Empirical Mechanism Design: Designing Mechanisms from Data

Enrique Areyan Viqueira
Brown University
Providence, RI, USA
NEC-AIST AI Cooperative
Research Laboratory
Tokyo, Japan

Cyrus Cousins
Brown University
Providence, RI, USA

Yasser Mohammad
NEC-AIST AI Cooperative
Research Laboratory
Tokyo, Japan

Amy Greenwald
Brown University
Providence, RI, USA
NEC-AIST AI Cooperative
Research Laboratory
Tokyo, Japan

Abstract

We introduce a methodology for the design of parametric mechanisms, which are multi-agent systems inhabited by strategic agents, with knobs that can be adjusted to achieve specific goals. We assume agents play approximate equilibria, which we estimate using the probably approximately correct learning framework. Under this assumption, we further learn approximately optimal mechanism parameters. We do this both theoretically, assuming a finite design space, and heuristically, using Bayesian optimization (BO). Our BO algorithm incorporates the noise associated with modern concentration inequalities, such as Hoeffding's, into the underlying Gaussian process. We show experimentally that our search techniques outperform standard baselines in a stylized but rich model of advertisement exchanges.

1 INTRODUCTION

Mechanism design is concerned with the design of multi-agent systems that achieve certain objectives, assuming strategic behavior on the part of the participating agents. As mechanisms are in effect games, a standard assumption is that agents exhibit equilibrium behavior.

This paper is concerned with **parameterized mechanism design**, where the designer has at their disposal a parameterized mechanism, each parameter setting of which induces a different game, and is interested in setting the parameters so that the ensuing equilibria exhibit certain properties. For example, an online auctioneer (the mechanism designer) wishing to maximize revenue (the objective) can set reserve prices (the parameters), prices below which a bid cannot win, to try to maximize their revenue, assuming, for example, Nash-equilibrium bidding.

But an application domain like online auctions (e.g., advertising auctions) comes with a twist: an analytical description of the game is not available, because the game dynamics are too complicated (e.g., users of all different demographics are coming and going, stochastically). Only a simulator of these so-called **simulation-based** [1], or **black-box** [2], games is available, but simulations are expensive. Still, the mechanism designer can query the simulator a few times to determine the quality of the current parameter settings, before moving on to try some others. The question we address in this paper is: how should a mechanism designer (e.g., an auctioneer) set parameters (e.g., reserve prices), given access only to *data* (or to a *simulator* capable of generating data) about the game under different settings.

A key source of the expense in most mechanism design simulators is the internal computation characterizing the agents' strategic behavior. It is generally assumed that agents' play an equilibrium, but equilibrium computation is notoriously difficult [3]; and adding insult to injury, an analytical description of these games is not available, so the games themselves must first be learned from data. The construction of **empirical games** (i.e., games learned from data) that well estimate the corresponding simulation-based games, and the search for equilibrium within is called **empirical game-theoretic analysis** (EGTA) [4, 5]. When mechanism design depends on EGTA, it is called **empirical mechanism design** (EMD).

The contributions of this paper comprise a novel methodology for learning equilibria from data using the probably-approximately correct (PAC) learning framework [6], and thus contribute to the literature on EGTA. We also note that EMD, where the parameter-induced games are accessible only via a black box, is an instance of black-box optimization. Hence, we review standard search routines for black-box optimization problems, such as Bayesian optimization [7], and then instantiate them using the **piecewise constant noise** that characterizes PAC learners. In total, our methodology makes it possible to carry out

EMD on very rich games, at times obtaining theoretical guarantees about the quality of the mechanisms learned.

To demonstrate, we apply our methods to two mechanism design problems. First, we study a relatively simple environment—symmetric first-price auctions—where analytical solutions are available. We show that we can, more quickly than our baselines, recover a near-revenue-maximizing reserve price, assuming all bidders play the unique Bayes-Nash equilibrium. Second, we study a rich model of electronic advertisement exchanges, where bidders bid for the opportunity to display advertisements to users that arrive stochastically. This domain is much more difficult to analyze, because an equilibrium must be computed empirically. We use our novel EGTA methodology to construct empirical games and their equilibria, and then we wrap our search heuristics around this equilibrium computation. We show that our inner equilibrium-finding methodology is effective, and further that our methodology discovers better solutions to the outer mechanism design problem more quickly than our baselines.

Related Work. Mechanism design is concerned with the design of games in which the strategic behavior of participants leads to desired outcomes. Traditional mechanism design lies within the realm of game theory, itself an area of economics. However, with the advent of strategic autonomous agents, particularly in e-commerce settings, there has been a surge of interest in mechanism design within the computer science community [8, 9]. In AI, at least two variants of mechanism design have been defined. The first is **automated mechanism design** [9, 10, 11], where a mechanism is automatically created by an algorithm that searches through a space of mechanisms constrained by standard mechanism design criteria, such as individual rationality and incentive compatibility. The second is EMD, where the designer is interested in optimizing a mechanism's parameters relative to some objective of interest, under the assumption of equilibrium play. One distinguishing feature of our EMD work *vis à vis* the existing literature is that we formulate the search for a mechanism's optimal parameters as a black-box optimization problem, and leverage Bayesian optimization techniques to perform the ensuing search.

2 EMPIRICAL GAME-THEORETIC ANALYSIS

Before we tackle the mechanism design problem, we focus on a related search problem, the search for equilibria, which is a key source of the expense in evaluating potential mechanisms. Generally speaking, an equilibrium in a game is achieved when agents' strategies are in a steady state; none has any incentive to deviate from their prescribed behavior. There are many such equilibrium concepts, one very notable one being that of Nash [12]. Nash equilibria, however, are not guaranteed to exist (in finite games), except in mixed strategies (i.e., by allowing for the possibility randomization), and mixed strategy equilibria are notoriously difficult to compute [3].

An alternative to Nash that is easier to compute are **sink** equilibria [13]. The sink equilibria are the sinks (i.e., strongly connected components without any outgoing edges) of what is called the game's **better-response graph** (BRG). This is a directed graph whose nodes are strategy profiles (one strategy per agent), and where each edge indicates that an agent would deviate from that node to the one to which it points. It turns out that sink equilibria are not readily amenable to approximation,[1] so we take as our solution concept the larger set of all **strongly connected components** (SCCs) of a game's BRG.

We begin this section by formally defining games, better-response graphs, and the set of SCCs as an equilibrium. We restrict our attention to finite games (and graphs).

Definition 2.1 (Normal-Form Game)**.** *A **normal-form game** $\Gamma \doteq \langle P, \{S_p\}_{p \in P}, \boldsymbol{u}(\cdot)\rangle$ consists of a set of agents P, with **strategy set** S_p for agent $p \in P$. Let $\boldsymbol{S} \doteq S_1 \times \cdots \times S_{|P|}$ be the strategy profile space of Γ, and then $\boldsymbol{u} : \boldsymbol{S} \to \mathbb{R}^{|P|}$ is a vector-valued utility function (equivalently, a vector of $|P|$ scalar utility functions $\boldsymbol{u}_p$). Given such a game Γ, we define its size $|\Gamma| \doteq |P| \prod_{p=1}^{|P|} |S_p|$.*

Example 2.1. *The Prisonsers' Dilemma game and its corresponding better-response graph are shown in Figure 1. Nodes are labelled by strategy profiles (e.g., CC), and edges are labelled by color, with red (blue) corresponding to the row (column) player.*

	C	D
C	3, 3	0, 5
D	5, 0	1, 1

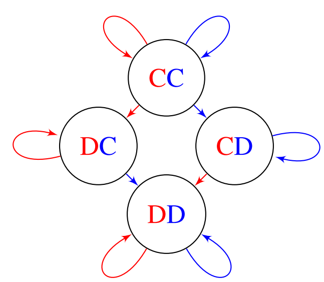

Figure 1: Prisoners' Dilemma's better-response graph.

