Hypothesis Testing Approach to Detecting Collusion in Competitive Environments

Pedro Hespanhol
Industrial Engineering and Operations Research
University of California, Berkeley, CA
pedrohespanhol@berkeley.edu

Anil Aswani
Industrial Engineering and Operations Research
University of California, Berkeley, CA
aaswani@berkeley.edu

ABSTRACT

There is growing concern about tacit collusion using algorithmic pricing, and regulators need tools to help detect the possibility of such collusion. This paper studies how to design a hypothesis testing framework in order to decide whether agents are behaving competitively or not. In our setting, agents are utility-maximizing and compete over prices of items. A regulator, with no knowledge of the agent's utility function, has access only to the agents' strategies (i.e., pricing decisions) and external shock values in order to decide if agents are behaving in competition according to some equilibrium problem. We leverage the formulation of such a problem as an inverse variational inequality and design a hypothesis test under a minimal set of assumptions. We demonstrate our method with computational experiments of the Bertrand competition game (with and without collusion) and show how our method performs.

KEYWORDS

hypothesis testing, collusion, algorithmic pricing, competition

1 INTRODUCTION

Algorithmic pricing [10, 13] is increasingly used due to the growth of internet sales channels, but there is concern that these algorithms will lead to tacit collusion that harms consumers [11, 12, 15, 31]. The current situation is unique in that, though it poses challenges for regulators because of the difficulty in detecting tacit collusion by algorithms, there is a large amount of real-time pricing and purchase data available for analysis by regulators. This paper is motivated by these pressing issues, and represents an initial step towards answering the question of how a regulator could detect algorithmic collusion from a large corpus of financial data.

The most closely related literature looks at collusion in auctions [6, 7, 26, 27]. Some work [26, 27] conducted a statistical analysis of bidding data from situations where collusion is known to have occurred, and found that collusion leads to less aggressive bidding, higher prices for consumers, and increased correlation in bids. Other work [6, 7, 27] uses econometrics to detect collusion, and is derived using analyses that assume (perfect) equilibrium behavior. However, assuming perfect equilibrium behavior is too stringent because it requires no model mismatch between the econometrics method and the bidders, and it requires bidders to be exactly accurate in the optimality of their bids. So these approaches generally lead to too many false positives when trying to detect collusion.

This paper proposes a hypothesis testing framework for detecting collusion between agents in competitive environments. The significant difference in our work is we allow the agents to not be in perfect equilibrium. Instead, we presuppose that agents typically choose actions close (but not exactly equal) to equilibrium but that they will also occasionally choose actions far from equilibrium. This weaker assumption partially mitigates model mismatch because it does not require any data without collusion to exactly match the equilibrium, and it also eliminates the need for agents to be exactly accurate in the optimality of their strategies. These ideas will be made more precise when we present our mathematical model.

1.1 Estimation in Equilibrium Models

Many competitive and cooperative environments can be analyzed using equilibrium models, where concepts like the Nash equilibrium are used to study agents' strategic behavior [17, 25]. Those models typically contain primitives – such as agents' private information, utility functions, and strategy spaces – that often are not known to an outside regulator or designer who then needs to estimate those elements in order to, for example, design a mechanism of interaction in order to induce a particular behavior or outcome. In this paper we use *strategy space* and *action space* interchangeably.

However, the estimation of such primitives in equilibrium problems is quite challenging. In this work, we consider the case where the agents' actions are observed by the regulator and their strategy spaces are known. Hence the estimation problem lies solely on the agents' utility functions and personal information. One line of work is of structural estimation methods [2, 5, 28], which estimate parametric utility functions by observing agents acting in equilibrium. Those approaches assume a "ground-truth" form of the utility function, and derive necessary conditions based on constrained optimization in order to formally derive estimators of the parameters. Another related line of work are surrogate methods [16, 21, 34] that, while not strictly estimation methods, elicit information from the agents themselves about sensitivities over strategies (i.e., about derivatives of their utility functions with respect to strategies) by providing the agents with a common surrogate function. Those methods can induce the appropriate equilibrium behavior even though the agents' utility functions are not known.

Lastly, estimation can be formulated as an inverse variational inequality [9, 20]. Such methods use the fact that equilibrium problems are a special case of variational inequalities, in order to pose an inverse optimization problem [1] where the solutions (i.e., the equilibrium strategies) are provided via samples, and the goal is to estimate the problem's parameters that generate such solutions. This approach is powerful as it does not require a "ground-truth" model, and under some conditions can produce good approximate equilibrium behavior even when the parametric form of the utility functions is misspecified. However, these methods based on inverse optimization can encounter difficulties when the data gathered is noisy [3, 4], which is what happens in most applications.

In this paper, we build on the framework of [9]. However, we will not assume that equilibrium actions are provided to the regulator via samples. Instead, we develop a method that given some arbitrary samples of actions is able to identify whether these samples came from agents acting in (approximate) equilibrium or not. Hence our work is more tied to coalition-detection in equilibrium games, in the field of economics, and to hypothesis testing in statistics.

