

---

# A Learning Framework for Distribution-Based Game-Theoretic Solution Concepts

---

**Tushant Jha**

International Institute of Information Technology, Hyderabad  
particle.mania@gmail.com

**Yair Zick**

National University of Singapore, Singapore  
zick@comp.nus.edu.sg

## Abstract

The past few years have seen several works on learning economic solutions from data; these include optimal auction design, function optimization, stable payoffs in cooperative games and more. In this work, we provide a unified learning-theoretic methodology for modeling such problems, and establish tools for determining whether a given economic solution concept can be learned from data. Our learning theoretic framework generalizes a notion of function space dimension — the graph dimension — adapting it to the solution concept learning domain. We identify sufficient conditions for the PAC learnability of solution concepts, and show that results in existing works can be immediately derived using our methodology. Finally, we apply our methods in other economic domains, yielding a novel notion of PAC competitive equilibrium and PAC Condorcet winners.

## 1 Introduction

Recent years have seen widespread application of learning-theoretic notions in economic domains. Rather than assuming full knowledge of the underlying domain (or a prior over the domain space), one assumes access to a dataset of past instances, and employs learning-theoretic tools in order to obtain approximate solutions. Consider the following simple example of learning a solution from data: we wish to find the maximum of a function  $f : \mathbb{R}^n \rightarrow \mathbb{R}$ ; we do not have access to  $f$ , but rather to a dataset of the form  $\langle \vec{x}_1, f(\vec{x}_1) \rangle, \dots, \langle \vec{x}_m, f(\vec{x}_m) \rangle$ . One way to find a likely candidate point would be to use classic learning-theoretic tools [1], learn an approximation  $f^*$  of  $f$ , and compute the maximum of  $f^*$ ; however, this goes above and beyond the problem requirement: the approximability of  $f$  depends on its *hypothesis class* (whether  $f$  is a linear function, a two-layer neural network etc.), and on the approximation robustness. A much simpler solution is available: if the number of samples is sufficiently large, taking the *empirical maximum* of  $f$  over the dataset — i.e.  $\vec{x}^* \in \operatorname{argmax}_j \{f(\vec{x}_j) : j \in [m]\}$  — yields a point that is likely to be greater in value than any future point sampled from the same distribution as the original dataset.

The same reasoning applies to other economic solutions: a naive approach to inferring solutions from data would be to learn an approximate model (e.g. learn a function  $f^*$  which approximates  $f$  in the maximization example above), and then try generating solutions for the approximate model. However, as has been shown in the literature, learning an approximate model may:

1. be insufficient for generating ‘good’ solutions (this is the case in [32])
2. require an exponential number of samples, whereas directly learning solutions is easy. Indeed, finding a payoff in the core of TU games is easy [4, 11], while PAC learning

cooperative games requires an exponential number of samples [7]; this is also the case for finding an empirical maximum in the example above.

Recent works directly learn solutions to classic optimization problems, as well as solutions in game-theoretic domains. These lines of work have progressed more or less independently, proving that solutions in a specific problem domain can (or cannot) be efficiently inferred from data; however, there has been no attempt to provide a unified theory of learning solution concepts from data. This is where our work comes in.

## 1.1 Our Contributions

We begin by establishing a learning-theoretic framework for learning solution concepts from data. Unlike classic learning problem spaces, solution concepts do not inhabit the same space as the observed samples (e.g. when learning an approximate maximum, the function space is  $\mathbb{R}^n \rightarrow \mathbb{R}$ , whereas the solution space is  $\mathbb{R}^n$ ). In Section 2, we define the *solution dimension*: this quantity depends on both the hypothesis class of the underlying game and the solution space. The solution dimension generalizes the graph dimension [18] in PAC learning, and serves a similar purpose: if the solution dimension is low, then a distribution-based solution can be efficiently learned from samples. Drawing on notions of shattering from VC dimension, we introduce *solution concept shattering* which is used to bound the solution dimension in various domains. We also show that the existence of a consistent solution and a low graph dimension are sufficient conditions for PAC learning solutions, simplifying technical learnability arguments in existing works, as well as paving the way for a straightforward learnability approach of other solution concepts. In Section 3, we apply our methodology to immediately derive sample complexity bounds on learning solutions for hedonic games, as well as for two novel domains: market equilibria, and Condorcet winners in voting.

## 1.2 Related Work

Several recent works study learning solutions from data; these include solutions in cooperative games [7, 11, 24, 32], combinatorial auctions [6, 8, 14, 16, 19, 25, 33], voting and judgment aggregation [12, 35], envy-free allocation [5], and optimization [9, 10, 30]. Some of these works offer low-error approximation guarantees with respect to an optimal solution, such as estimating the maximum [9, 10], maximizing revenue in mechanism design [16], or finding an election winner [12]; our work focuses on solutions that minimize *expected loss with respect to sampled data*, as is the case when learning the core of a cooperative game [7, 24, 32], reserve prices in auctions [6, 26], or approximately efficient allocations [14]. While some of the above works explicitly explore the dimension of the solution space, they do not offer the full generality of our model.

Our analysis of the underlying solution space utilizes recent learning-theoretic tools [18], yielding an extension of classical function dimension measures such as the VC dimension [1, 34], and the graph dimension [18]. Our results generalize the General Learning problem discussed in [31], as learning solution concepts can also involve some global properties playing a role in the loss function.

## 1.3 Classic PAC Learning

For the sake of completeness, we provide a brief overview of the PAC learning model. A *learning problem* is defined over an instance space  $\mathcal{X}$  and a set of functions (the *hypothesis class*)  $\mathbb{H} \subseteq \mathcal{Y}^{\mathcal{X}}$  ( $\mathcal{Y}$  is the *label space*). Let  $\mathcal{D}$  be a distribution over  $\mathcal{X} \times \mathcal{Y}$ ; we let the *loss* of  $h \in \mathbb{H}$  given  $\mathcal{D}$  be  $L_{\mathcal{D}}(h) = \Pr_{(x,y) \sim \mathcal{D}}[h(x) \neq y]$ . Given a set of  $m$  i.i.d. samples  $T = \langle x_j, y_j \rangle_{j \in [m]}$  from  $\mathcal{D}$ , the *empirical loss* of  $h \in \mathbb{H}$  is  $\hat{L}_T(h) = \frac{1}{m} \sum_{j=1}^m \mathbb{1}(h(x_j) \neq y_j)$ . We assume that  $\mathcal{D}$  is a distribution over  $\mathcal{X}$ , where every  $x \in \mathcal{X}$  is evaluated by some unknown  $c \in \mathbb{H}$ ; this is referred to as the *realizable case*, in which there is some  $h \in \mathbb{H}$  for which  $L_{\mathcal{D}}(h) = 0$ , and there is at least one hypothesis  $h \in \mathbb{H}$  for which  $\hat{L}_T(h) = 0$  for any  $T \subseteq \mathcal{X}$ . An algorithm  $\mathcal{A}$  is a PAC learner for  $\mathbb{H}$  if there is some  $m_0$  polynomial in  $\frac{1}{\varepsilon}, \frac{1}{\delta}$  and the natural problem parameters, such that for any distribution  $\mathcal{D}$  and any set of  $m \geq m_0$  samples  $T$  sampled i.i.d. from  $\mathcal{D}$ ,  $\mathcal{A}$  outputs a hypothesis  $h^* \in \mathbb{H}$  (which is a function of  $T$ , but not of  $\mathcal{D}$ ) such that  $\Pr_{T \sim \mathcal{D}^m}[L_{\mathcal{D}}(h^*) \geq \varepsilon] < \delta$ .