For the next few definitions, we fix a game Γ.

Definition 2.2 (ϵ-Better Response)**.** *An ϵ-better response for agent p at strategy profile $\boldsymbol{s}$ is a strategy profile $\boldsymbol{s}^* = (s_1, \ldots, s_p^*, \ldots, s_{|P|})$ where agent p plays $s_p^* \in S_p$, and all agents other than p play s_j, the strategy played by agent $j \neq p$ at $\boldsymbol{s}$, such that $\boldsymbol{u}_p(\boldsymbol{s}^*) + \epsilon \geq \boldsymbol{u}_p(\boldsymbol{s})$.*

[1] A counterexample appears in the supplemental material.

Definition 2.3 (ϵ-Better Response Graph). *An **ϵ-better response graph**, $B_\epsilon(\Gamma) = (\mathcal{V}, \mathcal{E}_\epsilon)$, is a directed graph with a node for each strategy profile $\boldsymbol{s} \in \boldsymbol{S}$, and an edge $(\boldsymbol{s}, \boldsymbol{s}')$ iff $\boldsymbol{s}'$ is an ϵ-better response for some agent at $\boldsymbol{s}$.*

Definition 2.4 (Strongly Connected Component). *Let $G = (V, E)$ be a directed graph. A **strongly connected component** (SCC) of G is a set of nodes $U \subseteq V$ s.t. $\forall u, v \in U, \exists$ a path from u to v, and one from v to u.*

An **ϵ-SCC equilibrium** is a SCC of an ϵ-better response graph. We denote by $\mathrm{SCC}_\epsilon(\Gamma)$ the set of all ϵ-SCC equilibria of Γ. Similarly, an **ϵ-sink equilibrium** is a SCC of an ϵ-better response graph *without any outgoing edges*.

Approximating Equilibria Recall that our goal is to learn equilibria from data. To do so, we start by approximating a simulation-based game from data. But given an approximation of one game by another, there is not necessarily a connection between their equilibria. In particular, there may be equilibria in one game with no corresponding equilibria in the other, as small changes to the utility functions can add or remove equilibria. Nonetheless, we now show the SCC equilibria of games that are close enough to one another are themselves close. We formalize this idea using the concept of a uniform approximation.

Given compatible games—games with the same agents sets P and strategy profile spaces $\boldsymbol{S}$—with utility functions $\boldsymbol{u}$ and $\boldsymbol{u}'$, respectively, define $\|\Gamma - \Gamma'\|_\infty \doteq \|\boldsymbol{u}(\cdot) - \boldsymbol{u}'(\cdot)\|_\infty \doteq \max_{p \in P, \boldsymbol{s} \in \boldsymbol{S}} |\boldsymbol{u}_p(\boldsymbol{s}) - \boldsymbol{u}'_p(\boldsymbol{s})|$. Now Γ' is said to be a **uniform ϵ-approximation** of Γ whenever $\|\Gamma - \Gamma'\|_\infty \leq \epsilon$. Uniform approximations bound between utility deviations in Γ and Γ' *uniformly* over *all* players and strategy profiles.

Theorem 2.1 (Approximate Equilibria). *Let $\epsilon > 0$. If Γ' is a uniform approximation of Γ, then $\mathrm{SCC}_0(\Gamma) \rightsquigarrow \mathrm{SCC}_{2\epsilon}(\Gamma') \rightsquigarrow \mathrm{SCC}_{4\epsilon}(\Gamma)$, where $\mathcal{A} \rightsquigarrow \mathcal{B}$ means that for all $A \in \mathcal{A}$, there exist $B \in \mathcal{B}$, such that $A \subseteq B$.*

We show Thm. 2.1 via a lemma concerning the edge set of the better-responses graphs of uniform approximations.

Lemma 2.1 (ϵ-BRG edges containment.). *Let $\epsilon > 0$. If $\|\Gamma - \Gamma'\|_\infty \leq \epsilon$, then $\mathcal{E}_0(\Gamma) \subseteq \mathcal{E}_{2\epsilon}(\Gamma') \subseteq \mathcal{E}_{4\epsilon}(\Gamma)$.*

Proof. If $(\boldsymbol{s}, \boldsymbol{t}) \in \mathcal{E}_0(\Gamma)$, then there exists p such that $\boldsymbol{u}_p(\boldsymbol{t}) \geq \boldsymbol{u}_p(\boldsymbol{s})$. The following chain of reasoning then holds: $\boldsymbol{u}'_p(\boldsymbol{t}) + \epsilon \geq \boldsymbol{u}_p(\boldsymbol{t}) \geq \boldsymbol{u}_p(\boldsymbol{s}) \geq \boldsymbol{u}'_p(\boldsymbol{s}) - \epsilon$, where the first and last inequalities follow from the uniform approximation assumption. Hence, $\boldsymbol{u}'_p(\boldsymbol{t}) \geq \boldsymbol{u}'_p(\boldsymbol{s}) - 2\epsilon$, and thus $(\boldsymbol{s}, \boldsymbol{t}) \in \mathcal{E}_{2\epsilon}(\tilde{\Gamma})$. Now, starting from the assumption that $(\boldsymbol{s}, \boldsymbol{t}) \in \mathcal{E}_{2\epsilon}(\tilde{\Gamma})$, the following chain of reasoning also holds: $\boldsymbol{u}_p(\boldsymbol{t}) + \epsilon \geq \boldsymbol{u}'_p(\boldsymbol{t}) \geq \boldsymbol{u}'_p(\boldsymbol{s}) - 2\epsilon \geq (\boldsymbol{u}_p(\boldsymbol{s}) - \epsilon) - 2\epsilon = \boldsymbol{u}_p(\boldsymbol{s}) - 3\epsilon$. Hence, $\boldsymbol{u}_p(\boldsymbol{t}) \geq \boldsymbol{u}_p(\boldsymbol{s}) - 4\epsilon$, and thus $(\boldsymbol{s}, \boldsymbol{t}) \in \mathcal{E}_{4\epsilon}(\Gamma)$. ☐

Proof of Theorem 2.1. We must show that any SCC of $B_0(\Gamma)$ remains strongly connected in $B_{2\epsilon}(\Gamma')$. Consider $Z \in B_0(\Gamma)$. Since all the edges in $\mathcal{E}_0(\Gamma)$ are also present in $\mathcal{E}_{2\epsilon}(\Gamma')$, it follows that any path connecting two nodes in Z is preserved in $\mathcal{E}_{2\epsilon}(\Gamma')$. Consequently, Z remains strongly connected in $B_{2\epsilon}(\Gamma')$. Similarly, any SCC of $B_{2\epsilon}(\Gamma')$ remains strongly connected in $B_{4\epsilon}(\Gamma)$. ☐

This theorem establishes perfect recall by the approximate game, in the sense that the approximate game contains all true positives: i.e., all (exact) equilibria of the original game. It also establishes approximately perfect precision, in the sense that all false positives in the approximate game are approximate equilibria in the original game.

Learning Games We now move on from approximating games to estimating them with guarantees. We are interested not only in estimating the game's parameters, by which we mean its utility functions; we are also interested in estimating the value of a mechanism designer's objective function, defined on these utilities. We use f to refer to the designer's objective, with $f(\boldsymbol{s})$ denoting the value of this objective at strategy profile $\boldsymbol{s}$.