1.2 Coalition Detection

It is common to assume agents behave in equilibrium; however this is not observed in several applications [8]. In fact, agents often collude, exchange information, and form sub-groups – instead of cooperating as a whole [19, 33]. This poses a problem for estimation methods, which often assume the observed strategies come from an specific type of behavior (e.g., total cooperation or competition). The work [29, 30] characterizes coalitions in games and provides conditions where the establishment or not of coalitions can be tested. However the problem of coalition formation can also appear where agents play imperfectly and are learning the primitives of the environment while acting on it. Work done by [22] gives evidence that when agents employ learning algorithms in a competitive environment, such algorithms "learn" to implicitly collude, even if collusion is not part of the agents' plans.

In this paper, we formulate the problem of identifying whether or not the observed actions come from agents in total competition or not. Instead of identifying the coalition itself based on structural properties of the game, we instead use a data-driven approach where a Kolmogorov-Smirnov hypothesis test [23, 24] is used in order to accept or reject the hypothesis that the strategies observed by the regulator come from agents in equilibrium. This novel approach is powerful in the sense that is independent of the type of utility functions considered and makes very mild assumptions on the *a priori* behavior of the agents, leveraging both the variational inequality formulation and the power of hypothesis testing.

1.3 Outline

In Section 2, we describe our (approximate) equilibrium model and develop the corresponding hypothesis testing framework. In Section 3, we present computational experiments based on a constrained Bertrand-type game to showcase the performance of our method when agents are acting in competition and in collusion. Lastly, we conclude in Section 4 with a discussion on the potential future applications of our model and method. We also discuss open theoretical questions that are motivated by our work in this paper.

2 PROBLEM SETTING

We describe the agents' model, the estimation setting the regulator faces as it observes the actions taken by agents, and present our hypothesis testing framework for detecting collusion.

2.1 Model Description

Suppose two agents are engaged in a Bertrand-type game, competing over prices of certain items (e.g., products or airline tickets). Let $p_i \in \mathbb{R}^q$ be the price vectors of agent $i \in \{1, 2\}$ over $\{1, \ldots, q\}$ items, and note we use the terms *prices* and *strategies* interchangeably, as the strategy of each agent consists solely on the prices over the items. We assume the strategy space of agent i is

$$\mathcal{P}_i = \{p \in \mathbb{R}^q : Ap = b, p \in K\} \tag{1}$$

where A is a $m \times q$ matrix, b is a m-vector, and K is a closed convex cone. Hence the strategy space is a cone given in standard form. In addition, each agent has their own utility function

$$\mathbf{U}_i(p_1, p_2, \mu; \theta_i) = p_i D_i(p_1, p_2, \mu; \theta_i) \tag{2}$$

where $D_i(p_1, p_2, \mu; \theta_i)$ is the agent's demand function, which depends on both price vectors and is parametrized by the vector $\theta_i \in \mathbb{R}^d$, which we assume to be the private information of each agent. Lastly, the demand function also depends on μ, which is a common shock value disturbance that we assume to be a bounded disturbance with small magnitude (to be made precise in next section). The goal of this shock vector is to represent uncertainties that may affect the demand and are outside of the agents' control.

The goal of each agent is to select a feasible price $p_i \in \mathcal{P}_i$ such that its utility function is maximized. We focus our analysis to the Nash equilibrium of the resulting game:

Definition 2.1. A strategy profile (p_1^*, p_2^*) is a Nash equilibrium if each agent plays the "best-response" to the other, namely if

$$p_i^* \in \arg\max_{p_i \in \mathcal{P}_i} \mathbf{U}_i(p_1, p_2^*, \mu; \theta_i), \text{for } i \in \{1, 2\} \tag{3}$$

The (pure-strategy) Nash Equilibrium may not be an adequate solution concept for this constrained environment of the Bertrand-type game, as it may not exist [14]. However, we will focus on the case where each agent plays imperfectly. Namely, we assume that before playing the game, a gap value ϵ is sampled from a (known to the agents and regulator) parametric distribution $\mathcal{D}(\phi)$ with (unknown to the regulator but known to the agents) parameter ϕ. Next, both agents pick a strategy vector (p_1, p_2) that is an ϵ-approximate Nash equilibrium of the Bertrand game. To formalize this notion of approximate equilibrium, we will use the characterization based on variational inequalities presented in [9].

Definition 2.2. Given a function $\mathbf{f} : \mathbb{R}^q \to: \mathbb{R}^q$ and a non-empty set $\mathcal{F} \in \mathbb{R}^q$, the problem of finding the point p^* such that

$$\mathbf{f}(p^*)^\top (p - p^*) \geq 0, \text{for } p \in \mathcal{F} \tag{4}$$

is called the variational inequality problem $VI(\mathbf{f}, \mathcal{F})$.