For binary hypothesis classes (where the label space is  $\mathcal{Y} = \{\pm 1\}$ ), the VC dimension [34] characterizes the sample complexity of  $\mathbb{H}$ . The sample complexity required by any PAC learning algorithm

for  $\mathbb{H}$  is upper and lower-bounded by the VC dimension of  $\mathbb{H}$ . This is achieved by an algorithm that outputs a *consistent* hypothesis, i.e. one which minimizes the empirical loss  $\hat{L}_T(h)$  w.r.t. a sample  $T$ .

**Definition 1.1.** Given a hypothesis class  $\mathbb{H}$ , a set  $C \subseteq \mathcal{X}$  is said to be shattered if for any binary labeling  $b : C \rightarrow \{0, 1\}$  there exists some  $h \in \mathbb{H}$  such that  $h(x) = b(x)$  for all  $x \in C$ . The VC dimension of  $\mathbb{H}$ , or  $VC(\mathbb{H})$ , is the size of the largest set  $C \subseteq \mathcal{X}$  that is shattered by  $\mathbb{H}$ .

For example, if the hypothesis class is the set of all linear classifiers over  $\mathbb{R}^n$ , its VC dimension is  $\mathcal{O}(n)$  [1]. Theorem 1.2 relates the VC dimension and the PAC learnability of  $\mathbb{H}$ .

**Theorem 1.2.** *There exists absolute constants  $\alpha_1$  and  $\alpha_2$ , such that for a hypothesis class  $\mathbb{H}$ , the sample complexity of  $\mathbb{H}$  with respect to  $\varepsilon$  and  $\delta$  (denoted  $m(\varepsilon, \delta)$ ) is*

$$\frac{\alpha_1}{\varepsilon} \left( VC(\mathbb{H}) + \log \left( \frac{1}{\delta} \right) \right) \leq m(\varepsilon, \delta) \leq \frac{\alpha_2}{\varepsilon} \left( \log \left( \frac{1}{\varepsilon} \right) VC(\mathbb{H}) + \log \left( \frac{1}{\delta} \right) \right)$$

Theorem 1.2 can be slightly generalized to the following claim: for any two functions  $f, g \in \mathbb{H}$ , the empirical loss on an i.i.d. sample of more than  $\frac{\alpha_2}{\varepsilon} \left( \log \left( \frac{1}{\varepsilon} \right) VC(\mathbb{H}) + \log \left( \frac{1}{\delta} \right) \right)$  points, is close within  $\varepsilon$  to the statistical loss (ie.  $\Pr_{x \sim \mathcal{D}} [f(x) \neq g(x)]$ ).

The case where samples are labelled by some arbitrary function  $c$  (not necessarily in  $\mathbb{H}$ ) is also known as the *agnostic* case; however, as a result of the uniform convergence results, if an algorithm  $\mathcal{A}$  outputs a hypothesis  $h^*$  that minimizes empirical risk —  $\forall h \in \mathbb{H} : \hat{L}_T(h^*) \leq \hat{L}_T(h)$  — the statistical error is  $\leq \varepsilon$ :  $\Pr_{T \sim \mathcal{D}^m} [L_{\mathcal{D}}(h^*) \geq \min_{h \in \mathbb{H}} L_{\mathcal{D}}(h) + \varepsilon] < \delta$ . Therefore, as discussed in [18], uniform convergence is a powerful tool for bounding statistical loss.

## 1.4 Game-Theoretic Solution Concepts

In what follows, we briefly introduce the solution concepts discussed in this work. In all scenarios below, we have a set of *players*  $N = \{1, \dots, n\}$ , with preferences over outcomes induced in some manner; our objective is to obtain a *solution* with some desirable properties.

### 1.4.1 Hedonic Games

In hedonic games [13, Chapter 15], each player  $i \in N$  has a complete, transitive preference order  $\succ_i$  over coalitions in  $N$  that contain it. Solutions are partitions (also referred to as *coalition structures*) of  $N$ ; a coalition structure  $\pi$  is *blocked* by a coalition  $S \subseteq N$  if all members of  $S$  prefer  $S$  over the coalition they are in (denoted  $\pi(i)$ ), i.e.  $S \succ_i \pi(i)$  for all  $i \in S$ . The *core* of a hedonic game is the set of *stable* coalition structures: they cannot be blocked by any coalition  $S \subseteq N$ . It is often assumed that players' preferences over subsets are induced by a *cardinal* utility function  $v_i : 2^N \rightarrow \mathbb{R}_+$ ; in this case,  $S \succ_i T$  if and only if  $v_i(S) > v_i(T)$ .

### 1.4.2 Competitive Equilibria in Fisher Markets

We are given a set of  $k$  indivisible goods  $G = \{g_1, \dots, g_k\}$ . Each player  $i \in N$  values bundles of goods in  $G$  according to  $v_i : 2^G \rightarrow \mathbb{R}_+$ , where  $v_i(\emptyset) = 0$  for all  $i \in N$ . A *market outcome* is a tuple  $\langle \pi, \vec{p} \rangle$ , where  $\pi$  is a partition of  $G$  into  $n$  disjoint bundles (some of them may be empty), with  $\pi(i)$  assigned to player  $i$ ;  $\vec{p} \in \mathbb{R}^k$  is a price vector, denoting the price of each item in  $G$ . In these markets, known as *Fisher markets* [15], we assume that each player  $i$  has a *budget*  $\beta_i \in \mathbb{R}_+$ . Given a price vector  $\vec{p} \in \mathbb{R}^k$ , the *affordable set* of player  $i$  is the set of all bundles whose total price is less than  $\beta_i$ :

$$\mathcal{A}_i(\vec{p}, \beta_i) = \left\{ S \subseteq G : \sum_{g_j \in S} p_j \leq \beta_i \right\}.$$

An outcome  $\langle \pi, \vec{p} \rangle$  is a *competitive equilibrium* if for all  $i \in N$ ,  $\pi(i) \in \mathcal{A}_i(\vec{p}, \beta_i)$ , and  $\forall S \in \mathcal{A}_i(\vec{p}, \beta_i)$ ,  $v_i(\pi(i)) \geq v_i(S)$ .