Definition 2.5 (Empirical Normal-Form Game). *Consider a black-box game Γ whose utilities may depend on random draws over $\mathcal{X}$ from a distribution $\mathscr{D}$. Assuming samples $\boldsymbol{X} \sim \mathscr{D}^m$, a black-box game simulator would output $\boldsymbol{u}(x_j, \boldsymbol{s})$, for each x_j, based on which a mechanism designer could compute $f(\boldsymbol{u}(x_j, \boldsymbol{s}))$. We define the **empirical utility function** $\hat{\boldsymbol{u}}_{\boldsymbol{X}}(\boldsymbol{s}) \doteq \frac{1}{m}\sum_{j=1}^{m} \boldsymbol{u}(x_j, \boldsymbol{s})$, the corresponding **empirical normal-form game** $\hat{\Gamma}_{\boldsymbol{X}} \doteq \langle P, \{S_p\}_{p \in P}, \hat{\boldsymbol{u}}_{\boldsymbol{X}}(\cdot)\rangle$, and the **empirical objective function** $\hat{f}_{\boldsymbol{X}}(\boldsymbol{s}) \doteq \frac{1}{m}\sum_{j=1}^{m} f(\boldsymbol{u}(x_j, \boldsymbol{s}))$.*

Assume that the range of f is $[c_-, c_+]$ and take $\Delta \doteq c_+ - c_-$. Under these assumptions, Hoeffding's inequality [14] upper bounds the probability that the absolute difference between the empirical and expected mean exceeds ϵ as $\mathbb{P}_{\boldsymbol{X} \sim \mathscr{D}}\left(\left|\mathbb{E}_{x \sim \mathscr{D}}\left[f(\boldsymbol{u}(x, \boldsymbol{s}))\right] - \frac{1}{m}\sum_{j=1}^{m} f(\boldsymbol{u}(x_j, \boldsymbol{s}))\right| \geq \epsilon\right) \leq 2e^{-2\epsilon^2 m/\Delta^2}$. To obtain a uniform guarantee, we can apply a union bound to Hoeffding's inequality:

Theorem 2.2 (Hoeffding Finite-Sample Uniform Convergence Bounds for Designers' Objectives). *Consider a simulator $\mathcal{S}$ of a game Γ whose utilities may depend on random draws over $\mathcal{X}$ from a distribution $\mathscr{D}$ such that for all $x \in \mathcal{X}$ and $\boldsymbol{s} \in \boldsymbol{S}$, it holds that $f(\boldsymbol{u}(x, \boldsymbol{s})) \in [c_-, c_+]$ with $\Delta \doteq c_+ - c_-$, and take $\epsilon \doteq \Delta\sqrt{\ln(2|\Gamma|/\delta)/2m}$. Then $\mathrm{Pr}_{\boldsymbol{X} \sim \mathscr{D}^m}\left(\left\|f_{\mathscr{D}} - \hat{f}_{\boldsymbol{X}}\right\|_\infty \leq \epsilon\right) \geq 1 - \delta$.*

The exact same argument can be used to obtain uniform convergence bounds on utilities, as well as objective function values. This theorem, and its utility counterpart, thus establish that black-box games can be uniformly well-approximated with high probability, and that objective functions are likewise well-approximable.

All of these estimations can be accomplished very simply, by simulating the game m times, and then averaging the ensuing utilities across simulations. The requisite number of samples m is a function of a user-specified desired accuracy ϵ and failure probability δ.

Note that Hoeffding's is one of many possible choices of concentration inequalities that can be used learn empirical games and their properties. Hoeffding requires bounded noise, but this is not an inherent limitation of our methodology. We could obtain similar results under varied noise assumptions; e.g., we could assume (unbounded) *subgaussian* or *subexponential* noise, and substitute the appropriate Chernoff bounds.

3 BLACK-BOX OPTIMIZATION WITH NOISY MEASUREMENTS

In this section, we define two generic black-box optimization problems, and two corresponding algorithmic solutions. The first algorithm is primarily of theoretical interest; the second is heuristic, but more practical. Later, using the fact that empirical mechanism design is an instance of black-box optimization, we apply our heuristic approach to two EMD applications.

Definition 3.1 (Optimization with noisy measurements (OwNM))**.** *Given a design space* Θ*, an objective function* $F : \Theta \to \mathbf{R}$*, a noise model* $\mathscr{D}$*, and a measurement operator* $M : \Theta \to P_F$*, where* P_F *is the space of all possible probability distributions over the range of* F*, find* $\theta^* \in \arg\max_{\theta \in \Theta} F(\theta)$.

In this paper, we are concerned with a specific form of measurements, which produce **piecewise constant uniform** noise. This noise model is that which results from probably approximately correct (PAC)-style guarantees of the form, "accuracy is achieved up to additive error ϵ with probability $1-\delta$" [6].

More formally, we assume the measurement operator M returns $\hat{F}(\theta)$ along with an additive error bound ϵ that holds with probability $1-\delta$. In other words, the algorithm outputs a $1-\delta$ confidence interval $[c_1, c_2]$ of width 2ϵ centered at $\hat{F}(\theta)$. Now assuming the range of $F(\theta)$ is $[c_-, c_+]$, and letting $\Delta \doteq c_+ - c_-$, we take as a sample measurement of the pdf $p_F(x)$ the following:

$$\hat{p}_F(x) = \begin{cases} \frac{\delta}{\Delta - 2\epsilon} & c_- \leq x < c_1 \\ \frac{1-\delta}{2\epsilon} & c_1 \leq x \leq c_2 \\ \frac{\delta}{\Delta - 2\epsilon} & c_2 < x \leq c_+ \\ 0 & \text{otherwise} \end{cases}, \qquad (1)$$

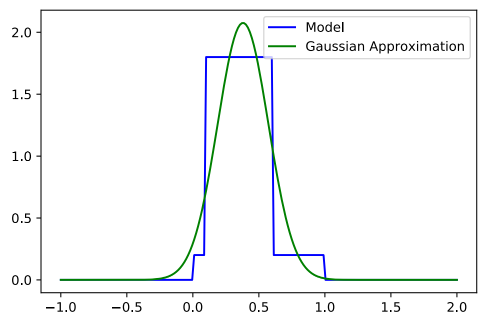

Figure 2: The Gaussian approximation of a 90% confidence interval on $[0.1, 0.6]$ where F ranges over $[0, 1]$.

Intuitively, this distribution captures our complete ignorance about the value of the objective function, except that it lies somewhere in the interval $[c_1, c_2]$ with probability $1-\delta$, and elsewhere with probability δ. This model is only valid if both lower and upper bounds of the objective function are known and finite, but we use this same information to achieve our PAC guarantees anyway.

Algorithm 1 PAC-OwNM

procedure PAC-OwNM$(\Theta, F, \mathscr{D}, \epsilon, \delta, \Delta) \to (\hat{\theta}^*, \hat{F}(\hat{\theta}^*))$
  input: Finite design space Θ, designer's parameterized black-box objective function F, distribution $\mathscr{D}$, error tolerance ϵ, failure probability δ, range Δ.
  output: Maximizing parameter $\hat{\theta}^*$ and designer's approximate black-box objective value $\hat{F}(\hat{\theta}^*)$.
  $m \leftarrow \left\lceil (\Delta/\epsilon)^2 \ln(2|\Theta|/\delta)/2 \right\rceil$ ▷ Hoeffding bound
  $\boldsymbol{X} \leftarrow \mathscr{D}^m$ ▷ Draw m samples from $\mathscr{D}$
  for $\theta \in \Theta$ **do**
    $\hat{F}[\theta] \leftarrow \textsc{Measure}(\theta, F, \boldsymbol{X})$
  end for
  $\hat{\theta}^* \leftarrow \arg\max_{\theta \in \Theta} \hat{F}[\theta]$
  return $(\hat{\theta}^*, \hat{F}[\hat{\theta}^*])$
end procedure

Exhaustive Search In the special case where the parameter space is finite and is searched exhaustively, it is straightforward to extend PAC guarantees on multiple independent measurements to a global guarantee across the search space. Algorithm 1 presents such an exhaustive search, and Theorem 3.1, which again invokes Hoeffding's inequality, describes the guarantee it achieves.