It turns out that several problems can be formulated as variational inequalities (We refer to [20] for an in-depth characterization). In particular, if we let $\mathcal{F} = \mathcal{P}_1 \times \mathcal{P}_2$ and we let

$$\mathbf{f}(p) = \begin{bmatrix} \mathbf{f}_1(p_1, p_2) \\ \mathbf{f}_2(p_1, p_2) \end{bmatrix} = \begin{bmatrix} -\nabla_1 U_1(p_1, p_2, \mu; \theta_1) \\ -\nabla_2 U_2(p_1, p_2, \mu; \theta_2) \end{bmatrix} \tag{5}$$

where ∇_i is the gradient w.r.t. p_i, then solving $VI(\mathbf{f}, \mathcal{F})$ from (4) is equivalent to finding the Nash equilibrium (3). We thus have the following definition for approximate Nash equilibrium:

Definition 2.3. A strategy profile $(\bar{p}_1, \bar{p}_2)$ is an ϵ-approximate Nash equilibrium if and only if

$$\mathbf{f}(\bar{p})^\top (p - \bar{p}) \geq -\epsilon, \text{for } p \in \mathcal{F}. \tag{6}$$

It will suit our purposes to formulate the above approximate variational inequality problem as a (convex) optimization problem. This can be done under technical regularity conditions that ensure constraint qualification holds (e.g., Slater's condition).

Theorem 2.4. *[9] Let $\mathcal{F} = \mathcal{P}_1 \times \mathcal{P}_2$, where $\mathcal{P}_i$ is given by (1), for $i \in \{1, 2\}$. Let $\boldsymbol{f}$ be given by (5). If $\mathcal{F}$ satisfies constraint qualification (e.g., Slater's condition), then a strategy profile $(\bar{p}_1, \bar{p}_2)$ is an ϵ-approximate Nash equilibrium if and only if*

$$\exists y_1, y_2 \in \mathbb{R}^m : \begin{cases} A_i^\top y_i \leq_C \boldsymbol{f}_i(\bar{p}_1, \bar{p}_2), \text{for } i \in \{1, 2\} \\ \sum_{i=1}^2 \boldsymbol{f}_i(\bar{p}_1, \bar{p}_2)^\top \bar{p}_i - b_i^\top y_i \leq \epsilon \end{cases} \tag{7}$$

where we use the symbol "$\leq_C$" to denote conic inequalities.

Next we assume that given some $\epsilon \sim \mathcal{D}(\phi)$, the agents solve the above feasibility problem in order to select the prices. In particular, we assume that the agents solve the above problem where the second inequality is replaced by an equality constraint – that is the selected strategies satisfy condition (6) with equality. We will not focus on how such prices are achieved, that is, how the agents learn to play the ϵ-approximate Nash Equilibrium strategies (we refer to [22] for a discussion about learning in cooperative games). Instead, we will focus on the following estimation problem faced by an external regulator: Given a sequence of observed prices and shocks $\{(p_1^j, p_2^j, \mu^j)\}_{j=1}^N$, the regulator would like to ascertain whether or not agents are playing according to ϵ-approximate Nash Equilibrium or not. In this setup, the private information vectors (θ_1, θ_2) of each agent are so-called *nuisance parameters* for the regulator (i.e., they require estimation even though they are not of primary interest). To that end, the regulator will construct estimates $(\hat{\theta}_1, \hat{\theta}_2)$ of the private information vectors and residual estimates $\hat{\epsilon}_j$ for each observation tuple $j \in \{1, ..., N\}$ by solving the inverse variational problem given by

$$\min_{\hat{\theta}, y, \hat{\epsilon}} \; L(\hat{\epsilon}^1, ..., \hat{\epsilon}^N) \tag{8}$$

$$\text{s.t. } A_i^\top y_i^j \leq_C \mathbf{f}_i(p_1^j, p_2^j), \text{for } i \in \{1, 2\}, j \in \{1, ..., N\} \tag{9}$$

$$\sum_{i=1}^2 \mathbf{f}_i(p_1^j, p_2^j)^\top p_i^j - b_i^\top y_i^j = \hat{\epsilon}^j, \text{for } j \in \{1, ..., N\} \tag{10}$$

where $L(\hat{\epsilon}^1, ..., \hat{\epsilon}^N)$ is some loss function over the residual estimates. We assumed that the regulator knows the distribution $\mathcal{D}(\phi)$, but does not know ϕ. Hence the loss function can be written, for example, as the negative log-likelihood as a function of ϕ [32]. We note in this optimization problem, the prices are given by our N samples, and we seek to select a $\hat{\theta}$ such that the resulting utilities form an approximate Nash equilibrium for every sample collected, where the computed $\hat{\epsilon}_j$ are our residual estimates of ϵ.

Lastly, in order to make a decision as to whether or not the observed prices are in approximate Nash equilibrium, the regulator will formulate a hypothesis test over the computed residuals.