### 1.4.3 Condorcet Winners

Consider a set of *voters*  $N = \{1, \dots, n\}$ , each with a *preference order*  $\succ_i$  over some finite set of *candidates*  $C$ . Given two candidates  $c, c' \in C$ , we define  $B(\succ, c', c) = 1$  iff a majority of voters

prefer  $c'$  to  $c$  under  $\succ$ . A candidate  $c^*$  is a *Condorcet winner* iff  $B(\succ, c^*, c) = 1$  for every other candidate  $c \in C$ .

## 2 A PAC Framework for Distribution-Based Solution Concepts

As described in Section 1.4, a *solution concept* or an *equilibrium concept* characterizes a subset of its solution space satisfying some natural desiderata. *Games* are mappings from some domain  $\mathcal{X}$  to a *label space*  $\mathcal{Y}$ . For example, in hedonic games, a game is a set of functions  $v_i : 2^N \rightarrow \mathbb{R}$  (for every  $i \in N$ ), mapping from subsets of players to real values; thus,  $\mathcal{X} = 2^N$  and  $\mathcal{Y}$  consists of vectors of the form  $(v_i(S))_{i \in S}$  for every  $S \subseteq N$ . We assume no knowledge of the actual game  $g$ , except for the hypothesis class it belongs to; we only observe samples of the game's evaluation on points in  $\mathcal{X}$ .

Constraints characterizing solution concepts are often universal quantifiers over a *local loss function*  $\lambda$ . In hedonic games, we define  $\lambda : 2^N \times \mathbb{G} \times \Pi(N) \rightarrow \{0, 1\}$ , where  $\lambda(S, \vec{v}, \pi) = 0$  iff  $v_i(S) \leq v_i(\pi(i))$  for all  $i \in N$ . Thus,  $\pi$  is in the core iff  $\lambda(S, \vec{v}, \pi) = 0$  for all  $S \subseteq N$ . Similarly, the maxima of a function satisfy  $x^* \in \operatorname{argmax} f(x) \iff \forall x : f(x) \leq f(x^*)$ ; in particular,  $\lambda(x, f, x^*) = 0$  iff  $f(x^*) \geq f(x)$ . Solution concepts that can be defined via a local loss function  $\lambda$  readily admit a distributional variant: we require that the expected loss as measured by  $\lambda$  is low, with respect to a distribution  $\mathcal{D}$  over the domain  $\mathcal{X}$ ; i.e.  $\Pr_{x \sim \mathcal{D}} [\lambda(x, g, s) = 0] \geq 1 - \varepsilon$ , for some  $\varepsilon \in (0, 1)$ .

More formally, an instance of the STATISTICALSOLUTION problem is a tuple  $\Psi = (\mathcal{X}, \mathcal{Y}, \mathbb{G}, \mathbb{S}, \lambda)$ . Here  $\mathcal{X}$  is the instance space;  $\mathcal{Y}$  is the codomain (or label) space;  $\mathbb{G} \subseteq \mathcal{Y}^{\mathcal{X}}$  is the class of games;  $\mathbb{S}$  is the solution space; finally,  $\lambda : \mathcal{X} \times \mathbb{G} \times \mathbb{S} \rightarrow \{0, 1\}$  measures *local loss*. In standard PAC learning (Section 1.3),  $\mathbb{G} = \mathbb{S}$  and  $\lambda(x, g_0, g_1) = 0 \iff g_0(x) = g_1(x)$ .

Given a game  $g \in \mathbb{G}$  and  $m$  points  $T = \langle (x_j, g(x_j)) \rangle_{j=1}^m$ , the *empirical error* (or empirical risk) of  $s \in \mathbb{S}$  is

$$\hat{L}_T(g, s) = \frac{1}{m} \sum_{(x_j, g(x_j)) \in T} \lambda(x_j, g, s),$$

and the statistical error (or statistical risk) as  $L_{\mathcal{D}}(g, s) = \mathbb{E}_{T \sim \mathcal{D}^m} [\lambda(x_j, g, s)]$ . A PAC solver for a STATISTICALSOLUTION  $\Psi$  is an algorithm  $\mathcal{L}$  whose input is a list  $T = \langle x_j, g(x_j) \rangle_{j=1}^m$  of  $m$  values  $x_j \in \mathcal{X}$  labelled by some unknown  $g \in \mathbb{G}$ , and whose output is a solution  $s^* \in \mathbb{S}$ ; its sample complexity, denoted  $m_{\mathcal{L}}(\varepsilon, \delta)$ , is the minimal number of samples required such that for any  $m \geq m_{\mathcal{L}}(\varepsilon, \delta)$ ,  $\Pr_{T \sim \mathcal{D}^m} [L_{\mathcal{D}}(g, s^*) > \varepsilon] < \delta$ . We let  $m_{\Psi}^{\text{PAC}}(\varepsilon, \delta)$  be the minimal sample complexity  $m_{\mathcal{L}}(\varepsilon, \delta)$  required by any PAC solver  $\mathcal{L}$  for  $\Psi$ .

### 2.1 Consistent Solvers and Barriers of Indistinguishability

In standard PAC learning, *consistent* or *empirical risk minimizing (ERM)* solvers play an important role; these are algorithms that minimize empirical error ( $\hat{L}_T(h)$ ) on the training sample. As discussed in Section 1.3, for binary functions, consistent algorithms are PAC learners whose sample complexity is bounded by the VC dimension. We first define a notion of consistency for solution concepts.

**Definition 2.1.** An algorithm  $\mathcal{A}_m : (\mathcal{X} \times \mathcal{Y})^m \rightarrow \mathbb{S}$  is said to be a *consistent solver* for the STATISTICALSOLUTION problem  $\Psi$  if for all  $g \in \mathbb{G}$ , and for any set of  $m$  samples  $T_m = \langle x_j, g(x_j) \rangle_{j=1}^m$ ,  $\mathcal{A}_m$  takes as input  $T_m$  and outputs a solution  $s^* = \mathcal{A}_m(T_m) \in \mathbb{S}$  such that the empirical loss of the solution  $s^*$  over  $T_m$  is 0:  $\hat{L}_{T_m}(g, s^*) = 0$ . In other words, for any input batch  $T_m$  labelled by some underlying function  $g \in \mathbb{G}$ , the algorithm returns a solution that has zero loss w.r.t  $g$  on all points in the input sample.