Consider a design space Θ and an objective function $F : \Theta \mapsto [c_-, c_+]$, with $\Delta \doteq c_+ - c_-$. Let $\hat{F}_1(\theta), \ldots, \hat{F}_m(\theta)$ be a sequence of m i.i.d. samples of $F(\theta)$ drawn from distribution $\mathscr{D}$. Hoeffding's inequality [14] upper bounds the probability that the absolute difference between the empirical mean and its expected value exceeds ϵ as $\mathbb{P}_{\boldsymbol{X} \sim \mathscr{D}}\left(\left|\mathbb{E}_{\mathscr{D}}[F(\theta)] - \frac{1}{m}\sum_{j=1}^m \hat{F}_j(\theta)\right| \geq \epsilon\right) \leq 2e^{-2\epsilon^2 m/\Delta^2}$.

Theorem 3.1. *Consider an OwNM problem s.t.* $\theta^* \in \arg\max_{\theta\in\Theta} F(\theta)$*, and assume the measurement noise is uniform piecewise constant. Algorithm 1 applied to such a problem outputs parameter* $\hat{\theta}^*$ *and value* $\hat{F}[\hat{\theta}^*]$ *such that* $\left|F(\theta^*) - \hat{F}[\hat{\theta}^*]\right| \leq \epsilon$ *with probability at least* $1-\delta$*, where* $\delta = \sum_{\theta\in\Theta} 2e^{-2\epsilon^2 m/\Delta^2}$.

Proof. Algorithm 1 explores the entire parameter space. By applying a union bound to Hoeffding's inequality, it follows that all confidence intervals (CI) hold simultaneously, with probability $1-\delta$.

The algorithm then returns the maximum measurement, namely $\hat{F}[\hat{\theta}^*]$. We can bound the difference between this output and the optimal value $F(\theta)^*$ as follows: $-\epsilon \leq F(\theta^*) - (F(\hat{\theta}^*) + \epsilon) \leq F(\theta^*) - \hat{F}[\hat{\theta}^*] \leq (\hat{F}[\theta^*] + \epsilon) - \hat{F}[\hat{\theta}^*] \leq \epsilon$. On the left, we used the CI around $\hat{F}[\hat{\theta}^*]$ and the fact that $F(\theta^*)$ is optimal; on the right, we used the CI around $\hat{F}[\theta^*]$ and the fact that $\hat{F}[\hat{\theta}^*] \geq \hat{F}[\theta^*]$. Therefore, $\left|F(\theta^*) - \hat{F}[\hat{\theta}^*]\right| \leq \epsilon$. □

To summarize, given ϵ, δ, and Δ, Algorithm 1 calculates the requisite number of samples m to ensure that the output is accurate up to ϵ with probability $1-\delta$.

Heuristic Search Algorithm 1 *exhaustively* searches the whole design space. But this is impossible for uncountable and continuous spaces, and becomes computationally prohibitive very fast even for finite spaces. Hence, we seek a methodology that can find a good approximation of θ^* using limited computational resources. That is the search problem we address presently.

Definition 3.2 (Budget-constrained optimization with noisy measurements (BCOwNM))**.** *Given a design space* Θ*, an objective function* $F : \Theta \to \mathbf{R}$*, a noise model* $\mathscr{D}$*, and a measurement operator* $M : \Theta \to P_F$*, where* P_F *is the space of all possible probability distributions over the range of F, approximate* $\theta^* \in \arg\max_{\theta\in\Theta} F(\theta)$ *invoking* M *no more than some budget* $B \in \mathbf{N}$ *times.*

Bayesian optimization (BO) is a common tool used to solve BCOwNM problems. BO works by constructing and maintaining a probabilistic model of the objective function. This model is used to decide where to take the next measurement, until the budget is exhausted.

Most implementations of BO employ a **Gaussian Process** (GP) to model the uncertainty surrounding the objective function. Technically, a GP is a collection of random variables, any finite number of which are jointly distributed by a Gaussian [15]. BO with a GP model has been shown to outperform state-of-the-art global optimization methods in a number of benchmark problems (e.g., [7]).

Standard GPs can handle measurements with i.i.d. Gaussian noise by adding a diagonal term to the co-variance matrix [15]. But there is no easy way to incorporate general noise models into GPs. We incorporate piecewise constant uniform noise into a GP, heuristically, using the Gaussian that best approximates $\hat{p}_F(x)$, by minimizing the Kullback-Leibler divergence, D_{KL}.

Definition 3.3 (Best Approximating Gaussian)**.** *Given a continuous (discrete) random variable* x *with pdf (pmf)* p*, the best approximating Gaussian* q^* *is one s.t.* $q^*(x) \in \arg\min_{q(x)} D_{\mathrm{KL}}\left[p(x), q(x)\right]$.

The following proposition is straightforward.

Proposition 3.1. *Given any distribution* $\hat{p}_F(x)$ *in the form of Equation 1, the best approximating Gaussian has mean* $\mu^* = c_1$ *and variance* $\sigma^* = 0$ *when* $c_1 = c_2$*, otherwise they are given by:*

$$\mu^* = \frac{\alpha}{2}\left(c_1^2 + c_+^2 - c_-^2 - c_2^2\right) + \frac{\beta}{2}\left(c_2^2 - c_1^2\right)$$

$$\sigma^* = \sqrt{\frac{\alpha}{2}\left(c_1^3 + c_+^3 - c_-^3 - c_3^2\right) + \frac{\beta}{3}\left(c_2^3 - c_1^3\right) - \mu^{*2}}$$

where $\alpha \doteq \frac{\delta}{\Delta - 2\hat{\epsilon}}$ *and* $\beta \doteq \frac{1-\delta}{2\hat{\epsilon}}$*. It is easy to show that* μ^* *and* σ^* *are precisely the mean and variance of* $\hat{p}_F(x)$.

For any valid confidence interval (i.e., one that lies completely within the range of the function F), the square root operand is necessarily positive, which means that a real-valued solution always exists. The singular case where $\hat{\epsilon} = 0$ occurs only if $c_1 = c_+$ or $c_2 = c_-$.

There are (at least) two ways to utilize Proposition 3.1 within BO. The first is to assume Gaussian noise with mean μ^* and variance σ^*. This is the best possible white noise approximation. We refer to this approach as $\mathcal{GP}$-$\mathcal{N}$.

One shortcoming of $\mathcal{GP}$-$\mathcal{N}$ is that repeated measurements at the same value of θ are likely to produce very similar confidence intervals that do not significantly improve the accuracy of the probabilistic model. Nonetheless, we consider $\mathcal{GP}$-$\mathcal{N}$ as a basline. As an alternative heuristic approach, we assume Gaussian noise with mean μ^* and variance 0, thereby avoiding repeated measurements. We refer to this approach as $\mathcal{GP}$-$\mathcal{M}$.

$\mathcal{GP}$-$\mathcal{M}$ may seem incorrect as compared to $\mathcal{GP}$-$\mathcal{N}$, but the width of the confidence interval is independent of θ and known, given the mean of the GP, so it is possible to recover the *correct* posterior at any point by adding to the posterior at that point a diagonal matrix that encodes the disregarded variance. But as we use the GP only to guide the search, this correction is not required.

