_i^* \rangle\|^2 = \sigma^2 \quad a.s. \quad (26)$$

The proof follows a similar technique to that presented in [9, Chapter 13.4]; hence, we leave it to Appendix A. We remark that if $\beta(\theta_i) = \frac{1}{2}\|\theta_i\|_2^2$, then Theorem 3 implies $V_k(\theta_i^*) = \frac{1}{2}\|\theta_i^* - \theta_i^k\|_2^2$ converges a.s.

Corollary 4. Suppose agents play according to the myopic update rule (11) and that the assumptions of Theorem 3 hold. Then,

$$\lim_{k \rightarrow \infty} \frac{1}{k} \sum_{t=0}^{k-1} \mathbb{E}[\|x_i^{t+1} + w_i^{t+1} - x_i^d\|^2 | \mathcal{F}_t] = \sigma^2 \quad a.s.$$

Proof. Agent i 's response is $x_i^{k+1} = \langle \xi_i^k, \theta_i^* \rangle + \langle \Psi(x^k), \alpha^k \rangle$ and the planner designs α_i^k to satisfy $x_i^d = \langle \xi_i^k, \theta_i^k \rangle + \langle \Psi(x^k), \alpha^k \rangle$. Hence, replacing $\langle \xi_i^t, \theta_i^t - \theta_i^* \rangle$ in (25) completes the proof. \square

The results of Theorem 3 imply that the average mean square error between the observations and the predictions converges to σ^2 a.s. and, if we recall, the observations are derived from noisy versions of the first-order conditions for Nash. Indeed, we have shown that

$$\lim_{k \rightarrow \infty} \frac{1}{k} \sum_{t=0}^{k-1} \mathbb{E}[\|\langle D_i \Phi(x^{t+1}), \theta_i^* \rangle + \langle D_i \Psi(x^{t+1}), \alpha_i^t \rangle + w_i^{t+1}\|^2 | \mathcal{F}_t] = \sigma^2 \quad a.s.$$

or, equivalently,

$$\lim_{k \rightarrow \infty} \frac{1}{k} \sum_{t=0}^{k-1} \mathbb{E}[\|\langle D_i \Phi(x^{t+1}), \theta_i^* - \theta_i^t \rangle + w_i^{t+1}\|^2 | \mathcal{F}_t] = \sigma^2 \quad a.s.$$

On the other hand, in the Nash-play case, it is difficult to say much about the observed Nash equilibrium except in expectation. In particular, we can consider a modified version of Theorem 2 where we consider the differential game form in expectation—i.e. at iteration k , the differential game form for the induced game is

$$\tilde{\omega}_{\mathcal{G}_\gamma}(\theta, x) = \sum_{i=1}^n \mathbb{E}[\langle D_i \Phi(x), \theta_i \rangle + \langle D_i \Psi(x), \alpha_i^k \rangle + w_i^{k+1} | \mathcal{F}_{k-1}].$$

Proposition 1. Suppose that $D_2 \tilde{\omega}_{\mathcal{G}_\gamma}(\theta, x^d)$ is an isomorphism. Then there exists an $\epsilon > 0$, such that for all $\theta^k \in B_\epsilon(\theta^*)$,

$$\|x^* - x^d\| \leq (\sup_{0 \leq \lambda \leq 1} \|Dg((1 - \lambda)\theta^* + \lambda\theta^k)\|) \|\theta^k - \theta\|$$

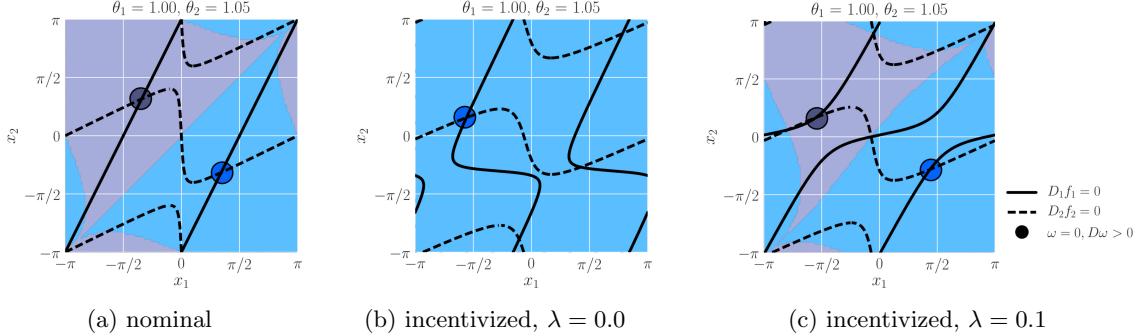


Figure 1: Visualization of Nash equilibria for the (a) nominal, (b) incentivized (without regularization), and (c) incentivized (with regularization $\lambda = 0.1$) games. Curves indicate the set of points where $D_i f_i \equiv 0$ for each $i \in \{1, 2\}$; we have $\omega = (D_1 f_1, D_2 f_2) = 0$ wherever the curves intersect. The stable, differential Nash equilibria are indicated by a dark circle of which there are two for the nominal game. Given an initial condition, Nash is calculated via a steepest descent algorithm [15]. The empirical basin of attraction for each Nash equilibrium is illustrated by the filled region containing the point. In the incentivized game, the coordinator aims to drive the agents to $(x_1^d, x_2^d) = (-1.8, 0.5)$. For $\lambda = 0.0$, the unique Nash equilibrium of the incentivized game with parameters $(\alpha_1, \alpha_2) = (0.16, -0.94, 0.29, -0.11)$ is (x_1^d, x_2^d) ; however, for the regularized incentivized game using with parameters $(\alpha_1, \alpha_2) = (0.13, 0, 0.15, 0)$, there are two stable Nash equilibria at $(1.4, -0.9)$ and $(-1.8, 0.5)$.