The definition of consistent solving presents a subtle yet crucial departure from the corresponding result in standard PAC learning. In PAC learning, since  $\lambda(x, g, h) = \mathbb{1}[g(x) \neq h(x)]$ , if there are two functions  $g_0, g_1 \in \mathbb{G}$  such that  $g_0(x) = g_1(x)$  for a point  $x \in \mathcal{X}$ , then for any hypothesis  $h$ ,  $\lambda(x, g_0, h) = \lambda(x, g_1, h)$ . Therefore, even if two functions  $g_0, g_1 \in \mathbb{G}$  generate an equivalent sample  $T = \langle x_j, y_j \rangle_{j=1}^m = \langle x_j, g_0(x_j) \rangle_{j=1}^m = \langle x_j, g_1(x_j) \rangle_{j=1}^m$ , if  $h$  is consistent with samples  $(x_j, y_j)$  in  $T$ , then it is consistent with both  $g_0$  and  $g_1$ . In fact, this implies that, time complexity considerations aside, a consistent solution always exists in standard learning, and can be found via exhaustive search. This is not the case in solution concept learning; two functions  $g_0, g_1 \in \mathbb{G}$  may generate an equivalent

sample  $T = \langle x_j, y_j \rangle_{j=1}^m$ , yet disagree on a solution (this is noted in prior works [24, 32]). Intuitively, this occurs since game-theoretic solutions treat unobserved regions of players' preferences. For example, in PAC market equilibria, one must inevitably set prices for unobserved goods, and assign bundles to players without knowing what their value might be; in hedonic games, a partition of players may contain subsets completely unobserved in the sample data. It is often useful to think of domains where this issue does not occur, as captured in the following definition. Given a labelled sample of  $m$  points  $T \in (\mathcal{X} \times \mathcal{Y})^m$ , let  $\mathbb{G}|_T$  be the set of games in  $\mathbb{G}$  which agree with  $T$ .

**Definition 2.2.** A STATISTICALSOLUTION  $\Psi$  is said to satisfy the *consistent solvability criterion* if for all  $m$  and all  $T \subseteq (\mathcal{X} \times \mathcal{Y})^m$ , there exists some  $s \in \mathbb{S}$  such that for all  $g \in \mathbb{G}|_T$ , the empirical loss  $\hat{L}_T(g, s)$  is 0.

## 2.2 A Dimension Theory for Game-Theoretic Solutions

We now present a novel definition of dimension for the PAC solution setting, and use it to bound the sample complexity for finding solutions to problem domains.

**Definition 2.3** (Solution-based Dimension). Given some  $C \subseteq \mathcal{X}$ , we say the set  $C$  is S-shattered in  $\Psi$  if there exists a game  $g \in \mathbb{G}$ , such that for every binary labelling  $b : C \rightarrow \{0, 1\}$  there exists a solution  $s \in \mathbb{S}$  (that may depend on  $b$ ) such that for all  $x \in C$ ,  $\lambda(x, g, s) = b(x)$ .

The Solution-based dimension of  $\Psi$ , denoted  $Sd(\Psi)$ , is the size of the largest set S-shattered in  $\Psi$ , and  $(C, g)$  as the corresponding shattering witness.

$Sd(\Psi)$  bounds the sample complexity of consistent solutions for  $\Psi$  (i.e.  $m^{PAC}(\varepsilon, \delta)$ ); however, we first prove a stronger claim, using the idea of uniform convergence discussed in Section 1.3. If we define the sample complexity for uniform convergence  $m^{UC}(\varepsilon, \delta)$  as the number of samples required such that the empirical loss of any solution is  $\varepsilon$ -close to its statistical loss, then  $m^{UC}(\varepsilon, \delta)$  is polynomially dependent on the solution dimension of the problem.

**Theorem 2.4.** *There are universal constants  $\alpha_1$  and  $\alpha_2$ , such that if  $Sd(\Psi) = d$ , then for a sample of  $m \geq \alpha_1 \frac{d + \log(\frac{1}{\delta})}{\varepsilon^2}$  points  $T = \langle x_j, y_j \rangle_{j=1}^m$ ,*

$$\Pr_{T \sim \mathcal{D}^m} [\exists g \in \mathbb{G}|_T, s \in \mathbb{S} : |\hat{L}_T(g, s) - L_{\mathcal{D}}(g, s)| > \varepsilon] < \delta.$$

*Furthermore, if a solution  $s^*$  is consistent, i.e.  $\hat{L}_T(g, s^*) = 0$ , then for any  $m$  greater than  $\frac{\alpha_2}{\varepsilon} (\log(\frac{1}{\varepsilon})d + \log(\frac{1}{\delta}))$ , we have that  $\Pr_{T \sim \mathcal{D}^m} [L_{\mathcal{D}}(g, s^*) > \varepsilon] < \delta$ .*

Note that in particular,  $m^{PAC}(\varepsilon, \delta) \leq m^{UC}(\varepsilon, \delta)$ , and both are polynomially dependent on  $Sd(\Psi)$ ,  $\frac{1}{\varepsilon}$  and  $\log \frac{1}{\delta}$ .

As a useful sanity check, we observe that  $Sd$  collapses to the classic VC dimension when learning classifiers: when  $\mathbb{S} = \mathbb{G} = \mathbb{H} \subseteq 2^{\mathcal{X}}$ , then  $Sd(\Psi) = VC(\mathbb{H})$ . Similarly, when  $\mathbb{S} = \mathbb{G} = \mathbb{H} \subseteq \mathcal{Y}^{\mathcal{X}}$  (i.e. for multiclass learning problems with a general domain  $\mathcal{Y}$ ),  $Sd$  collapses to the graph dimension [18]. We now observe few immediate corollaries of the above uniform convergence result.

**Corollary 2.5.** *Given a STATISTICALSOLUTION problem  $\Psi = \langle \mathcal{X}, \mathcal{Y}, \mathbb{G}, \mathbb{S} \rangle$ :*

**Simultaneous Constraints:** *if multiple local loss functions  $\lambda_1, \dots, \lambda_k$  need to be simultaneously approximated within  $\varepsilon$ , i.e.  $\forall i \in [k] : |\hat{L}_i(g, s) - L_{i\mathcal{D}}(g, s)| < \varepsilon$ , then the sample complexity of finding a solution satisfying all of them is in  $\mathcal{O}(\max_{i \in [k]} \{m_i^{UC}(\varepsilon, \delta)\})$ .*

**Separable Conjunctions:** *if there are local constraints  $\lambda_1$  over  $\mathbb{S}_1$ , and  $\lambda_2$  over  $\mathbb{S}_2$ , where  $\mathbb{S}_{\Psi} = \mathbb{S}_1 \times \mathbb{S}_2$ , such that we need to bound their conjunction within  $\varepsilon$ , i.e.  $\Pr[\lambda_1(x, g, s_1) \wedge \lambda_2(x, g, s_2)]$ , then  $m^{UC}(\varepsilon, \delta)$  is in  $\mathcal{O}(\max_{i \in \{1, 2\}} \{Sd(\Psi_i)\})$ .*

The proof of Corollary 2.5 is relegated to the appendix. The following claim (whose proof is also relegated to the appendix) is also useful

**Corollary 2.6** ( $Sd$  for Argmax). *Let  $\Psi_{\max}$  be defined by  $\mathbb{G} = \{f : \mathcal{X} \rightarrow \mathcal{Y}\}$  and  $\mathbb{S} = \mathcal{X}$ , where  $\mathcal{Y}$  is endowed with a total order  $\succ$ , and  $\lambda(x, g, x^*) = \mathbb{1}[g(x) \succ g(x^*)]$ . Then,  $Sd(\Psi_{\max}) = 1$ .*

To conclude, in order to establish an efficient PAC algorithm for a problem  $\Psi$ , it suffices to upper-bound  $m(\Psi)$  by its solution dimension  $Sd(\Psi)$ .