2.2 Hypothesis Testing Framework

In order to formalize the hypothesis testing framework, we begin by describing the temporal sequence of events under consideration:

(1) Both agents and regulator observe μ, the shock variable.
(2) The agents solve the feasibility problem (7) for some $\epsilon \sim \mathcal{D}(\phi)$. The strategies (p_1, p_2) are selected to exactly be an ϵ-approximate Nash Equilibrium.
(3) The regulator observes the strategies (p_1, p_2) and records it.
(4) Steps 1-3 are repeated N times and the regulator collects the sample tuples $\{(p_1^j, p_2^j, \mu^j)\}_{j=1}^N$.
(5) The regulator solves the inverse variational problem (8) for some parametric utility functions and computes the estimated residuals $\hat{\epsilon}^1, ..., \hat{\epsilon}^N$.
(6) The regulator uses those residuals to perform a Kolmogorov-Smirnov test (to be defined next).

We note that the regulator does not know the true utility functions of the agents. Importantly, our approach is partially amenable to parametric form misspecification because non-colluding agents are not required to be in perfect equilibrium. In other words, some amount of the ϵ^j are meant to capture model misspecification.

Step 6 is conducted as follows: The regulator will use the computed estimated residuals to perform a hypothesis test to determine if the $\hat{\epsilon}^1, ..., \hat{\epsilon}^N$ come from the distribution $\mathcal{D}(\phi)$. However, even though the regulator knows the distribution's parametric form, they do not know the underlying parameter ϕ. Hence, hypothesis tests such as the standard Kolmogorov-Smirnov test are not applicable since they require knowing the true underlying parameters of the distribution under the null hypothesis. Therefore, we we will resort to the Lilliefors variation of the Kolmogorov-Smirnov test [23]. We first compute the empirical cumulative distribution function $\hat{F}_N(d) = \frac{1}{N} \sum_{j=1}^N \mathbb{I}(\hat{\epsilon}_j \leq d)$, where $\mathbb{I}(\cdot)$ is an indicator function. Then the regulator computes some estimate $\hat{\phi} = g(\hat{\epsilon}^1, ..., \hat{\epsilon}^N)$ and computes the cumulative distribution function $\bar{F}_N(d) = F_{\mathcal{D}(\hat{\phi})}(d)$, where $F_{\mathcal{D}(\hat{\phi})}(d)$ is the cumulative distribution function of a random variable of distribution $\mathcal{D}(\hat{\phi})$. Lastly, the regulator computes the test statistic $D^* = \max_d |\hat{F}_N(d) - \bar{F}_N(d)|$. The null H_0 and alternative H_1 hypotheses for our test are

$$\mathcal{H} : \begin{cases} H_0 : & \text{The agents are behaving in an } \epsilon\text{-approximate} \\ & \text{equilibrium where } \epsilon \sim \mathcal{D}(\phi) \\ H_1 : & \text{Otherwise} \end{cases} \tag{11}$$

And the decision of whether to accept or reject the null hypothesis is made using the decision-rule

$$\begin{cases} \text{reject } H_0 : & \text{if } D^* \geq \tau(N) \\ \text{accept } H_0 : & \text{if } D^* < \tau(N) \end{cases} \tag{12}$$

where $\tau(N)$ is some threshold from the Lilliefors variation of the Kolmogorov-Smirnov test [23] and which is based on the number of samples collected and the desired significance level α.

3 COMPUTATIONAL EXPERIMENTS

Here, we analyze the performance of our approach in a Bertrand competition environment. We first detail the experiment setting and then proceed to the numerical experiments and analysis.

3.1 Experiment Setting

We showcase our method in a setting where two agents compete over a single item and need to set their respective prices in the Bertrand-game environment. Each agent's true demand function has the following form:

$$\bar{D}_i(p_1, p_2, \mu; \bar{\theta}_i) = \bar{\theta}_{i,0} + \sum_{j=1}^2 p_j \bar{\theta}_{i,j} + \bar{\theta}_{i,3} \mu + \eta_i \tag{13}$$

where $\bar{\theta}_i$ is the agent's private information vector, and we use the term η_i to encompass unmodeled terms of the dynamics. Furthermore, we assume the set of feasible price vectors belong to the

polyhedral set

$$\mathcal{P} = \{(p_1, p_2) \in \mathbb{R}^2 : \ 0 \leq p_1 \leq \bar{p}, 0 \leq p_2 \leq \bar{p}\} \tag{14}$$

where $\bar{p}$ is an upper-bound on each price. We consider the case where ϵ is drawn from an exponential distribution $\epsilon \sim \exp(\bar{\lambda})$.

We assume that the regulator observes the shock μ but does not observe η_i. Hence, the regulator forms the following demand estimate given some estimate $\hat{\theta}$:

$$\mathbf{D}_i(p_1, p_2, \mu; \hat{\theta}) = \hat{\theta}_{i,0} + \textstyle\sum_{j=1}^{2} p_j \hat{\theta}_{i,j} + \hat{\theta}_{i,3}\mu \tag{15}$$

Following the steps described in the previous section, the regulator collects the sample tuples $\{(p_1^j, p_2^j, \mu^j)\}_{j=1}^N$ and forms the optimization problem (8) with the loss function $L(\hat{\epsilon}^1, ..., \hat{\epsilon}^N)$ being the negative log-likelihood of the underlying exponential distribution. Note that in the negative log-likelihood the λ term is decoupled from the other terms because of the particular mathematical form of the density of an exponential distribution. As a result, we do not need to include λ in the inverse variational problem.