where $Dg(\theta^*) = -(D_2 \tilde{\omega}_{\mathcal{G}_\gamma})^{-1}(\theta^*, x^d) \circ D_1 \tilde{\omega}_{\mathcal{G}_\gamma}(\theta^*, x^d)$. and x^* is a (local) Nash equilibrium of the incentivized game $\mathcal{G}_\gamma = (f_1^{\gamma_1}(x; \theta_1^*), \dots, f_n^{\gamma_n}(x, \theta_n^*))$ with $\gamma_i(x) = \langle \Psi(x), \alpha_i^k \rangle$ for each $i \in \mathcal{I}$. Furthermore, if $\|Dg(\theta^*)\|$ is uniformly bounded by $M > 0$ on $B_\epsilon(\theta^*)$, then

$$\|x^* - x^d\| \leq M \|\theta^k - \theta^*\|.$$

To apply Proposition 1, we need a result ensuring that the parameter estimate θ^k converges to the true parameter value θ^* . One of the consequences of Theorem 3 is that $V_k(\theta_i^*)$ converges a.s. and when $\beta(\theta_i) = \frac{1}{2} \|\theta_i\|_2^2$, $\|\theta_i^* - \theta_i^k\|_2^2$ converges a.s. If it is the case that it converges a.s. to a value less than ϵ , then Proposition 1 would guarantee that a local Nash equilibrium of the incentivized game is near the desired nondegenerate differential Nash equilibrium in expectation. We leave further exploration of the convergence of the parameter estimate θ_i^k as future work.

7. Numerical Examples

In this section we present several examples to illustrate the theoretical results of the previous sections³.

7.1 Two-Player Coupled Oscillator Game

The first example we consider is a game between two players trying to control oscillators that are coupled. Coupled oscillator models are used widely for applications including

3. Code for examples can be found at github.com/fiezt/Online-Utility-Learning-Incentive-Design.

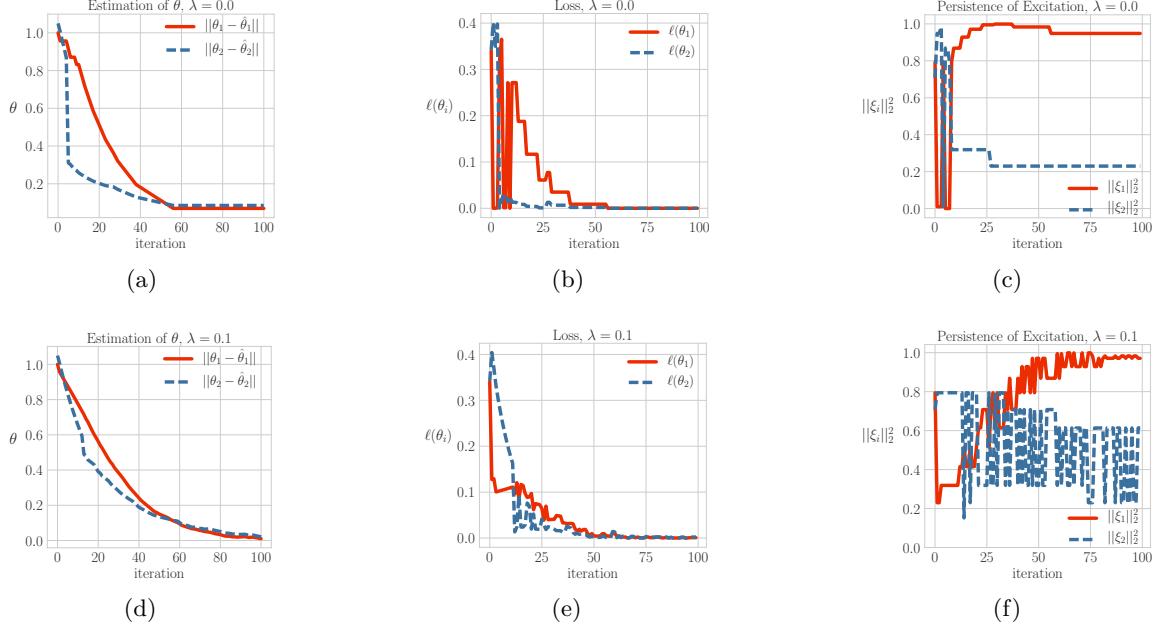


Figure 2: Estimation error for (a) $\lambda = 0.0$ and (d) $\lambda = 0.1$. Loss for (b) $\lambda = 0.0$ and (e) $\lambda = 0.1$. Persistence of excitation (as measured by $\|\xi\|^2$) for the coupled oscillator with regularization (c) $\lambda = 0.0$ and (f) $\lambda = 0.1$ on the incentive parameters.

power [25], traffic [26], and biological [27] networks, and in coordinated motion control [28] among many others. Furthermore, it is often the case that coupled oscillators are viewed in a game-theoretic context in order to gain further insight into the system properties [29, 30]. While the example is simple, it demonstrates that even in complex games on non-convex strategy spaces (in particular, smooth manifolds), Algorithm 1 still achieves the desired result.

Consider a game of coupled oscillators in which we have two players each one aiming to minimize their nominal cost

$$f_i(x_i, x_{-i}) = -\theta_i^* \cos(x_i) + \cos(x_i - x_{-i}). \quad (27)$$

Essentially, the game takes the form of a location game in which each player wants to be near the origin while also being as far away as possible from the other player's phase. The coordinator selects $\alpha_i \in \mathbb{R}^2$, $i \in \{1, 2\}$ such that (x_1^d, x_2^d) is the Nash equilibrium of the game with incentivized costs

$$f_i^\gamma(x_i, x_{-i}) = f_i(x_i, x_{-i}) + \alpha_{i,1} \sin(x_i - x_i^d) + \alpha_{i,2} \cos(x_i - x_i^d). \quad (28)$$

Using Algorithm 1, we estimate θ_1 and θ_2 while designing incentives that induce the players to play (x_1^d, x_2^d) . In this example, we set $\theta_1 = 1.0$ and $\theta_2 = 1.05$ and set the desired Nash at $(x_1^d, x_2^d) = (1.8, -0.5)$.