### 2.3 Uniform Convergence Beyond Consistency

As discussed in Section 1.3, uniform convergence complexity bounds can bound the sample complexity for Agnostic PAC learning via Empirical Risk Minimizers (ERM learners). However, agnostic solution learning can be defined in many ways. We discuss two definitions, for which the corresponding notion of *ERM solving* has a sample complexity that follows from uniform convergence.

**Definition 2.7.** For a given STATISTICALSOLUTION  $\Psi$ ,  $\mathcal{A}$  is a worst-case Agnostic PAC Solver if

$$\Pr_{T \sim \mathcal{D}^m} [L_{\mathcal{D}}(g, \mathcal{A}(T)) \leq \min_s \max_{g' \in \mathbb{G}|_T} L_{\mathcal{D}}(g', s) + \varepsilon] \geq 1 - \delta;$$

it is a Bayesian agnostic PAC Solver, for a prior over games  $\tilde{\mathcal{D}}$ , if

$$\Pr_{T \sim \mathcal{D}^m} [\mathbb{E}_{g' \sim \tilde{\mathcal{D}}} [L_{\mathcal{D}}(g', \mathcal{A}(T)) | g' \in \mathbb{G}|_T] \leq \min_s \mathbb{E}_{g' \sim \tilde{\mathcal{D}}} [L_{\mathcal{D}}(g', s) | g' \in \mathbb{G}|_T] + \varepsilon] \geq 1 - \delta.$$

**Corollary 2.8.** *Given some STATISTICALSOLUTION, the sample complexity for Worst-Case and Bayesian agnostic PAC solving is in  $\mathcal{O}(Sd(\Psi))$ , and is achievable by an empirical risk minimizer.*

The proof of Corollary 2.8 is relegated to the appendix.

## 3 Learning Game-Theoretic Distribution-Based Solution Concepts

Let us now apply our theory for learning solution concepts in game-theoretic domains; all problems described below follow a common theme: rather than learning preferences, we learn solutions using the sampled dataset. While the focus of this paper is on game-theoretic solutions, our theory applies for other types of solution concepts as long as one can define a local loss function  $\lambda$  that depends only on a given point in  $x \in \mathcal{X}$ ,  $g \in \mathbb{G}$  and the solution  $s \in \mathbb{S}$  (see Section 2).

### 3.1 The PAC Core for Hedonic Games

Let us begin with hedonic games (Section 1.4.1); we analyze another type of cooperative game (TU cooperative games) in the appendix. A partition  $\pi^*$  of  $N$  PAC stabilizes a hedonic game w.r.t. a distribution  $\mathcal{D} \in \Delta(2^N)$  (where  $i \in \pi^*(i)$  for every  $i \in N$ ), if  $\Pr_{S \sim \mathcal{D}} [\forall i \in S : v_i(S) \geq v_i(\pi^*(i))] < \varepsilon$ . The local loss function  $\lambda$  takes as input a coalition  $S \subseteq N$ , players' valuations  $\vec{v} = \langle v_1 \dots, v_n \rangle$  and a partition  $\pi \in \Pi(N)$ ;  $\lambda(S, \vec{v}, \pi) = 1$  iff  $S$  can block  $\pi$  under  $\vec{v}$ . Our key result here is that the sample complexity of PAC stabilizing hedonic games is linear in  $n$ , for any class  $\mathcal{H}$  of games.

**Lemma 3.1.** *For any class of Hedonic Games  $\mathcal{H}$  over  $n$  players, the solution dimension of PAC stabilizing  $\mathcal{H}$  is  $\leq n$ .*

*Proof.* By definition, for a given hedonic game  $h \in \mathcal{H}$ , a partition  $\pi$ , and a coalition  $S \subseteq N$ , the local loss  $\lambda(S, h, \pi) = 0$  if and only if there exists a player in  $S$  that does not prefer it over her assigned coalition in  $\pi$ , i.e.  $v_i(S) < v_i(\pi(i))$ . If a set of  $m$  coalitions  $\mathcal{S} = \{S_1, \dots, S_m\}$  is  $\mathcal{S}$ -shattered by a witness  $h \in \mathcal{H}$ , then for each  $S_j \in \mathcal{S}$ , there exists a coalition structure  $\pi_j$  such that  $\lambda(S_j, h, \pi_j) = 0$ , but  $\lambda(S_k, h, \pi_j) = 1$  for all  $k \neq j$ . In other words, under  $\pi_j$ , there exists some  $i \in S_j$  such that  $v_i(S_j) < v_i(\pi_j(i))$ , and for all  $k \neq j$  and for all  $i \in S_k$ ,  $v_i(S_k) \geq v_i(\pi_j(i))$ . We conclude that for every  $S_j \in \mathcal{S}$ , there exists a player  $i$  who strictly prefers all coalitions that she belongs to in  $\mathcal{S} \setminus \{S_j\}$  over  $S_j$ . More formally, we let  $\mathcal{T}(i)$  be the set of coalitions which are least preferred by player  $i$  in  $\mathcal{S}$ ; note that  $\mathcal{T}(i)$  must be a singleton, or else we arrive at a contradiction (the least liked coalition must be unique). Therefore, if  $\mathcal{S}$  is  $\mathcal{S}$ -shattered, the number of coalitions in  $\mathcal{S}$  is bounded by  $n$ , and we are done.  $\square$

Applying Theorem 2.4 and leveraging Lemma 3.1 we obtain the following result:

**Theorem 3.2.** *A class of Hedonic games  $\mathcal{H}$  is efficiently PAC stabilizable iff there exists an algorithm that outputs a partition consistent with samples evaluated by a game  $g \in \mathcal{H}$ ; the sample complexity in this case is  $\mathcal{O}(n)$ .*

In particular, Sliwinski and Zick [32] propose a consistent algorithm for *top-responsive* hedonic games [13, Chapter 15]; Igarashi et al. [24] present a consistent algorithm for hedonic games whose underlying interaction graph is a tree [23]. Indeed, given Theorem 3.2, it suffices to show that the algorithms they propose are consistent; their correctness is immediately implied by our results.