The main advantage of $\mathcal{GP}$-$\mathcal{M}$ over the more straightforward $\mathcal{GP}$-$\mathcal{N}$ approach is improved exploration. As a second heuristic approach in this same spirit, we consider a third variant, which we call $\mathcal{GP}$, in which the variance is again zero, but the mean is set to the mean of the confidence interval, namely $\hat{F}$. This approach, while simple and intuitive, is not the best possible Gaussian approximation, since μ^* need not equal $\hat{F}$.

Algorithm 2 EMD_Measure

procedure EMD_MEASURE$(\Gamma_\theta, f, \boldsymbol{X}) \to \hat{F}(\theta)$
    input: Parameterized black-box game Γ_θ, designer's objective function f, and m samples $\boldsymbol{X}$.
    output: Designer's approximate black-box, worst-case objective value $\hat{F}(\theta)$.
    for $p \in P$ and $\boldsymbol{s} \in \boldsymbol{S}_\theta$ **do**
        $\hat{\boldsymbol{u}}_p(\boldsymbol{s}) \leftarrow \frac{1}{m}\sum_{j=1}^{m} \boldsymbol{u}_p(x_j, \boldsymbol{s})$
    end for
    $\hat{\Gamma}_\theta \leftarrow \langle P, \boldsymbol{S}_\theta, \hat{\boldsymbol{u}}(\cdot)\rangle$
    $\mathrm{SCC}_\epsilon(\hat{\Gamma}_\theta) \leftarrow \textsc{FindSCCs}(\hat{\Gamma}_\theta, \epsilon)$
    $\hat{F}(\theta) \leftarrow \min\limits_{Z \in \mathrm{SCC}_\epsilon(\hat{\Gamma}_\theta)} \min\limits_{\boldsymbol{s} \in Z} f(\boldsymbol{s}; \hat{\Gamma}_\theta)$
    return $\hat{F}(\theta)$
end procedure

4 EMPIRICAL MECHANISM DESIGN

Next, we describe how EMD can be viewed as an instance of black-box optimization. We then proceed to apply the aforementioned BO heuristics to two EMD applications.

Let Θ be an abstract design space over which a mechanism designer is free to choose parameters $\theta \in \Theta$. Conditioned on θ, we denote by Γ_θ the ensuing θ-parameterized game where, in Definition 2.1, we augment all strategies and utilities to depend on θ. We thus define $\Gamma_\theta \doteq \langle |P|, \boldsymbol{S}_\theta, \boldsymbol{u}_\theta(\cdot)\rangle$, where $\boldsymbol{S}_\theta$ and $\boldsymbol{u}_\theta(\cdot)$ denote the dependency of strategies and utilities on θ.

Overloading notation, we write $f(\boldsymbol{s}(\theta); \Gamma_\theta)$ to denote the value of the designer's objective in game Γ_θ at strategy profile $\boldsymbol{s}(\theta)$. In optimizing f, the designer assumes that players will play at (or near) equilibrium. So, as θ varies, the value of f likewise varies, as $\boldsymbol{s}(\theta)$ potentially moves from one equilibrium to another.

While solution concepts are meant to be predictive—that is, to predict the outcome of a game—most yield sets of equilibria, rather than unique predictions. Accordingly, a mechanism designer often faces a choice. In this work, we assume they choose a worst-case outcome, minimizing the value of f over the set $\mathrm{E}(\Gamma_\theta)$ of equilibria of Γ_θ. We denote this worst-case objective by $F^{\mathrm{E}}(\theta; \Gamma_\theta) = \min_{\boldsymbol{s} \in \mathrm{E}(\Gamma_\theta)} f(\boldsymbol{s}; \Gamma_\theta)$. More generally, for $\epsilon \geq 0$, letting $\mathrm{E}_\epsilon(\Gamma_\theta)$ denote the set of equilibria of Γ_θ up to ϵ, we define $F^{\mathrm{E}}_\epsilon(\theta; \Gamma_\theta) = \min_{Z \in \mathrm{E}_\epsilon(\Gamma_\theta)} \min_{\boldsymbol{s} \in Z} f(\boldsymbol{s}; \Gamma_\theta)$. (N.B. We usually write $\boldsymbol{s} \in \boldsymbol{S}_\theta$, suppressing the dependency of strategies on θ, because strategies always depend on θ, and θ is usually clear from context.)

The **worst-case EMD problem** is to optimize the black-box, worst-case objective function $F^{\mathrm{E}}(\theta; \Gamma_\theta)$. Our approach to this problem is to learn the designer's objective f up to some additive error ϵ, and then optimize the corresponding objective $F^{\mathrm{E}}_\epsilon(\theta; \Gamma_\theta)$. We now argue that this approach is reasonable, assuming SCCs as the equilibria.

Definition 4.1 (((θ, ϵ)-approximable objective function)**.** *Let Γ_θ and Γ'_θ be two θ-parameterized games, and let f be a designer's objective. If $\max_{\theta \in \Theta, \boldsymbol{s} \in \boldsymbol{S}_\theta} |f(\boldsymbol{s}; \Gamma_\theta) - f(\boldsymbol{s}; \Gamma'_\theta)| \leq \epsilon$, then we say that the objective f is ϵ-**approximable**.*

The conditions of Definition 4.1 do hold always hold. For example, consider two simultaneous auctions with reserve prices for two complementary goods. Here, a small change to one reserve price could cause revenue to drop to zero, if bidders no longer value either good.

The next theorem states that, for $\epsilon > 0$, when f is ϵ-approximable, a solution to the ϵ-worst-case EMD problem is an ϵ-approximate solution to the exact problem, assuming SCC as the equilibrium.

Theorem 4.1. *Let θ^* optimize $F^{\mathrm{SCC}}(\theta; \Gamma_\theta)$, and let θ^*_ϵ optimize $F^{\mathrm{SCC}}_\epsilon(\theta; \Gamma_\theta)$, for $\epsilon > 0$. If f is ϵ-approximable, then $\left|F^{\mathrm{SCC}}(\theta^*; \Gamma_\theta) - F^{\mathrm{SCC}}_\epsilon(\theta^*_\epsilon; \Gamma_\theta)\right| \leq \epsilon$.*

We show Theorem 4.1 via an intermediate lemma, which shows that when f is ϵ-approximable, $F^{\mathrm{E}}_\epsilon(\theta; \Gamma)$ well approximates $F^{\mathrm{E}}(\theta; \Gamma)$, assuming $\mathrm{E} = \mathrm{SCC}$: i.e., $F^{\mathrm{E}}(\Gamma) \doteq F^{\mathrm{SCC}}(\theta; \Gamma) = \min_{Z \in \mathrm{E}(\theta; \Gamma)} \min_{\boldsymbol{s} \in Z} f(\boldsymbol{s}; \Gamma)$.