In order to make the presentation of the final optimization problem clear, we define the marginal utility function for each agent (as considered by the regulator) to be

$$m_i(p_1, p_2, \mu; \hat{\theta}_i) = p_i \frac{\partial}{\partial p_i}\mathbf{D}_i(p_1, p_2, \mu; \theta_i) + \mathbf{D}_i(p_1, p_2, \mu; \theta_i) = p_i\hat{\theta}_{i,i} + \hat{\theta}_{i,0} + \textstyle\sum_{j=1}^{2} p_j\hat{\theta}_{i,j} + \hat{\theta}_{i,3}\mu. \tag{16}$$

In addition, we impose some structure to the fitted utility functions: (1) we normalize the fitted utility functions; (2) we enforce that the marginal utilities of each agent decrease as they increase their own prices (on the observed data); and (3) we enforce an additional constraint that sets the dual variable y_i^j to zero if the observed price p_i^j is strictly less than the upper bound $\bar{p}$. Recalling the definition of $\mathbf{f}(p)$ in (4), the optimization problem becomes

$$\min_{\hat{\epsilon}, y, \theta_1, \theta_2} \ \textstyle\sum_{j=1}^{N} \hat{\epsilon}_j \tag{17}$$

$$\text{s.t. } y_i^j \geq m_i(p_1^j, p_2^j, \mu^j, \theta_1), \text{for } i \in \{1, 2\}, j \in \{1, ..., N\} \tag{18}$$

$$\bar{p}\textstyle\sum_{i=1}^{2}(y_i^j) - \sum_{i=1}^{2} p_i^j m_i(p_1^j, p_2^j, \mu^j, \theta_i) = \hat{\epsilon}_j, \quad \text{for } j \in \{1, ..., N\} \tag{19}$$

$$m_i(1, 1, 0, \theta_i) = m_i(1, 1, 0, \bar{\theta}_i), \text{for } i \in \{1, 2\} \tag{20}$$

$$y_i^j = 0, \text{for } i \in \{1, 2\}, j \in \{1, ..., N\} \text{ s.t. } p_i^j < \bar{p} \tag{21}$$

$$\theta_{i,i} \leq 0, \text{for } i \in \{1, 2\} \tag{22}$$

$$\hat{\epsilon}^j \geq 0, \text{for } j \in \{1, ..., N\} \tag{23}$$

$$y^j = (y_1^j, y_2^j) \geq 0, \text{for } \forall j \in \{1, ..., N\} \tag{24}$$

where (20) are the normalization constraints, in which the marginal utility of both agents when there is no external shock and the prices are set to unity is equal to the true marginal at that point. (Note we could have set these normalization constraints to any other suitable positive value without affecting the results. Different normalization may yield different models that can be used to explain the same observed data. This phenomenon is common in inverse optimization problems, as discussed in detail in [4, 9].) Equation (22) ensures that the fitted marginal functions decrease as the agents increase their own prices. (This constraint is obtained after some arithmetic by requiring that $m_1(p_1, \cdot, \cdot; \theta_i)$ and $m_2(\cdot, p_2, \cdot; \theta_i)$ decrease as p_1 and p_2 increase, respectively, on the observed data.) The "dual" vector y is associated with the constraints (14) and (23) ensures that if the the observed prices are not on the boundary of the feasible region $\mathcal{P}$ then the associated dual variable is set to zero. We note that (23) has a very subtle implication in the optimization problem above: The very natural notion that dual variables are zero once their associated constraint is non-binding is not enforced at all by the original formulation in (8). If the prices are sampled in perfect Nash equilibrium (that is $\epsilon = 0$), then as argued in [9] the formulation in (8) is able to recover exactly the true parameters θ_i and the computed residuals are exactly zero. However, in our scenario prices are obtained in approximate equilibrium (i.e., $\epsilon > 0$). Hence if complimentary slackness (23) is not enforced explicitly then the computed residuals will present bias – namely they will be "shrunk" since the formulation (8) could achieve smaller values for the residuals by setting the dual variables to be positive, even though the sampled prices are in the interior of the feasible region.

Recall that the objective function follows from the negative log-likelihood of exponential distribution, where we dropped the term N/λ since it does not impact the optimization. After solving the problem above, we compute the MLE estimate $\hat{\lambda}_{MLE} = (\frac{1}{N}\sum_{j=1}^{N}\hat{\epsilon}_j)^{-1}$ of λ. Then we let $\bar{F}_N(d) = F_{\exp(\hat{\lambda}_{MLE})}(d)$ and conduct the Lilliefors hypothesis test (12). To illustrate the performance of the hypothesis testing we will simulate the process under two scenarios:

Scenario 1: Agents are competing over prices, i.e.: they solve the feasibility problem (7) after observing the shock variable μ and the value of ϵ.

Scenario 2: Agents are colluding, i.e.: instead of solving the feasibility problem (7), they maximize the sum of both utility functions up to a ϵ optimality gap.