### 3.2 PAC Competitive Equilibria

Competitive equilibria (CE) readily admit a PAC variant: given an allocation  $\pi$ , let  $\mathcal{P}_i(\pi) = \{S \subseteq G : v_i(S) > v_i(\pi(i))\}$  be the set of bundles that are strongly preferred by  $i$  to  $\pi(i)$ . One can think of a CE as an outcome that ensures that  $\mathcal{P}_i(\pi) \cap \mathcal{A}_i(\vec{p}, \beta_i) = \emptyset$ , i.e.  $i$  cannot afford any bundle that it prefers to its assigned bundle. In the statistical variant, we wish to ensure that this intersection has a low measure under a distribution  $\mathcal{D}$  over  $2^G$ . We define a loss  $\lambda_i$  per player  $i$  as follows: given a bundle of goods  $S \subseteq G$ , player valuations  $\vec{v}$  and a market outcome  $\langle \pi, \vec{p} \rangle$ ,  $\lambda_i(S, \vec{v}, \langle \pi, \vec{p} \rangle) = 1$  iff  $S$  is both affordable (in  $\mathcal{A}_i(\vec{p}, \beta_i)$ ), and is preferred to  $\pi(i)$  (in  $\mathcal{P}_i(\pi)$ ). Our objective is to ensure that the overall error of player loss functions  $\lambda_1, \dots, \lambda_n$  are within an error of  $\varepsilon$ . Lemma 3.3 bounds the sample complexity for PAC learning this problem,  $m^{PAC}(\varepsilon, \delta)$ , by  $O(k)$ . Therefore, by Theorem 2.4 and Corollary 2.5, any algorithm that generates an outcome consistent against  $m$  sampled bundles would also be a PAC CE solver with a sample complexity in  $O(k)$ . We refer to an instance of the CE problem as  $\Psi_{CE}(N, G, \vec{v}, \vec{\beta})$ .

**Lemma 3.3.** *The solution dimension  $Sd(\Psi_{CE}(N, G, \vec{v}, \vec{\beta}))$  is  $O(k)$ , where  $k = |G|$ .*

*Proof.* For every player  $i \in N$ , the local constraint  $\lambda_i$  can be seen as a conjunction of  $\lambda_{1,i}(S, \vec{v}, \langle \pi^*, \vec{p}^* \rangle) = \mathbb{1} \left[ \sum_{g_j \in S} p_j^* > \beta_i \right]$ , and  $\lambda_{2,i}(S, \vec{v}, \langle \pi^*, \vec{p}^* \rangle) = \mathbb{1}[v_i(S) > v_i(\pi^*(i))]$ . Since  $\lambda_{1,i}$  is defined by a linear constraint set by  $\vec{p}^*$  and  $\beta_i$ , it can be S-shattered by  $O(k)$  samples (in a manner similar to linear separators in standard PAC learning), which bounds its S-dimension. On the other hand, every  $\lambda_{2,i}$  is a simple argmax constraint, which by Corollary 2.6, has a solution dimension of 1. By applying Corollary 2.5, the dimension of  $\lambda_i = \lambda_{1,i} \wedge \lambda_{2,i}$  is  $O(k)$ ; since the condition of  $\lambda_i$  must hold for each  $i \in N$ , the CE loss is given by  $\lambda = \bigwedge_i \lambda_i$ , which is  $O(k)$  by Corollary 2.5.  $\square$

Lemma 3.3 bounds the dimension of  $\Psi_{CE}$  by  $O(k)$ ; however, the challenge is to design algorithms that generate consistent market solutions: bundle assignments and prices that ensure that all observed goods have been allocated, with no excess demand or assignment. We show the existence of consistent solutions in two different settings; however, our solutions relax the market constraints. For Fisher markets with budgets  $\vec{\beta}$ , for any  $\zeta > 0$ , there exists a perturbed budget vector  $\vec{\beta}^*$  with  $\|\vec{\beta}^* - \vec{\beta}\|_\infty \leq \zeta$  for which there exists a consistent solution  $\langle \pi^*, \vec{p}^* \rangle$  w.r.t.  $\vec{\beta}^*$ ; this result holds for *any class of valuation functions*. Theorem 3.4 utilizes *inefficient* market outcomes, where a good may be allocated to more than one person; it is easy to think of an allocation  $\pi$  as a list of vectors in  $\{0, 1\}^k$ , where  $\pi_j(i) = 1$  iff the  $j$ -th good is allocated to player  $i$ . If all goods are allocated, then  $\sum_{i \in N} \pi(i) = \vec{1}$ ; if goods are over-allocated, then  $\sum_{i \in N} \pi(i) > \vec{1}$ .

**Theorem 3.4.** *We are given  $\Psi_{CE}(N, G, \vec{v}, \vec{\beta})$ , and  $m$  sampled bundles  $S_1, \dots, S_m \subseteq G$  evaluated by  $\vec{v}$ . For any  $\zeta > 0$ , there exists a perturbation on  $\vec{\beta}$ ,  $\vec{\beta}^*$  such that  $\|\vec{\beta} - \vec{\beta}^*\|_\infty < \zeta$ , for which there is an outcome  $\langle \pi^*, \vec{p}^* \rangle$  such that players with budget levels  $\vec{\beta}^*$  do not demand  $S_1, \dots, S_m$ ; moreover,  $\|\sum_{i \in N} \pi^*(i) - \vec{1}\|_2 \leq \frac{k}{2}$ , where  $\vec{1} = (1, 1, \dots, 1) \in [0, 1]^k$ .*

*Proof.* We restrict ourselves to finding an assignment using only the sampled bundles and the empty bundle, i.e. for all  $i \in N$ :  $\pi^*(i) \in \{\emptyset, S_1, \dots, S_m\}$ ; thus, we avoid making any assumptions about the structure of  $v_i$ . Budish [15, Theorem 1] shows that given  $\vec{\beta}$  such that  $\max_i \beta_i > \min_i \beta_i$ , for any  $\zeta > 0$  there exists a perturbed budget vector  $\vec{\beta}^*$  and an outcome  $\langle \pi^*, \vec{p}^* \rangle$  for which:  $\pi^*(i) \in \operatorname{argmax}_{S \in \mathcal{A}_i(\vec{p}^*, \beta_i)} v_i(S)$ ;  $\|\vec{\beta} - \vec{\beta}^*\|_\infty < \zeta$  and  $\|\sum_{i \in N} \pi^*(i) - \vec{1}\|_2 \leq \frac{k}{2}$ . Assuming that for every other  $S \notin \{\emptyset, S_1, \dots, S_m\}$ ,  $v_i(S) \leq 0$  for all  $i \in N$ , and applying the result by Budish, there exists a consistent outcome satisfying our requirements.  $\square$

While it makes no assumptions on player valuations, Theorem 3.4 is not constructive: it relies on a classic result from Budish [15], which utilizes a fixed-point theorem by Cromme and Diener [17] for discontinuous maps to bound excess demand. We analyze *exchange economies*, a market variant with divisible goods, in the appendix. In both cases, we are able to show that consistent market solutions exist. However, our results show the existence of solutions which only *partially satisfy* the equilibrium guarantees; moreover, both cases utilize non-constructive fixed-point theorems, rather than provide an efficient algorithm. There is little reason to believe that consistent solutions can be

easily computed in the general case; finding market solutions in settings similar to ours is PPAD complete [29].