Lemma 4.1. *If f is ϵ-approximable, then for all $\theta \in \Theta$, $\left|F^{\mathrm{SCC}}(\theta; \Gamma_\theta) - F^{\mathrm{SCC}}_\epsilon(\theta; \Gamma'_\theta)\right| \leq \epsilon$.*

Proof. First, note that the strongly connected components of a graph partition its vertices. Therefore, for all $\epsilon \geq 0$, it holds that $F^{\mathrm{SCC}}_\epsilon(\theta; \Gamma_\theta) = \min\limits_{Z \in \mathrm{SCC}_\epsilon(\Gamma_\theta)} \min\limits_{\boldsymbol{s} \in Z} f(\boldsymbol{s}; \Gamma_\theta) = \min\limits_{\boldsymbol{s} \in \boldsymbol{S}_\theta} f(\boldsymbol{s}; \Gamma_\theta)$. Now $F^{\mathrm{SCC}}(\theta; \Gamma_\theta) - F^{\mathrm{SCC}}_\epsilon(\theta; \Gamma'_\theta) = \min\limits_{\boldsymbol{s} \in \boldsymbol{S}_\theta} f(\boldsymbol{s}; \Gamma_\theta) - \min\limits_{\boldsymbol{s} \in \boldsymbol{S}_\theta} f(\boldsymbol{s}; \Gamma'_\theta) \leq \min\limits_{\boldsymbol{s} \in \boldsymbol{S}_\theta} f(\boldsymbol{s}; \Gamma_\theta) - \min\limits_{\boldsymbol{s} \in \boldsymbol{S}_\theta} (f(\boldsymbol{s}; \Gamma_\theta) - \epsilon) = \epsilon$. A similar argument shows that $F^{\mathrm{SCC}}(\theta; \Gamma_\theta) - F^{\mathrm{SCC}}_\epsilon(\theta; \Gamma'_\theta) \geq -\epsilon$. □

(Proof of Theorem 4.1). First, observe the following: $F^{\mathrm{SCC}}(\theta^*; \Gamma_\theta) \geq F^{\mathrm{SCC}}(\theta'; \Gamma_\theta) \geq F^{\mathrm{SCC}}_\epsilon(\theta'; \Gamma'_\theta) - \epsilon$. The first inequality follows from the fact that θ^* maximizes F, and the second, from Lemma 4.1. Via analogous reasoning, $F^{\mathrm{SCC}}_\epsilon(\theta'; \Gamma'_\theta) \geq F^{\mathrm{SCC}}_\epsilon(\theta^*; \Gamma'_\theta) \geq F^{\mathrm{SCC}}(\theta^*; \Gamma_\theta) - \epsilon$, assuming θ' maximizes $F^{\mathrm{SCC}}_\epsilon$. □

Note that Theorem 4.1 also holds for the best-case EMD problem, assuming SCC as the equilibrium, which is to optimize $F^{\mathrm{SCC}} = \max\limits_{Z \in \mathrm{SCC}_\epsilon(\Gamma_\theta)} \max\limits_{\boldsymbol{s} \in Z} f(\boldsymbol{s}; \Gamma_\theta)$.

In sum, it suffices to approximate f up to ϵ to obtain an ϵ-approximation of $F^{\mathrm{SCC}}(\theta; \Gamma)$. Furthermore, it suffices to optimize $F_{\epsilon}^{\mathrm{SCC}}(\theta; \Gamma_{\theta})$ up to ϵ to obtain an ϵ-approximate solution to the worst-case EMD problem. Such estimates can be computed with high probability via Algorithm 1, using the measurement procedure defined in Algorithm 2. But to obtain the usual guarantees, the number of samples m calculated in Line 4 of Algorithm 1 must be updated to $\left\lceil (\Delta/\epsilon)^2 \ln(2(|\Theta|+|\Gamma|)/\delta)/2 \right\rceil$ to account for a union bound over both the mechanism's and the game's parameters.

5 EXPERIMENTS

In this section, we present experimental results using our equilibrium estimation and BO search methodology. We note that all BO algorithms use a Mattern Kernel [15] in the underlying Gaussian Process and the expected improvement [7] acquisition function. We use $\mathcal{GP}$-$\mathcal{N}$, $\mathcal{GP}$, and uniform sampling over the range $[0.5, 1.5]$[8] as our baselines. Note that uniform sampling is a competitive search strategy for hyper-parameter optimization [16].

We experiment with our methodology in two settings. The first is in first-price auctions where analytical solutions are known, and the condition of Definition 4.1 holds, as the revenue curve is a known to be a continuous function of the reserve price. The second is in a setting with no known analytical solutions, and with no guarantees as to whether condition of Definition 4.1 holds or not. Nonetheless, we report anecdotally that an empirical analog of the condition is often met in our experiments.

First-price Auctions

As our first application domain, we consider **first-price auctions** in the standard independent, private value model [17]. There is one good up for auction, and n bidders, with bidders' values drawn from some joint distribution G. Bidders submit bids $b_1, \ldots, b_n$, and the highest bidder wins and pays their bid. As a proof of concept, we apply our BO search heuristics in this domain, in attempt to maximize revenue as a function of a **reserve price**, r, below which the good will not be sold.

In this controlled setting, analytical solutions are known in certain special cases. For example, when bidders' valuations are drawn i.i.d. uniformly on $[0, 1]$, the (unique) symmetric Bayes-Nash strategy is $s(v) \doteq \mathbf{1}(\mathbf{v} \geq \mathbf{r})\left[r^n/\mathbf{v}^{\mathbf{n}-\mathbf{1}} + (n-1)(v^n - r^n)/\mathbf{n}\mathbf{v}^{\mathbf{n}-\mathbf{1}}\right]$ [17]. Assuming $n = 2$ bidders play this equilibrium, the optimal reserve price is $1/2$, which yields a revenue of $5/12$.

Figure 3 depicts a comparison of $\mathcal{GP}$-$\mathcal{M}$ and $\mathcal{GP}$ against the baselines, $\mathcal{GP}$-$\mathcal{N}$ and uniform sampling, for $\epsilon \in \{0.02, 0.01\}$. We note that a similar behavior was observed for $\epsilon = 0.03$. Specifically, we plot a running maximum of revenue as a function of the number of reserve prices searched. Each point is an average over 30 trials, where a single trial consists of exploring 20 different reserve prices with different initial random points. All the BO heuristics are initialized with the same three initial random points.

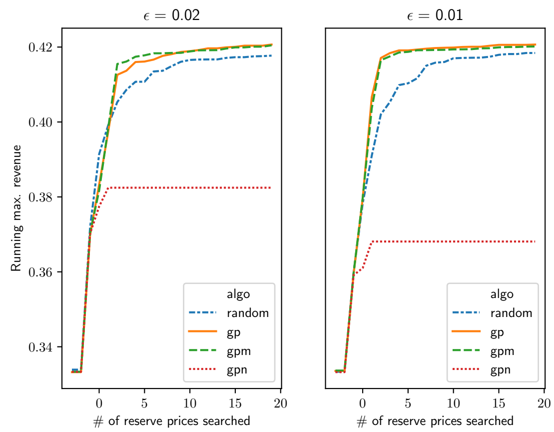

Figure 3: FPA search for optimal reserve price, $\delta = 0.1$

Our BO heuristic $\mathcal{GP}$-$\mathcal{M}$ consistently outperforms uniform sampling and $\mathcal{GP}$-$\mathcal{N}$ (whose predicted anomalous behavior was explained in Section 3), taking fewer measurements to achieve near-optimal values of revenue. The close performance between $\mathcal{GP}$-$\mathcal{M}$ and $\mathcal{GP}$ can be explained by the fact that δ is so small that the mean of a measurement's confidence interval is close to μ^*.

Advertisement Exchange

In this section, we illustrate our EMD methodology in a one-shot advertisement exchange game,[2] a game rich enough that analytic solutions are not readily available, but nonetheless amenable to analysis through sampling and simulation. The game is designed to model a scenario common in electronic advertisement platforms, such as Google's AdWords© and Amazon Sponsored Brands©.

We begin by describing, at a high-level,[3] the elements and dynamics of the game, and the strategies used by agents. We conclude with experiments showing that our algorithms more quickly accrue higher revenue than a baseline in an 8-dimensional parameter space.

AdX game in a nutshell. In the AdX game, agents play the role of advertising networks competing in an exchange for opportunities to show Internet users impressions (i.e.,

[2] A simplification of TAC AdX [18].