Hence for Scenario 2, prices are generated by solving

$$(p_1^j, p_2^j) = \arg\max_{(p_1, p_2) \in \mathcal{P}} \textstyle\sum_{i=1}^{2} p_i D_i(p_1, p_2, \theta_i, \mu_j) \tag{25}$$

for $j \in \{1, ..., N\}$. In the next subsection, we present numerical simulations of these two scenarios and show how the regulator rejects/does not reject the null hypothesis as the agents change their behavior from competition to collusion.

Table 1: Numerical Results for Scenario 1 (Competing)

N	D^*	$\tau(N)$	$\hat{\lambda}$	Decision
10	0.317	0.325	33.7	Competing
20	0.206	0.234	27.93	Competing
30	0.120	0.192	20.83	Competing
40	0.069	0.168	21.43	Competing
50	0.089	0.150	19.99	Competing
100	0.070	0.106	18.80	Competing
200	0.031	0.075	18.62	Competing
500	0.022	0.047	20.01	Competing

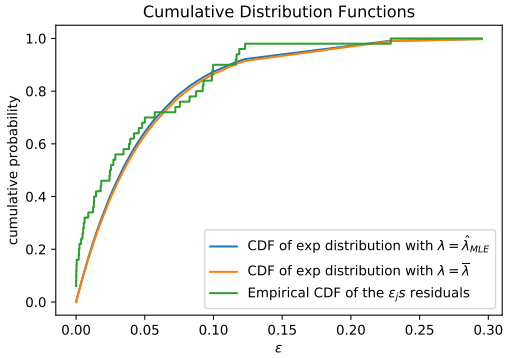

Figure 1: Comparing CDF's of Residuals For Scenario 1

3.2 Computational results

We let $\bar{\lambda} = 20$ and chose $\bar{\theta}_1 = [10, -1, 0.5, 1]$ and $\bar{\theta}_2 = [8, 0.4, -3.0, 1]$ to be the agents' true private information vectors. The shock values were generated by $\mathcal{N}(5, 1)$, and we fix the upper-bound $\bar{p} = 8.0$ on the prices. Furthermore, we fix our significance level $\alpha = 0.05$. The threshold $\tau(N)$ for the hypothesis testing is obtained by the table presented in [23]. Lastly, we let η_j for $j = \{1, 2\}$ be sampled from $\mathcal{N}(0, 1)$. For the first scenario, the approximate equilibrium prices need to be generated by solving (7). In our test case, we generate approximate equilibrium prices via the algorithm described in the appendix. The results for Scenario 1 are summarized in Table 1.

When agents are competing (i.e., acting under the specifications of the null hypothesis), a false positive (i.e., decision of collusion occurring) was not seen in the experiments. This is not surprising because we set $\alpha = 0.05$ and each row in the table corresponds to a single numerical experiment. If we repeated these experiments, we would expect to see a close to α fraction of them report a false positive. Also, note we are able to recover the correct estimate of λ for the underlying distribution generating the residuals. This is highlighted in Figure 1, where we plot the $\epsilon_j's$ samples from exp(20) and the computed residual estimates $\hat{\epsilon}_j's$ by the regulator after solving the optimization problem, for sample size equal to 50.

For the second scenario we generate prices by solving an aggregate problem where we sum both agents' utilities to compute the prices. In Table 2, we see that the null hypothesis is rejected (i.e., decision of collusion occurring) for moderate and large sample sizes. The MLE estimate of $\bar{\lambda}$ is inaccurate as well since the agents are not behaving in approximate equilibrium. In Figure 2, for $N = 50$ we plot the empirical CDF of residual estimates $\hat{\epsilon}_j's$ by the regulator in this scenario. Observe that when agents are cooperating instead of competing, the computed residuals are vastly different then their true values (we omit plotting the true cdf of exp(20) since the computed residuals are very large for this scenario). The null hypothesis that the agents are competing in equilibrium is rejected for almost all sample sizes, indicating that our method is able to identify when agents are not behaving in competition. We stress that rejecting the null hypothesis is not proof that agents are colluding, but rather gives some statistical evidence that suggests collusion is occurring.

Table 2: Numerical Results for Scenario 2 (Colluding)

N	D^*	$\tau(N)$	$\hat{\lambda}$	Decision
10	0.261	0.325	0.12	Competing
20	0.263	0.234	0.11	Colluding
30	0.222	0.192	0.22	Colluding
40	0.300	0.168	0.26	Colluding
50	0.301	0.150	0.22	Colluding
100	0.322	0.106	0.28	Colluding
200	0.301	0.075	0.30	Colluding
500	0.335	0.047	0.30	Colluding