### 3.3 PAC Condorcet Winners

We conclude with a discussion of statistical solution concepts in voting (see Section 1.4.3 above). A PAC Condorcet winner is a candidate  $c^*$  such that  $\Pr_{c \sim \mathcal{D}}[B(\succ, c, c^*)] < \varepsilon$  (recall that  $B(\succ, c, c^*) = 1$  iff a majority of voters prefer  $c$  to  $c^*$ ). We refer to the problem of finding a Condorcet winner as  $\Psi_{Cond}$ . We require that given a sample  $T \subseteq C$  of candidates, we can infer voters' preferences w.r.t.  $T$ . This can be encoded as a valuation function of  $i$  over the candidates (as is the case for hedonic games, see Section 3.1), or the truncated ranking  $\succ_i$  over the sampled candidates for every  $i \in N$ . Given a class of preference profiles  $\mathbb{H}$ , let  $\Psi_{Cond}(\mathbb{H})$  be the problem of finding Condorcet winners for profiles in  $\mathbb{H}$ . We define the *tournament graph*: this is a directed graph where candidates are nodes; given a preference profile  $\succ$ , there is an edge from  $a$  to  $b$  if  $a$  beats  $b$  in a pairwise election under  $\succ$ .

**Theorem 3.5.** *Given a class of preference profiles  $\mathbb{H}$  over  $C$  such that  $|C| > 1$ , and a sample of candidates  $T \subseteq C$ , the following are equivalent: (a) There exists a consistent solver for  $\Psi_{Cond}(\mathbb{H})$  that returns a PAC Condorcet winner  $c^* \in T$ . (b)  $Sd(\Psi_{Cond}) = 1$ . (c) for every preference profile  $\succ \in \mathbb{H}$ , the tournament graph is transitive.*

*In particular, if  $\mathbb{H}$  satisfies the above, there exists a PAC solver for  $\Psi_{Cond}(\mathbb{H})$  whose sample complexity is  $\frac{1}{\varepsilon} \log \frac{1}{\delta}$ .*

*Proof.* If there is a preference profile  $h \in \mathbb{H}$  for which the tournament graph contains a 3-cycle, then there immediately exist two vertices of that cycle that can be S-shattered. This is true since for every  $\mathbb{H}$  with more than one candidate, every singleton is shattered. Therefore,  $Sd(\Psi_{Cond}) = 1$  if and only if there are no preference profiles with Condorcet 3-cycles, which is equivalent to transitivity. Similarly, the existence of Condorcet winner for every  $C' \in C$  is equivalent to absence of any cycles, which is equivalent to transitivity of the tournament graph.  $\square$

Two notable families of voter preferences exhibit transitive preferences: single peaked preferences [13, Chapter 2] and single-crossing preferences [21] (see [20] for an overview); thus, if  $\mathbb{H}$  is any of the former, a Condorcet winner can be PAC learned using  $\frac{1}{\varepsilon} \log \frac{1}{\delta}$  samples. Whenever the Condorcet winner is known to exist within a sample  $C' \in C$ , the problem is equivalent to the argmax problem discussed in Corollary 2.6. However, as shown in Section 2.3, the graph dimension is still useful as a means to estimate (within  $\pm\varepsilon$  with high confidence) the behavior of a candidate in pairwise elections using a small empirical sample, even when no Condorcet winner exists. Theorem 3.6 bounds  $Sd(\Psi_{Cond})$  in the case where Condorcet winners do not exist. The result bounds the solution dimension in terms of the underlying structure of the tournament graph, and is based on Corollary 2.8; the full proof is in the appendix.

**Theorem 3.6.** *Let  $k$  be the largest number of candidates, such that for some tournament graph in  $\mathbb{H}$ , every pair among them is part of some 3-cycle. Then  $Sd(\Psi_{Cond}(\mathbb{H})) \leq \log_2(k + 2)$ .*

## 4 Conclusions and Future Work

We propose a formal, general framework for learning solution concepts from data, and apply it to several problems in economic domains. While several solution concepts have a polynomial sample complexity, efficiently computing a consistent solution remains a challenging open problem. In the case of market equilibria, we believe that there exist consistent algorithms for specific valuation classes, such as gross-substitutes [22] or submodular valuations. While we mostly focus on the realizable case, solving the non-realizable case is an interesting open problem. Our model easily accommodates *approximate solutions* (as we do for market equilibria) by assimilating the approximation guarantee into the loss function; this can be done generally by adopting the PMAC learning framework [3]. Our work upper-bounds the solution dimension using a generalization of the graph dimension; however, we offer no lower bounds. Daniely et al. [18] use the *Natarajan dimension* [27] to establish lower bounds in multiclass learning; using the Natarajan dimension to lower-bound the solution dimension is a promising direction for future work.

## References

- [1] M. Anthony and P. Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, 1999.
- [2] Y. Bachrach, E. Elkind, E. Malizia, R. Meir, D. Pasechnik, J. S. Rosenschein, J. Rothe, and M. Zuckerman. Bounds on the cost of stabilizing a cooperative game. *Journal of Artificial Intelligence Research*, 63:987–1023, 2018.
- [3] M. F. Balcan and N. Harvey. Learning submodular functions. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 793–802, 2011.
- [4] M. F. Balcan, A. D. Procaccia, and Y. Zick. Learning cooperative games. *CoRR*, abs/1505.00039, 2015.
- [5] M. F. Balcan, T. Dick, R. Noothigattu, and A. D. Procaccia. Envy-free classification. *CoRR*, abs/1809.08700, 2018.
- [6] M. F. Balcan, T. Sandholm, and E. Vitercik. A general theory of sample complexity for multi-item profit maximization. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, pages 173–174, 2018.
- [7] M.F. Balcan, A.D. Procaccia, and Y. Zick. Learning cooperative games. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 475–481, 2015.
- [8] M.F. Balcan, T. Sandholm, and E. Vitercik. Sample complexity of automated mechanism design. In *Proceedings of the 29th Annual Conference on Neural Information Processing Systems (NIPS)*, pages 2083–2091, 2016.
- [9] E. Balkanski and Y. Singer. The sample complexity of optimizing a convex function. In *Proceedings of the 30th Conference on Computational Learning Theory (COLT)*, pages 275–301, 2017.
- [10] E. Balkanski, A. Rubinstein, and Y. Singer. The power of optimization from samples. In *Proceedings of the 29th Annual Conference on Neural Information Processing Systems (NIPS)*, pages 4017–4025, 2016.
- [11] E. Balkanski, U. Syed, and S. Vassilvitskii. Statistical cost sharing. In *Proceedings of the 30th Annual Conference on Neural Information Processing Systems (NIPS)*, pages 6222–6231, 2017.
- [12] A. Bhattacharyya and P. Dey. Sample complexity for winner prediction in elections. In *Proceedings of the 14th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1421–1431, 2015.
- [13] F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A.D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- [14] G. Brero, B. Lubin, and S. Seuken. Combinatorial auctions via machine learning-based preference elicitation. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 128–136, 2018.
- [15] E. Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- [16] R. Cole and T. Roughgarden. The sample complexity of revenue maximization. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC)*, pages 243–252, 2014.
- [17] L. J. Cromme and I. Diener. Fixed point theorems for discontinuous mappings. *Mathematical Programming*, 51:257–267, 1991.
- [18] A. Daniely, S. Sabato, S. Ben-David, and S. Shalev-Shwartz. Multiclass learnability and the ERM principle. *Journal of Machine Learning Research*, 16:2377–2404, 2011.
- [19] N. Devanur, Z. Huang, and C. Psomas. The sample complexity of auctions with side information. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC)*, pages 426–439, 2016.