[3] A complete mathematical formalization of this game is provided in the supplemental material.

advertisements) needed to fulfill advertising campaigns. These impression opportunities (henceforth, impressions) are sold in real-time sequential auctions, as users arrive, but agents submit their bids in advance, placing different bids for different classes of users.

Users are characterized by *attributes*: e.g. GENDER, INCOME, AGE, etc., each of which is characterized by a set of *attribute values*: e.g., $\{\text{YOUNG}, \text{OLD}\}$. Each user belongs to a *market segment*, which is a list of values for (not necessarily all) attributes: e.g., $\langle\text{FEMALE}, \text{YOUNG}\rangle$.

A market segment M *matches* another market segment M' if all the attribute values in M' are present in M. For example, the market segments $\langle\text{FEMALE}, \text{YOUNG}, \text{HIGH INCOME}\rangle$ and $\langle\text{FEMALE}, \text{YOUNG}, \text{LOW INCOME}\rangle$ both match $\langle\text{FEMALE}, \text{YOUNG}\rangle$; however, $\langle\text{FEMALE}\rangle$ does not match $\langle\text{FEMALE}, \text{YOUNG}\rangle$, since attribute value YOUNG is missing from the former.

In the AdX game, agents target some market segments, but not others, as described by their (single) advertising campaign. Each advertising campaign $C_j = \langle I_j, M_j, R_j\rangle$ demands $I_j \in \mathbb{N}$ impressions in total, procured from users belonging to any market segment M' that matches the campaign's desired market segment M_j. A campaign's budget $R_j \in \mathbb{R}_+$ is the maximum amount the advertiser is willing to spend on those impressions. From the agent's point of view, the budget maps to its potential revenue.

To a first approximation then, agent j's goal, is to procure at least I_j impressions matching market segment M_j to fulfill C_j's so that it can earn revenue, which depends on R_j. More specifically, the value of a number of procured impressions z is determined via a sigmoidal function that maps z to a percentage of the budget: i.e., small values of z yield a small percentage of R_j, while values close to I_j yield values close to R_j. The non-linearity inherent in this function models complementarities, because it incentivizes agents to focus either on completely satisfying a campaign's demand, or not to bother satisfying it at all.

The AdX game is a one-shot game. To play, agents submit bids and spending limits for each market segment. Then, simulated users arrive at random, from the various market segments. For each user that arrives from market segment M, a second-price auction with a publicly known reserve price $r_M \in \mathbb{R}_+$ is held among all agents whose bids match M, and who have not yet reached their spending limit in M. Agent j's utility is computed as the difference between its revenue and its expenditure for all auctions.

Strategies. Designing bidding strategies for electronic ad exchanges (and ad auctions, more generally) is an active research area (e.g. [19, 20, 21]). The goal of these experiments is not to investigate the performance of state-of-the art bidding strategies, though this is certainly an interesting future research direction, but rather to test the methodology developed in this paper. Toward this end, we devised two heuristics which, we call Walrasian Equilibrium (WE) and Waterfall (WF). Detailed descriptions of these heuristics are provided in the supplemental material. Here, we present only their main ideas.

At a high level, both heuristics work by building an offline model of the market induced by the AdX game, which is then used to compute an allocation (assignment of impressions to campaigns) and prices (for impressions), based on which bids and limits are determined. The strategies differ in how this outcome—the allocation and prices—are computed. In a nutshell, the WE strategy searches for an outcome which forms a near-Walrasian equilibrium, i.e., an outcome whose prices provide little incentive for agents to relinquish their allocation. The study of equilibria and near-equilibrium computation in combinatorial markets (ad exchanges, being one example) is an active research area (e.g., [22, 23, 24]) with promising applications (e.g., [25, 26, 27]). The WF strategy works by simulating the arrival of impression opportunities in a fixed order that is endogenously determined by the campaigns present in the market. We call this strategy Waterfall because impressions are allocated *to* campaigns in descending order of budget-per-impression, and *from* market segments in ascending order of second-highest budget-per-impression. Hence, the final bid prices can be visualized as a descending waterfall-like structure.

Experimental setup. We assumed three user attributes: gender, age, and income level, each with two values. We then simulated a fixed number of impression opportunities, namely $K = 500$, distributed over 8 market segments, corresponding to all the possible combinations of attribute values: $\{\text{MALE}, \text{FEMALE}\} \times \{\text{YOUNG}, \text{OLD}\} \times \{\text{LOW INCOME}, \text{HIGH INCOME}\}$. The distribution π over these impression opportunities was constructed from statistical data at `www.alexa.com` [18].

Each agent's campaign $C_j = \langle I_j, M_j, R_j\rangle$ is determined as follows: A market segment M_j is drawn uniformly over all 20 possible market segments corresponding to combinations of user attributes of size 2 (e.g., $\langle\text{MALE}, \text{YOUNG}\rangle$) and 3 (e.g., $\langle\text{MALE}, \text{YOUNG}, \text{LOW INCOME}\rangle$). Given M_j, the demand $I_j := {}^{K\pi_{M_j}}/_N$, where $K\pi_{M_j}$ is the expected size of market segment M_j, and N is the number of agents in the game. Given I_j, the budget R_j is a noisy signal of the demand modeled by a beta distribution: $R_j \sim I_j(\mathcal{B}(\alpha = 10, \beta = 10) + 0.5)$.

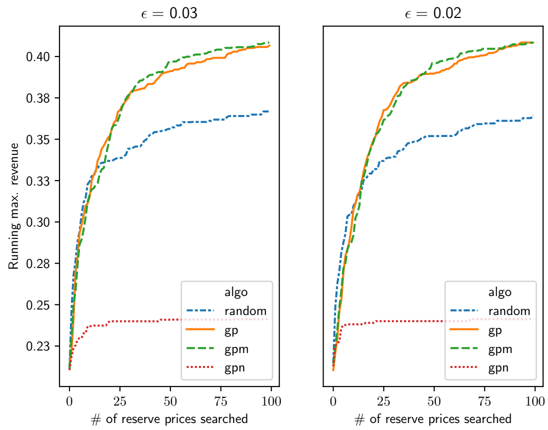

Figure 4: AdX search for optimal reserve prices, $\delta = 0.1$

The task is to find an 8-dimensional vector of reserve prices $\boldsymbol{r} \in \mathbb{R}^8$, consisting of one reserve price per market segment, that maximizes the auctioneer's revenue, up to a desired accuracy ϵ and specified failure probability δ. We experiment with $N = 4$ agents, and two possible strategies; in particular, for all p, $S_p = \{\text{WE, WF}\}$. Since by construction these strategies never bid higher than a campaign's budget-per-impression, $R_j/I_j = \mathcal{B}(\alpha = 10, \beta = 10) + 0.5 \in [0.5, 1.5]$, we bound the search for each reserve price to this same range, and hence search the space $[0.5, 1.5]^8$.

Even in this bounded region, the complexity of this task is substantial: for each candidate vector of reserve prices, we must first learn the corresponding game it induces, and then solve for the equilibria of this game. In this paper, we use Algorithm 2 for this purpose, which finds the set SCC_ϵ using Tarjan's algorithm [28], and then returns the minimum revenue among all SCC_ϵ equilibria. This is a computationally intensive task. Indeed, it took approximately 5 days to obtain these results using 4 machines, each with an E5-2690v4 (2.60GHz, 28Core) Xeon processor, and 1,536 GB of memory.