In Fig. 1a, we show a visualization of the Nash equilibria for the game where the square represents the unfolded torus with player 1's choice along the x-axis and player

2's choice along the y-axis. The solid and dashed black lines depict the zero sets for the derivatives of player 1 and 2 respectively. The intersections of these lines are all the candidate Nash equilibrium (all the Nash equilibria of the game must be in the set of intersection points of these lines). There are four intersection points: $\{x|\omega(x) = 0\} = \{(0, \pi), (\pi, \pi), (1.1, -1.0), (-1.1, 1.0)\}$. Of the four, the first two are not Nash equilibria which can be verified by checking the necessary conditions for Nash. However, $(0, \pi)$ is a saddle point and (π, π) is an unstable equilibrium for the dynamics $\dot{x} = -\omega(x)$. The last two points are stable differential Nash equilibria. We indicate each of them with a circle in Fig. 1a and the colored regions show the empirical basin of attraction for each of the points.

In Fig. 1b and 1c, we show similar visualizations of the Nash equilibria but for the incentivized game. In the incentive design step of the algorithm, we can add a regularization term to the constrained least squares for selecting the incentive parameters. The regularization changes the problem and may be added if it is desirable for the incentives to be small or sparse for instance. Fig. 1b and 1c are for the case without and with regularization, respectively, where we regularize with $\lambda\|\alpha\|_2^2$.

As we adjust the regularization parameter λ , it has the effect of reducing the norm of the incentive parameters and as a consequence, what happens is that the number of stable Nash equilibria changes. In particular, a larger value of λ results in small incentive values and, in turn, non-uniqueness in the number of stable Nash equilibria in the incentivized game. There is a tradeoff between uniqueness of the desired Nash equilibrium and how much the coordinator is willing to expend to incentivize the agents.

Moreover, in Fig. 2, we see that in the unregularized case ($\lambda = 0$) the estimation error for θ_2 appears not to converge to zero while in the regularized case ($\lambda = 0.1$) it does. This is likely due to the fact that Algorithm 1 is not persistently exciting for the unregularized problem (see Fig. 2c where $\|\xi_1\|_2$ drops to zero) while it is in the regularized case (see Fig. 2f where $\|\xi_i\|_2$, $i = 1, 2$ are always non-zero).

We remark that this example highlights that the technical conditions of the theoretical results in the preceding sections—in particular, persistence of excitation—matter in practice. Indeed, regularization allows for the problem to be persistently exciting. There are other methods for inducing persistence of excitation such as adding noise or ensuring the input is *sufficiently rich* [31].

7.2 Bertrand Nash Competition

In this section, we explore classical Bertrand competition. Following [32], we adopt a stylized model inspired by [33]. We consider a two firm competition. Firms are assumed to have full information and they choose their prices (x_1, x_2) to maximize their revenue $R_i(x_i, x_{-i}, \tau) = x_i F_i(x_i, x_{-i}, \tau)$ where F_i is the demand and τ is a normally distributed i.i.d. random variable that is common knowledge and represents economic indicators such as gross domestic product.

When firms compete with their nominal revenue functions R_i , the random variable τ changes over time causing *demand shocks* which, in turn, cause prices to shift. In our framework, we introduce a *regulator* that wants to push the prices chosen by firms towards some desired pricing level by introducing *incentives*. It is assumed that the regulator also has knowledge of τ .

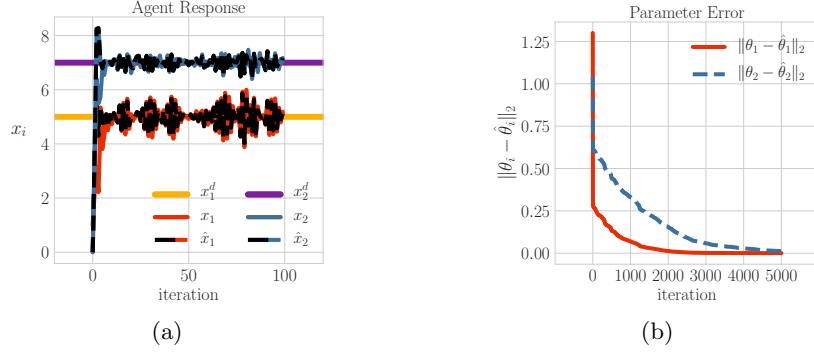


Figure 3: (a) Agent response and (b) parameter estimation error. While the agents' responses converge relatively quickly to the desired responses (within a range dictated by the statistics of the noise term τ), the parameter estimation takes quite a bit longer to converge within $1e-2$. This is in part due to the fact that the response converges quickly and the *information* that can be attained from the responses is limited; hence, it takes a while to improve the estimate.

We explore several variants of this problem to highlight the theoretical results as well as desirable empirical results we observe under relaxed assumptions on the knowledge the planner has of agent behavior.

As noted in the previous example, the assumption of persistence of excitation, which also appears in the adaptive control literature, is hard to verify *a priori*; however, selecting basis functions from certain classes such as radial basis functions which are known to be persistently exciting [34, 35] can help to overcome this issue. We note that this is not a perfect solution; it may make the algorithm persistently exciting, yet result in a solution which is not asymptotically incentive compatible, i.e. $x_i^k \not\rightarrow x_i^d$. Since we seek classes of problems that are persistently exciting and asymptotically incentive compatible, this exposes an interesting avenue for future research.