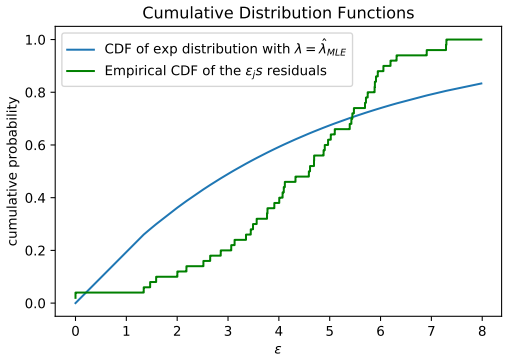

Figure 2: Comparing CDF's of Residuals For Scenario 2

4 CONCLUSION AND FUTURE WORK

We proposed a hypothesis testing framework to decide whether agents are behaving competitively or not. In our setting, a regulator formulates an inverse variational problem in order to estimate the unknown private information vectors as well as estimate the residuals of the approximate equilibrium that arises from the agents' competition. Our setting is flexible as the regulator only requires access to prices and shock values. A future direction of work is to study the theoretical properties of our framework. We demonstrated our method in a simple two-player game with a polyhedral feasible action space. We stress our setting is more general and allows for any number of players with arbitrarily conic-representable sets, as long as they satisfy some regularity conditions. Another direction of future research is to apply our testing framework to the setting in [11, 22], where groups of agents "learn" to collude instead of competing. This problem is more challenging but can be explored in the light of inverse variational problems and our framework.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant CMMI-1847666, and by the UC Berkeley Center for Long-Term Cybersecurity.

REFERENCES

[1] Ravindra K Ahuja and James B Orlin. 2001. Inverse optimization. *Operations Research* 49, 5 (2001), 771–783.

[2] Gad Allon, Awi Federgruen, and Margaret Pierson. 2011. How much is a reduction of your customers' wait worth? An empirical study of the fast-food drive-thru industry based on structural estimation methods. *Manufacturing & Service Operations Management* 13, 4 (2011), 489–507.
[3] Anil Aswani. 2019. Statistics with set-valued functions: applications to inverse approximate optimization. *Mathematical Programming* 174, 1-2 (2019), 225–251.
[4] Anil Aswani, Zuo-Jun Shen, and Auyon Siddiq. 2018. Inverse optimization with noisy data. *Operations Research* 66, 3 (2018), 870–892.
[5] Patrick Bajari, C Lanier Benkard, and Jonathan Levin. 2007. Estimating dynamic models of imperfect competition. *Econometrica* 75, 5 (2007), 1331–1370.
[6] Patrick Bajari and Lixin Ye. 2003. Deciding between competition and collusion. *Review of Economics and statistics* 85, 4 (2003), 971–989.
[7] Laura H Baldwin, Robert C Marshall, and Jean-Francois Richard. 1997. Bidder collusion at forest service timber sales. *Journal of Political Economy* 105, 4 (1997), 657–699.
[8] B Douglas Bernheim, Bezalel Peleg, and Michael D Whinston. 1987. Coalition-proof nash equilibria i. concepts. *Journal of Economic Theory* 42, 1 (1987), 1–12.
[9] Dimitris Bertsimas, Vishal Gupta, and Ioannis Ch Paschalidis. 2015. Data-driven estimation in equilibrium using inverse optimization. *Mathematical Programming* 153, 2 (2015), 595–633.
[10] Severin Borenstein. 2004. Rapid price communication and coordination: The airline tariff publishing case (1994). *The Antitrust Revolution: Economics, Competition, and Policy* 4 (2004).
[11] Emilio Calvano, Giacomo Calzolari, Vincenzo Denicolò, and Sergio Pastorello. 2018. Artificial intelligence, algorithmic pricing and collusion. (2018).
[12] Emilio Calvano, Giacomo Calzolari, Vincenzo Denicolò, and Sergio Pastorello. 2019. Algorithmic Pricing What Implications for Competition Policy? *Review of Industrial Organization* 55, 1 (2019), 155–171.
[13] Le Chen, Alan Mislove, and Christo Wilson. 2016. An empirical analysis of algorithmic pricing on Amazon marketplace. In *International Conference on World Wide Web*. 1339–1349.
[14] Francis Y Edgeworth. 1925. The pure theory of monopoly. *Papers relating to political economy* 1 (1925), 111–142.
[15] Ariel Ezrachi and Maurice E Stucke. 2017. Artificial intelligence & collusion: When computers inhibit competition. *U. Ill. L. Rev.* (2017), 1775.
[16] Farzaneh Farhadi, S Jamaloddin Golestani, and Demosthenis Teneketzis. 2018. A surrogate optimization-based mechanism for resource allocation and routing in networks with strategic agents. *IEEE TAC* 64, 2 (2018), 464–479.
[17] Joseph Farrell. 1988. Communication, coordination and Nash equilibrium. *Economics Letters* 27, 3 (1988), 209–214.
[18] Franco Giannessi, Antonino Maugeri, and Panos M Pardalos. 2006. *Equilibrium problems: nonsmooth optimization and variational inequality models*. Vol. 58. Springer Science & Business Media.
[19] Joseph Greenberg. 1994. Coalition structures. *Handbook of game theory with economic applications* 2 (1994), 1305–1337.
[20] Patrick T Harker and Jong-Shi Pang. 1990. Finite-dimensional variational inequality and nonlinear complementarity problems: a survey of theory, algorithms and applications. *Mathematical programming* 48, 1-3 (1990), 161–220.
[21] Ramesh Johari and John N Tsitsiklis. 2009. Efficiency of scalar-parameterized mechanisms. *Operations Research* 57, 4 (2009), 823–839.
[22] Xin Li and Leen-Kiat Soh. 2004. Investigating reinforcement learning in multiagent coalition formation. In *AAAI*. 22–28.
[23] Hubert W Lilliefors. 1969. On the Kolmogorov-Smirnov test for the exponential distribution with mean unknown. *J. Amer. Statist. Assoc.* 64, 325 (1969), 387–389.
[24] Frank J Massey Jr. 1951. The Kolmogorov-Smirnov test for goodness of fit. *Journal of the American statistical Association* 46, 253 (1951), 68–78.
[25] Reinhard Neck and Engelbert Dockner. 1987. Conflict and cooperation in a model of stabilization policies: A differential game approach. *Journal of Economic Dynamics and Control* 11, 2 (1987), 153–158.
[26] Martin Pesendorfer. 2000. A study of collusion in first-price auctions. *The Review of Economic Studies* 67, 3 (2000), 381–411.
[27] Robert H Porter and J Douglas Zona. 1993. Detection of bid rigging in procurement auctions. *Journal of political economy* 101, 3 (1993), 518–538.
[28] John Rust. 1994. Structural estimation of Markov decision processes. *Handbook of econometrics* 4 (1994), 3081–3143.
[29] Walid Saad, Zhu Han, Mérouane Debbah, Are Hjørungnes, and Tamer Basar. 2009. Coalitional game theory for communication networks: A tutorial. *arXiv preprint arXiv:0905.4057* (2009).
[30] Walid Saad, Zhu Han, and H Vincent Poor. 2011. Coalitional game theory for cooperative micro-grid distribution networks. In *IEEE ICC*. 1–5.
[31] Bruno Salcedo. 2015. Pricing algorithms and tacit collusion. *Manuscript, Pennsylvania State University* (2015).
[32] Thomas A Severini. 2000. *Likelihood methods in statistics*. Oxford University Press.
[33] John M Staatz. 1983. The cooperative as a coalition: a game-theoretic approach. *American Journal of Agricultural Economics* 65, 5 (1983), 1084–1089.
[34] Sichao Yang and Bruce Hajek. 2007. VCG-Kelly mechanisms for allocation of divisible goods: Adapting VCG mechanisms to one-dimensional signals. *IEEE Journal on Selected Areas in Communications* 25, 6 (2007), 1237–1243.