Experimental results. Figure 4 summarizes our AdX results. Specifically, we plot a running maximum of revenue as a function of the number reserve prices evaluated so far. Each plot is an average over 30 trials, where a single trial consists of exploring 100 different vectors of reserve prices, initialized at random. To ensure a fair comparison, all search algorithms are fed the same 10 initial points during each trial. These results are compared to uniform sampling. As in the first-price auction experiment, our BO heuristic $\mathcal{GP}$-$\mathcal{M}$ outperforms uniform sampling and $\mathcal{GP}$-$\mathcal{N}$, taking fewer measurements to achieve higher values of revenue. $\mathcal{GP}$-$\mathcal{M}$'s performance is on par with that of $\mathcal{GP}$, which once again can be explained by a sufficiently small value of δ. Note, however, the apparent difficulty in optimizing revenue in this game; $\mathcal{GP}$-$\mathcal{M}$ and $\mathcal{GP}$ take significantly more measurements to outperform the baselines. Nonetheless, these results show that, even under a fairly constrained budget, our BO heuristics are effective as compared to our baselines.

6 CONCLUSION

Our methodology tackles the fundamental problem of optimizing a designer's objective function in parametric systems inhabited by strategic agents. The fundamental assumptions are: 1) players play equilibria among an *a priori* known set of strategies, which in general depend on the parameters of the system; and 2) there is no analytical description available of either the strategies or the system, but only a simulator capable of producing data about game play. This framework captures modern, computationally intensive systems, such as electronic advertisement exchanges. The challenge then is two-fold: first, one must *learn equilibria* of a game for a fixed setting of parameters, and second, one must *search* the space of parameters for those that maximize the designer's function.

Our main contribution is a PAC-style framework to solve the former, while for the latter we enhance standard search routines for black-box optimization problems to include piecewise constant noise, precisely the kind of noise that characterizes PAC learners. We prove theoretical guarantees on the quality of the learned parameters when the parameter space is finite. We also demonstrate the practical feasibility of our methodology, first in a setting with known analytical solutions, and then in a stylized but rich model of advertisement exchanges with no known analytical solutions—precisely the kind of setting for which we devised our methodology—and show that we can find solutions of higher quality than standard baselines.

One potential criticism of this work is that SCC is not a well-motivated solution concept. Our main motivation in using SCC was to demonstrate our rich methodology, end-to-end. For this purpose, we sought a solution concept that was both approximable and relatively easy to compute. SCC is both approximable (sink is not) and amenable to fast computation (Nash is not). An interesting future research direction is to investigate the trade-offs between approximability and computability using other popular solution concepts, such as Nash equilibrium.

Acknowledgements

This work was supported by the NEC-AIST AI Cooperative Research Laboratory, and NSF Awards CMMI-1761546 and RI-1813444.

References

[1] Vorobeychik, Y. and Wellman, M. P. (2008) Stochastic search methods for Nash equilibrium approximation in simulation-based games. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 2* International Foundation for Autonomous Agents and Multiagent Systems pp. 1055–1062.

[2] Vorobeychik, Y., Wellman, M. P., and Singh, S. (2007) Learning payoff functions in infinite games. *Machine Learning,* **67**(1-2), 145–168.

[3] Daskalakis, C., Goldberg, P. W., and Papadimitriou, C. H. (2009) The complexity of computing a Nash equilibrium. *SIAM Journal on Computing,* **39**(1), 195–259.

[4] Wellman, M. P. (2006) Methods for empirical game-theoretic analysis. In *AAAI* pp. 1552–1556.

[5] Areyan Viqueira, E., Greenwald, A., Cousins, C., and Upfal, E. (2019) Learning Simulation-Based Games from Data. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems* International Foundation for Autonomous Agents and Multiagent Systems pp. 1778–1780.

[6] Valiant, L. G. (1984) A theory of the learnable. *Communications of the ACM,* **27**(11), 1134–1142.

[7] Snoek, J., Larochelle, H., and Adams, R. P. (2012) Practical bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems* pp. 2951–2959.

[8] Nisan, N. and Ronen, A. (2001) Algorithmic mechanism design. *Games and Economic behavior,* **35**(1-2), 166–196.

[9] Sandholm, T. (2003) Automated mechanism design: A new application area for search algorithms. In *International Conference on Principles and Practice of Constraint Programming* Springer pp. 19–36.

[10] Guo, M. and Conitzer, V. (2010) Computationally feasible automated mechanism design: General approach and case studies. In *Twenty-Fourth AAAI Conference on Artificial Intelligence*.

[11] Balcan, M.-F., Sandholm, T., and Vitercik, E. (2018) A general theory of sample complexity for multi-item profit maximization. In *Proceedings of the 2018 ACM Conference on Economics and Computation* ACM pp. 173–174.

[12] Nash, J. F. (1950) Equilibrium points in n-person games. *Proceedings of the national academy of sciences,*.

[13] Goemans, M., Mirrokni, V., and Vetta, A. (2005) Sink equilibria and convergence. In *Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on* IEEE pp. 142–151.

[14] Hoeffding, W. (1963) Probability inequalities for sums of bounded random variables. *Journal of the American statistical association,* **58**(301), 13–30.

[15] Rasmussen, C. E. (2003) Gaussian processes in machine learning. In *Summer School on Machine Learning* Springer pp. 63–71.

[16] Bergstra, J. and Bengio, Y. (2012) Random search for hyper-parameter optimization. *Journal of Machine Learning Research,* **13**(Feb), 281–305.

[17] Krishna, V. (2009) Auction theory, Academic press, .

[18] Schain, M. and Mansour, Y. (2013) Ad exchange–proposal for a new trading agent competition game. In *Agent-Mediated Electronic Commerce. Designing Trading Strategies and Mechanisms for Electronic Markets* pp. 133–145 Springer.

[19] Gerakaris, S. and Ramamoorthy, S. (2018) Learning Best Response Strategies for Agents in Ad Exchanges. In *European Conference on Multi-Agent Systems* Springer pp. 77–93.

[20] Even Dar, E., Mirrokni, V. S., Muthukrishnan, S., Mansour, Y., and Nadav, U. (2009) Bid optimization for broad match ad auctions. In *Proceedings of the 18th international conference on World wide web* ACM pp. 231–240.

[21] Cary, M., Das, A., Edelman, B., Giotis, I., Heimerl, K., Karlin, A. R., Mathieu, C., and Schwarz, M. (2007) Greedy bidding strategies for keyword auctions. In *Proceedings of the 8th ACM conference on Electronic commerce* ACM pp. 262–271.

[22] Feldman, M., Gravin, N., and Lucier, B. (2016) Combinatorial walrasian equilibrium. *SIAM Journal on Computing,* **45**(1), 29–48.

[23] Garg, R., Kapoor, S., and Vazirani, V. (2004) An auction-based market equilibrium algorithm for the separable gross substitutability case. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques* pp. 128–138 Springer.

[24] Blumrosen, L. and Nisan, N. (2007) Combinatorial auctions. *Algorithmic game theory,* **267**, 300.

[25] MacKie-Mason, J. K., Osepayshvili, A., Reeves, D. M., and Wellman, M. P. (2004) Price prediction strategies for market-based scheduling. *AAAI,*.

[26] Areyan Viqueira, E., Greenwald, A., and Naroditskiy, V. (2017) On Approximate Welfare- and Revenue-Maximizing Equilibria for Size-Interchangeable Bidders. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems* International Foundation for Autonomous Agents and Multiagent Systems pp. 1466–1468.

[27] Lehmann, B., Lehmann, D., and Nisan, N. (2006) Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior,* **55**(2), 270–296.

[28] Tarjan, R. (1972) Depth-first search and linear graph algorithms. *SIAM journal on computing,* **1**(2), 146–160.