7.2.1 TRUE MODEL KNOWN TO PLANNER

We consider each firm i to have a linear marginal revenue function given by

$$M_i(x^k, \tau^k) = \theta_{i,1}^* x_1^k + \theta_{i,2}^* x_2^k + \theta_{i,i}^* x_i^k + \tau^k, \quad (29)$$

and the firm price to evolve according to a gradient play update

$$x_i^{k+1} = x_i^k + \zeta_i^k (M_i(x_i^k, x_{-i}^k, \tau^k) + \langle \Psi_i(x_i^k, x_{-i}^k), \alpha_i^k \rangle) \quad (30)$$

where $\zeta_i^k > 0$ is the learning rate which we choose to ensure stability of the process. We let the incentive basis functions be the following set of Gaussian radial basis functions:

$$\begin{cases} \Psi_{i,1}(x_i^k, x_{-i}^k) &= \exp(-\kappa(x_i^k - x_i^d)^2) \\ \Psi_{i,2}(x_i^k, x_{-i}^k) &= \exp(-\kappa(x_i^k + x_i^d)^2) \\ \Psi_{i,3}(x_i^k, x_{-i}^k) &= \exp(-\kappa(x_i^k)^2) \end{cases} \quad (31)$$

with $\kappa = 0.01$. In order to test the theoretical results, in this example we assume the regulator has knowledge of the true nominal basis functions and assumes a gradient play

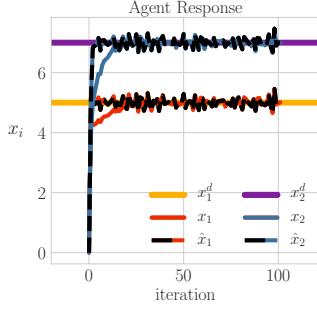


Figure 4: Prices for best response (gradient and fictitious play are similar).

model for the firms, $\Phi(x_1^k, x_2^k) = \{x_1^k, x_2^k\}$, and updates its estimate of the firm prices according to

$$x_i^{k+1} = x_i^k + \zeta_i^k (\langle \Phi_i(x_i^k, x_{-i}^k), \hat{\theta}_i^k \rangle + \tau^k + \langle \Psi_i(x_i^k, x_{-i}^k), \alpha_i^k \rangle).$$

For this simulation we use $\tau \sim \mathcal{N}(5, 0.25)$, $\theta_1^* = (-1.2, -0.5)$, and $\theta_2^* = (0.3, -1)$, and for all Bertrand examples we set the desired prices at $(x_1^d, x_2^d) = (5, 7)$.

In Fig. 3a, we show the firm prices and the regulator's estimates of the firm prices over the first 100 iterations of the simulation. The firm prices x_i almost immediately converge to the desired prices x_i^d . Likewise, the regulator's estimates of the firm prices \hat{x}_i rapidly converge to the true firm prices. We remark that in our simulations we observe that even with fewer incentive basis functions, the firm prices can still be pushed to the desired prices. Correspondingly, we find that increasing the number of incentive basis functions can increase the speed of convergence to the desired prices as well as mitigate the impact of noise due to more flexibility in the choice of incentives.

In Fig. 3b, we show that the regulator's estimates of the parameters of the marginal revenue function for each firm converge to the true parameters. While in this example the speed of estimation convergence is considerably slower than the convergence of the firm prices to the desired prices, we observe multiple factors that can significantly increase the speed of estimation convergence. Namely, as the variance of the economic indicator term τ increases, consequently providing more information at each iteration, the convergence speed also increases. Moreover, the stability of the process, dictated by the parameters of the marginal revenue functions, is positively correlated with the speed of convergence. We also note that the problem is persistently exciting and the expected loss decreases rapidly to zero.

7.2.2 PLANNER AGNOSTIC TO MYOPIC UPDATE RULE

We now transition to examples which demonstrate that our framework is suitable for problem instances that go beyond our theoretical results. In particular, we find even if the planner does not have access to the true nominal basis functions that govern the dynamics of agent play, agent responses can still be driven to a desired response for several forms of myopic updates. In these examples, we again consider firms to have the linear marginal revenue function from (29). We then allow the firm price to evolve according to either the

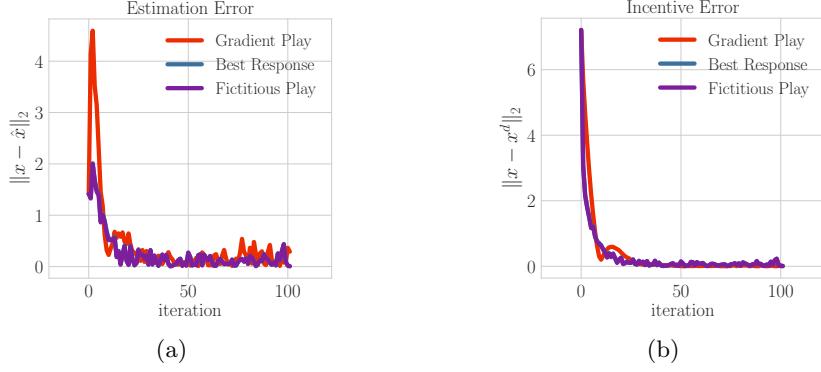


Figure 5: (a) Response estimation error for different update methods and (b) difference between desired and true response for different update methods. Note that in (a) and (b) the best response and fictitious play methods are nearly identical; this is due to the fact that fictitious play is best response to the historical average. We remark that it takes longer for fictitious play to converge given the incentive and noise at each iteration.

gradient play update from (30), a best response update given by

$$x_i^{k+1} = \arg \max_{x_i} R_i(x_i, x_{-i}^k, \tau^k) + \gamma_i^k(x_i, x_{-i}^k),$$

or a fictitious play update given by

$$x_i^{k+1} = \arg \max_{x_i} R_i(x_i, \bar{x}_{-i}^{(l,k)}, \tau^k) + \gamma_i^k(x_i, \bar{x}_{-i}^{(l,k)})$$

where $\bar{x}_{-i}^{(l,k)}$ is the average over the history from l to k with $0 \leq l \leq k$. For all simulations we choose to use the complete history, i.e. $l = 0$.