A ALGORITHM TO GENERATE APPROXIMATE EQUILIBRIUM PRICES

A key part of our numerical simulations is to generate prices that are ϵ-approximate equilibrium. In the general case, we need to solve the variational inequality formulation in (7). That problem is hard to solve in general, but tailored algorithms do exist [18]. However, for our setting the feasible region $\mathcal{P}$ contains only bounds on the prices. Hence the problem becomes to find prices (p_1, p_2) such that

$$\exists y_1, y_2 \geq 0 : \begin{cases} y_i \geq D_i(p_1, p_2.\mu, \bar{\theta}_i) + p_i\theta_{i,i}, \text{for } i \in \{1,2\} \\ \sum_{i=1}^{2} \bar{p} y_i - p_i(D_i(p_1, p_2.\mu, \bar{\theta}_i) + p_i\theta_{i,i}) = \epsilon \end{cases} \tag{26}$$

Hence we can generate samples of (p_1, p_2) by acceptation/rejection of samples based on the shock values μ and nuisance parameters (η_1, η_2). First, we sample μ and η_1, η_2 according to their specified distributions. Then we solve the following system of nonlinear equations (via, for example, Newton's Method): $p_i D_i(p_1, p_2, \mu, \bar{\theta}_i) + (p_i)^2\theta_{i,i} = \frac{-\epsilon}{2}$, for $i \in \{1,2\}$. After solving this system, if $(p_1, p_2) \in \mathcal{P}$ then it means they are ϵ-approximate solution to the variational inequality problem (since we can set both y_1 and y_2 to zero), and we accept the sample (p_1, p_2, μ). If $p_1 < 0$ or $p_2 < 0$, then we reject the sample. Now without loss of generality, suppose that $p_1 > \bar{p}$. Then we can set $p_1 = \bar{p}$ and let $y_1 = D_1(\bar{p}, p_2.\mu, \bar{\theta}_1) + \bar{p}\theta_{1,1}$. Then by letting $y_2 = 0$ we solve for p_2 by $p_2 D_2(\bar{p}, p_2, \mu, \bar{\theta}_i) + (p_2)^2\theta_{2,2} = -\epsilon$. Lastly if $p_2 \geq 0$ and $y_1 \leq 0$, then we accept the sample (p_1, p_2, μ). In all other cases, we reject the sample. With this simple method, we can generate sample prices that are ϵ-approximate equilibrium. By repeating the above N times for each sampled ϵ_j, we can generate all the samples necessary for the numerical simulation.