We choose the set of incentive basis functions to be same as in (31), but now consider the nominal basis functions to not include the true basis functions and instead be the following set of Gaussian radial basis functions:

$$\begin{cases} \Phi_{i,1}(x_i^k, x_{-i}^k) &= \exp(-\kappa(x_i^k)^2) \\ \Phi_{i,2}(x_i^k, x_{-i}^k) &= \exp(-\kappa(x_i^k - x_{-i}^k)^2) \\ \Phi_{i,3}(x_i^k, x_{-i}^k) &= \exp(-\kappa(x_{-i}^k)^2) \end{cases} \quad (32)$$

with $\kappa = 0.01$ as before. In this example, the regulator uses an update to its estimate of the prices, given by

$$x_i^{k+1} = \langle \Phi_i(x_i^k, x_{-i}^k), \hat{\theta}_i^k \rangle + \tau^k + \langle \Psi_i(x_i^k, x_{-i}^k), a_i^k \rangle, \quad (33)$$

which is agnostic to the firm update method. We use $\tau \sim \mathcal{N}(5, 0.1)$ and, as before, let the parameters of the marginal revenue functions be $\theta_1^* = (-1.2, -0.5)$, and $\theta_2^* = (0.3, -1)$.

In Fig. 5, we show the firm prices and the regulator's estimates of the firm prices when the firms use best response (each of the other methods have similar plots). As was the case when the regulator had knowledge of the true nominal basis functions, the regulator is able to quickly estimate the true firm prices accurately and push the firm prices to the desired

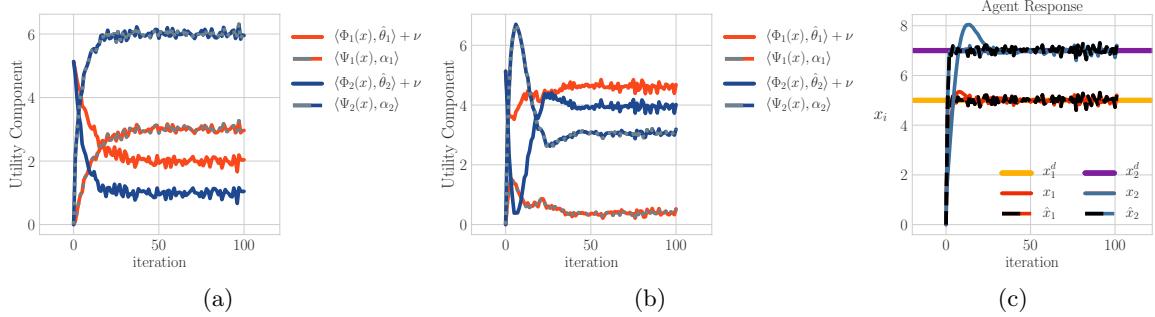


Figure 6: Estimated cost components for each firm when the firms use the best response update method: (a) linear marginal revenue and (b) non-linear marginal revenue. (c) Agent responses for gradient play with a non-linear marginal revenue functions.

prices. Figs. 5a and 5b show that the regulator's estimate of the firm prices converges to the true firm prices—i.e. $\|\hat{x}^k - x\|_2 \rightarrow 0$ —and the true firm responses converge to the desired responses—i.e. $\|x^k - x^d\|_2 \rightarrow 0$ —at nearly the same rates.

In Fig. 6a, we show the nominal and incentive cost components that form the regulator's estimate of the firm prices when the firms are use best response. Analogous to the firm prices, we note that these components are similar for each of the firms' update methods. We highlight this to point out that the nominal cost and incentive cost components do not simply cancel one another out, indicating that a portion of the true dynamics are maintained.

7.2.3 NONLINEAR MARGINAL REVENUE

To conclude our examples, we explore a Bertrand competition in which each firm has a nonlinear marginal revenue function given by

$$M_i(x^k, \tau^k) = \log(x_i^k) + \theta_{i,1}^* x_1^k + \theta_{i,2}^* x_2^k + \theta_{i,i}^* x_i^k + \theta_{i,3} + \tau^k + 1.$$

Here, we let the firm price evolve according to the gradient play update from (30) with the nonlinear marginal revenue function and let the regulator use the agnostic method to update its estimate of the firms' prices from (33) with the nominal basis functions from (32) and the incentive basis functions from (31). We use the same noise parameters as in the preceding examples and let the parameters of the marginal revenue functions be $\theta_1^* = (-1.2, -0.5, 7.5)$, and $\theta_2^* = (0.3, -1, 1.5)$.

In Fig. 6b, we show the estimated revenue and incentive components for the two firms. In Fig. 6c, we observe that the response of the firms to the regulator is marginally different from what we observed in previous examples when the firms had linear marginal revenue functions, but nonetheless the regulator's estimates of the firm prices converge to the true firm prices and the true firm prices converge to the desired prices.

8. Conclusion

We present a new method for adaptive incentive design when a planner faces competing agents of unknown type. Specifically, we provide an algorithm for learning the agents' decision-making process and updating incentives. We provide convergence guarantees on the algorithm. We show that under reasonable assumptions, the agents' true response is driven to the desired response and, under slightly more restrictive assumptions, the true preferences can be learned asymptotically. We provide several numerical examples that both verify the theory as well as demonstrate the performance when we relax the theoretical assumptions.

Appendix A. Proof of Lemma 1

The following proof uses Young's inequality:

$$v_1^T v_2 \leq \|v_1\|_* \|v_2\| \leq \frac{1}{2} \left(\frac{\|v_1\|_*^2}{\nu} + \nu \|v_2\|^2 \right) \quad (34)$$

for any $v_1, v_2 \in \mathbb{R}^m$.

Proof of Lemma 1. The proof is the essentially the same as the proof for [12, Lemma 2.1] with a few modifications.

Let $\theta^k \in (\Theta^k)^\circ$ and $\theta^{k+1} = P_{\theta^k}^{k+1}(g)$. Note that

$$\theta^{k+1} \in \operatorname{argmin}_{\theta' \in \Theta^{k+1}} \left\{ \langle g, \theta' - \theta^k \rangle + V(\theta^k, \theta') \right\} \quad (35)$$

or equivalently,

$$\theta^{k+1} \in \operatorname{argmin}_{\theta' \in \Theta^{k+1}} \left\{ \beta(\theta') - \langle \nabla \beta(\theta^k) - g, \theta' \rangle \right\} \quad (36)$$

where the latter form tells us that β is differentiable at θ^{k+1} and $\theta^{k+1} \in (\Theta^{k+1})^\circ$. Since $\nabla_2 V(\theta^k, \theta^{k+1}) = \nabla \beta(\theta^{k+1}) - \nabla \beta(\theta^k)$, the optimality conditions for (36) imply that

$$\langle \nabla \beta(\theta^{k+1}) - \nabla \beta(\theta^k) + g, \theta^{k+1} - \theta \rangle \leq 0, \quad \forall \theta \in \Theta^{k+1} \quad (37)$$

Note that this is where the proof of [12, Lemma 2.1] and the current proof are different. The above inequality holds here for all $\theta \in \Theta^{k+1}$ whereas in the proof of [12, Lemma 2.1]—using the notation of the current Lemma—the inequality would have held for all $\theta \in \Theta^k$. In particular, we need the inequality to hold for θ^* and it does since by assumption $\theta^* \in \Theta^{k+1}$ for each k .

First, $\langle \nabla \beta(\theta^k), \theta^{k+1} - \theta^k \rangle \leq \beta(\theta^{k+1}) - \beta(\theta^k)$. Hence, for $\theta^* \in \Theta^{k+1}$, we have that

$$\begin{aligned} & V(\theta^{k+1}, \theta^*) - V(\theta^k, \theta^*) \\ &= \beta(\theta^*) - \langle \nabla \beta(\theta^{k+1}), \theta^* - \theta^{k+1} \rangle - \beta(\theta^{k+1}) \\ &\quad - (\beta(\theta^*) - \langle \nabla \beta(\theta^k), \theta^* - \theta^k \rangle - \beta(\theta^k)) \\ &= \langle \nabla \beta(\theta^{k+1}) - \nabla \beta(\theta^k) + g, \theta^{k+1} - \theta^* \rangle + \langle g, \theta^* - \theta^{k+1} \rangle \\ &\leq \langle g, \theta^* - \theta^{k+1} \rangle - V(\theta^k, \theta^{k+1}) \end{aligned}$$

where the last inequality holds due to (37). By (34), we have that

$$\langle g, \theta^k - \theta^{k+1} \rangle \leq \frac{1}{2\nu} \|g\|_*^2 + \frac{\nu}{2} \|\theta^k - \theta^{k+1}\|^2. \quad (38)$$

Further, $\frac{\nu}{2} \|\theta^k - \theta^{k+1}\|^2 \leq V(\theta^k, \theta^{k+1})$ since $V(\theta^k, \cdot)$ is strongly convex. Thus,

$$\begin{aligned} & V(\theta^{k+1}, \theta^*) - V(\theta^k, \theta^*) \\ & \leq \langle g, \theta^* - \theta^{k+1} \rangle - V(\theta^k, \theta^{k+1}) \\ & = \langle g, \theta^* - \theta^k \rangle + \langle g, \theta^k - \theta^{k+1} \rangle - V(\theta^k, \theta^{k+1}) \\ & \leq \langle g, \theta^* - \theta^k \rangle + \frac{1}{2\nu} \|g\|_*^2 \end{aligned}$$

so that

$$V(P_{\theta^{(k)}}^{k+1}(g), \theta^*) \leq V(\theta^{(k)}, \theta^*) + \langle g, \theta^* - \theta^{(k)} \rangle + \frac{1}{2\nu} \|g\|_*^2.$$

□

Appendix B. Proof of Theorem 3

Proposition 2 ([36]). *Let $\{x_t\}$ be a zero conditional mean sequence of random variables adapted to $\{\mathcal{F}_t\}$. If $\sum_{t=0}^{\infty} \frac{1}{t^2} \mathbb{E}[x_t^2 | \mathcal{F}_{t-1}] < \infty$ a.s, then $\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{t=1}^N x_t = 0$ a.s.*

Proof of Theorem 3. Starting from Lemma 1, we have

$$\begin{aligned} \mathbb{E}[V_{t+1}(\theta_i^*) | \mathcal{F}_t] & \leq V_t(\theta_i^*) - \eta_t \langle \xi_i^t, \Delta\theta_i^t \rangle (\mathbb{E}[y_i^{t+1} - \langle \xi_i^t, \theta_i^t \rangle | \mathcal{F}_t]) \\ & \quad + \frac{\eta_t^2}{2\nu} \|\xi_i^t\|_*^2 ((\|\mathbb{E}[y_i^{t+1} - (\xi_i^t)^T \theta_i^t | \mathcal{F}_t]\|)^2 + \sigma^2) \\ & \leq V_t(\theta_i^*) - \left(\eta_t - \frac{\eta_t^2 \|\xi_i^t\|_*^2}{2\nu} \right) (\|\mathbb{E}[y_i^{t+1} - (\xi_i^t)^T \theta_i^t | \mathcal{F}_t]\|)^2 + \frac{\eta_t^2}{2\nu} \|\xi_i^t\|_*^2 \sigma^2 \\ & \leq V_t(\theta_i^*) - \left(\eta_t - \frac{\eta_t^2}{2\nu} \tilde{c}_1 \right) (\|\mathbb{E}[y_i^{t+1} - (\xi_i^t)^T \theta_i^t | \mathcal{F}_t]\|)^2 + \frac{\eta_t^2}{2\nu} \tilde{c}_1 \sigma^2. \end{aligned}$$

By the assumptions that $\eta_t - \frac{\eta_t^2}{2\nu} \tilde{c}_1 > 0$ and $\sum_{t=1}^{\infty} \eta_t^2 < \infty$, we can use the fact that $y_i^{t+1} = \langle \xi_i^t, \theta_i^* \rangle + w_i^{t+1}$ and apply the almost supermartingale convergence theorem [37] to get that

$$\sum_{t=1}^{\infty} \left(\eta_t - \frac{\eta_t^2}{2\nu} \tilde{c}_1 \right) (\|\mathbb{E}[w_i^{t+1} - \langle \xi_i^t, \theta_i^t - \theta_i^* \rangle | \mathcal{F}_t]\|)^2 < \infty$$

a.s. and that $V_t(\theta_i^*)$ converges a.s.

Now, we argue (25) holds; the argument follows that which is presented in [8, Chapter 8]. To do this, we first show that

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \|\langle \xi_i^t, \theta_i^t - \theta_i^* \rangle\|^2 = 0 \quad \text{a.s.} \quad (39)$$

Note that (B) implies that

$$\lim_{T \rightarrow \infty} \sum_{t=0}^{T-1} \frac{1}{r_t} \|\langle \xi_i^t, \theta_i^t - \theta_i^* \rangle\|^2 < \infty \quad \text{a.s.} \quad (40)$$

Where $r_t = (\eta_t - \frac{\eta_t^2}{2\nu} \tilde{c}_1)^{-1}$. Suppose that r_t is bounded—i.e. there exists K_3 such that $r_t < K_3 < \infty$. In this case, it is immediate from (40) that

$$\lim_{T \rightarrow \infty} \frac{1}{K_3} \sum_{t=0}^{T-1} \|\langle \xi_i^t, \theta_i^t - \theta_i^* \rangle\|^2 < \infty \quad \text{a.s.} \quad (41)$$

so that (39) follows trivially. On the other hand, suppose r_t is unbounded. Then we can apply Kronecker's Lemma [9] to conclude that

$$\lim_{T \rightarrow \infty} \frac{1}{r_T} \sum_{t=0}^{T-1} \|\langle \xi_i^t, \theta_i^t - \theta_i^* \rangle\|^2 = 0 \quad \text{a.s.} \quad (42)$$

Hence, from (24), we have

$$\lim_{T \rightarrow \infty} \frac{\frac{1}{T} \sum_{t=0}^{T-1} \|\langle \xi_i^t, \theta_i^t - \theta_i^* \rangle\|^2}{K_1 + \frac{K_2}{T} \sum_{t=0}^{T-1} \|\langle \xi_i^t, \theta_i^t - \theta_i^* \rangle\|^2} = 0 \quad \text{a.s.} \quad (43)$$

so that (39) follows immediately. Note that

$$\begin{aligned} \mathbb{E}[(y_i^{t+1} - (\xi_i^t)^T \theta_i^t)^2 | \mathcal{F}_t] &= \mathbb{E}[(y_i^{t+1} - \langle \xi_i^t, \theta_i^t \rangle - w_i^{t+1} + w_i^{t+1})^2 | \mathcal{F}_t] \\ &= \mathbb{E}[(y_i^{t+1} - \langle \xi_i^t, \theta_i^{(t)} \rangle - w_i^{t+1})^2 + (w_i^{t+1})^2 \\ &\quad + 2\langle y_i^{t+1} - \langle \xi_i^t, \theta_i^{(t)} \rangle - w_i^{t+1}, w_i^{t+1} \rangle | \mathcal{F}_t] \end{aligned}$$

Since $y_i^{t+1} - w_i^{t+1}$ and $\langle \xi_i^t, \theta_i^t \rangle$ are \mathcal{F}_t -measurable and $\mathbb{E}[w_i^{t+1} | \mathcal{F}_t] = 0$ a.s., we have

$$\mathbb{E}[(y_i^{t+1} - \langle \xi_i^t, \theta_i^t \rangle)^2 | \mathcal{F}_t] = (y_i^{t+1} - \langle \xi_i^t, \theta_i^t \rangle - w_i^{t+1})^2 - \mathbb{E}[(w_i^{t+1})^2 | \mathcal{F}_t].$$

Replacing $y_i^{t+1} = \langle \xi_i^t, \theta_i^* \rangle + w_i^{t+1}$ and using (39), we see that (25) holds since $\mathbb{E}[(w_i^{t+1})^2 | \mathcal{F}_t] = \sigma^2$ a.s.

Finally, if $\sup_t \mathbb{E}[(w_i^{t+1})^4 | \mathcal{F}_t] < +\infty$ almost surely, then by Proposition 2, we have

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \|w_i^{t+1} - \langle \xi_i^t, \theta_i^t - \theta_i^* \rangle\|^2 = \sigma^2 \quad \text{a.s.}$$

which concludes the proof. \square

References

- [1] G. A. Akerlof, "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism," *Quarterly J. Economics*, vol. 84, no. 3, pp. 488–500, 1970.
- [2] Y.-C. Ho, P. Luh, and R. Muralidharan, "Information structure, stackelberg games, and incentive controllability," *IEEE Trans. Automat. Control*, vol. 26, no. 2, pp. 454–460, Apr 1981.
- [3] Y. Ho and D. Teneketzis, "On the interactions of incentive and information structures," *IEEE Trans. Automat. Control*, vol. 29, no. 7, pp. 647–650, Jul 1984.
- [4] X. Liu and S. Zhang, "Optimal incentive strategy for leader-follower games," *IEEE Trans. Automat. Control*, vol. 37, no. 12, pp. 1957–1961, Dec 1992.
- [5] X. Zhou, E. Dall'Anese, L. Chen, and A. Simonetto, "An incentive-based online optimization framework for distribution grids," *IEEE Trans. Automat. Control*, 2017.
- [6] J. Barrera and A. Garcia, "Dynamic incentives for congestion control," *IEEE Trans. Automat. Control*, vol. 60, no. 2, pp. 299–310, Feb 2015.

- [7] D. Fudenberg and D. K. Levine, *The theory of learning in games*. MIT press, 1998, vol. 2.
- [8] G. C. Goodwin and K. S. Sin, *Adaptive filtering prediction and control*. Englewood Cliffs, NJ: Prentice-Hall, 1984.
- [9] P. R. Kumar and P. Varaiya, *Stochastic systems: estimation, identification and adaptive control*. Englewood Cliffs, NJ: Prentice-Hall, 1986.
- [10] S. Sastry, *Nonlinear Systems*. Springer New York, 1999.
- [11] N. Cesa-Bianchi and G. Lugosi, *Prediction, learning, and games*. Cambridge University Press, 2006.
- [12] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro, “Robust stochastic approximation approach to stochastic programming,” *SIAM J. Optimization*, vol. 19, no. 4, pp. 1574–1609, 2009.
- [13] M. Raginsky, A. Rakhlin, and S. Yüksel, “Online convex programming and regularization in adaptive control,” in *Proce. 49th IEEE Conf. Decision and Control*, 2010, pp. 1957–1962.
- [14] P. Bolton and M. Dewatripont, *Contract theory*. MIT press, 2005.
- [15] L. J. Ratliff, S. A. Burden, and S. S. Sastry, “On the Characterization of Local Nash Equilibria in Continuous Games,” *IEEE Trans. Automat. Control*, 2016.
- [16] ——, “Generictiy and Structural Stability of Non-Degenerate Differential Nash Equilibria,” in *Proc. 2014 Amer. Controls Conf.*, 2014.
- [17] L. J. Ratliff, “Incentivizing efficiency in societal-scale cyber-physical systems,” Ph.D. dissertation, University of California, Berkeley, 2015.
- [18] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.
- [19] L. M. Bregman, “The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming,” *USSR computational mathematics and mathematical physics*, vol. 7, no. 3, pp. 200–217, 1967.
- [20] J.-J. Moreau, “Proximité et dualité dans un espace hilbertien,” *Bulletin de la Société mathématique de France*, vol. 93, pp. 273–299, 1965.
- [21] P. Lötstedt, “Perturbation bounds for the linear least squares problem subject to linear inequality constraints,” *BIT Numerical Mathematics*, vol. 23, no. 4, pp. 500–519, 1983.
- [22] J. F. Bonnans and A. Shapiro, *Perturbation analysis of optimization problems*. Springer Science & Business Media, 2000.
- [23] R. Abraham, J. E. Marsden, and T. Ratiu, *Manifolds, Tensor Analysis, and Applications*, 2nd ed. Springer, 1988.

- [24] J. H. Hubbard and B. B. Hubbard, *Vector calculus, linear algebra, and differential forms: a unified approach*. Prentice Hall, 1998.
- [25] F. Dörfler and F. Bullo, “Synchronization and transient stability in power networks and nonuniform Kuramoto oscillators,” *SIAM J. Control Optim.*, vol. 50, no. 3, pp. 1616–1642, Jan 2012.
- [26] S. Coogan, G. Gomes, E. Kim, M. Arcak, and P. Varaiya, “Offset optimization for a network of signalized intersections via semidefinite relaxation,” in *Proc. 54th IEEE Conf. Decision and Control*, 2015.
- [27] Y. Wang, Y. Hori, S. Hara, and F. Doyle, “Collective oscillation period of inter-coupled biological negative cyclic feedback oscillators,” *IEEE Trans. Automat. Control*, vol. 60, no. 5, pp. 1392–1397, May 2015.
- [28] D. Paley, N. E. Leonard, R. Sepulchre, D. Grünbaum, J. K. Parrish *et al.*, “Oscillator models and collective motion,” *IEEE Control Syst. Mag.*, vol. 27, no. 4, pp. 89–105, 2007.
- [29] H. Yin, P. Mehta, S. Meyn, and U. Shanbhag, “Synchronization of coupled oscillators is a game,” *IEEE Trans. Automat. Control*, vol. 57, no. 4, pp. 920–935, April 2012.
- [30] T. Goto, T. Hatanaka, and M. Fujita, “Potential game theoretic attitude coordination on the circle: Synchronization and balanced circular formation,” in *IEEE Int. Symp. Intelligent Control*, Sept 2010, pp. 2314–2319.
- [31] S. Boyd and S. S. Sastry, “Necessary and sufficient conditions for parameter convergence in adaptive control,” *Automatica*, vol. 22, no. 6, pp. 629–640, 1986.
- [32] D. Bertsimas, V. Gupta, and I. C. Paschalidis, “Data-driven estimation in equilibrium using inverse optimization,” *Mathematical Programming*, vol. 153, no. 2, pp. 595–633, 2015.
- [33] S. T. Berry, “Estimating Discrete-Choice Models of Product Differentiation,” *RAND J. Economics*, vol. 25, no. 2, pp. 242–262, 1994.
- [34] A. J. Kurdila, F. J. Narcowich, and J. D. Ward, “Persistency of excitation in identification using radial basis function approximants,” *SIAM J. Control Optim.*, vol. 33, no. 2, pp. 625–642, 1995.
- [35] D. Gorinevsky, “On the persistency of excitation in radial basis function network identification of nonlinear systems,” *IEEE Tran. Neural Networks*, vol. 6, no. 5, pp. 1237–1244, 1995.
- [36] J. Neveu, *Discrete-Parameter Martingales*. North–Holland Publishing Company, 1975.
- [37] H. Robbins and D. Siegmund, “A convergence theorem for non negative almost supermartingales and some applications,” in *Herbert Robbins Selected Papers*, T. Lai and D. Siegmund, Eds. Springer New York, 1985, pp. 111–135.