- [20] E. Elkind, M. Lackner, and D. Peters. Structured preferences. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 10. AI Access, 2017.
- [21] J. Gans and M. Smart. Majority voting with single-crossing preferences. *Journal of Public Economics*, 59:219–237, 1996.
- [22] F. Gul and E. Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic theory*, 87(1):95–124, 1999.
- [23] A. Igarashi and E. Elkind. Hedonic games with graph-restricted communication. In *Proceedings of the 15th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 242–250, 2016.
- [24] A. Igarashi, J. Sliwinski, and Y. Zick. Forming probably stable communities with limited interactions. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 2019. To Appear.
- [25] J. Morgenstern and T. Roughgarden. On the pseudo-dimension of nearly optimal auctions. In *Proceedings of the 28th Annual Conference on Neural Information Processing Systems (NIPS)*, pages 136–144, 2015.
- [26] J. Morgenstern and T. Roughgarden. Learning simple auctions. In *Proceedings of the 29th Conference on Computational Learning Theory (COLT)*, pages 1298–1318, 2016.
- [27] B. K. Natarajan. On learning sets and functions. *Machine Learning*, 4(1):67–97, 1989.
- [28] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [29] A. Othman, C. Papadimitriou, and A. Rubinstein. The complexity of fairness through equilibrium. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*, pages 209–226, 2014.
- [30] N. Rosenfeld, E. Balkanski, A. Globerson, and Y. Singer. Learning to optimize combinatorial functions. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, pages 4374–4383, 2018.
- [31] S. Shalev-Shwartz, O. Shamir, N. Srebro, and K. Sridharan. Learnability, stability and uniform convergence. *Journal of Machine Learning Research*, 11:2635–2670, 2010.
- [32] J. Sliwinski and Y. Zick. Learning hedonic games. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2730–2736, 2017.
- [33] V. Syrgkanis. A sample complexity measure with applications to learning optimal auctions. In *Proceedings of the 30th Annual Conference on Neural Information Processing Systems (NIPS)*, pages 5352–5359, 2017.
- [34] V.N. Vapnik and A. Ya. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16(2):264–280, 1971.
- [35] H. Zhang and V. Conitzer. A PAC framework for aggregating agents’ judgments. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 2019. To Appear.

# Appendix

## A Missing Proofs for Section 4.2

We now present the proof for the Theorem 2.4 that provides an upper bound to the sample complexity for uniform convergence in terms of the *solution dimension*.

**Definition 2.3** (Solution-based Dimension). Given some  $C \subseteq \mathcal{X}$ , we say the set  $C$  is S-shattered in  $\Psi$  if there exists a game  $g \in \mathbb{G}$ , such that for every binary labelling  $b : C \rightarrow \{0, 1\}$  there exists a solution  $s \in \mathbb{S}$  (that may depend on  $b$ ) such that for all  $x \in C$ ,  $\lambda(x, g, s) = b(x)$ .

The Solution-based dimension of  $\Psi$ , denoted  $Sd(\Psi)$ , is the size of the largest set S-shattered in  $\Psi$ , and  $(C, g)$  as the corresponding shattering witness.

**Theorem 2.4.** *There are universal constants  $\alpha_1$  and  $\alpha_2$ , such that if  $Sd(\Psi) = d$ , then for a sample of  $m \geq \alpha_1 \frac{d + \log(\frac{1}{\delta})}{\varepsilon^2}$  points  $T = \langle x_j, y_j \rangle_{j=1}^m$ ,*

$$\Pr_{T \sim \mathcal{D}^m} [\exists g \in \mathbb{G}|_T, s \in \mathbb{S} : |\hat{L}_T(g, s) - L_{\mathcal{D}}(g, s)| > \varepsilon] < \delta.$$

Furthermore, if a solution  $s^*$  is consistent, i.e.  $\hat{L}_T(g, s^*) = 0$ , then for any  $m$  greater than  $\frac{\alpha_2}{\varepsilon} (\log(\frac{1}{\varepsilon})d + \log(\frac{1}{\delta}))$ , we have that  $\Pr_{T \sim \mathcal{D}^m} [L_{\mathcal{D}}(g, s^*) > \varepsilon] < \delta$ .

*Proof.* For any  $g \in \mathbb{G}$ , consider  $\Phi_g = \{\phi_{g,s} : \mathcal{X} \rightarrow \{0, 1\} | s \in \mathbb{S}, \phi_{g,s}(x) = 1 - \lambda(x, g, s)\}$ . Observe that,  $Sd(\Psi) = \max_g VC(\Phi_g)$ ; this implies that if  $Sd(\Psi) = d$ , then for every  $g \in \mathbb{G}$ ,  $|\Phi_g| \leq |\mathcal{X}|^d$ .

From Theorem 1.2, we know that for any  $g \in \mathbb{G}|_T$  and  $f \in 2^{\mathcal{X}}$ , if  $T$  contains at least  $m \geq \alpha \frac{d + \log(\frac{1}{\delta})}{\varepsilon^2}$  samples (for the appropriate constant  $\alpha$ , then (where  $L_T^\phi$  and  $L_{\mathcal{D}}^\phi$  denote the corresponding loss functions for binary functions):

$$\Pr_{T \sim \mathcal{D}^m} [\exists \phi_{g,s} \in \Phi_g : |L_T^\phi(\phi_{g,s}, f) - L_{\mathcal{D}}^\phi(\phi_{g,s}, f)| > \varepsilon] < \delta \quad (1)$$

Let  $\mathbb{1} : \mathcal{X} \rightarrow \{0, 1\}$ , such that  $\forall x : \mathbb{1}(x) = 1$ . Then, it follows that  $\lambda(x, g, s) = \mathbb{I}[\phi_{g,s}(x) \neq \mathbb{1}(x)]$ . By substituting the loss functions in Eq (1), and taking  $f = \mathbb{1}$ , we get:

$$\Pr_{T \sim \mathcal{D}^m} [\exists g \in \mathbb{G}|_T, s \in \mathbb{S} : |\hat{L}_T(g, s) - L_{\mathcal{D}}(g, s)| > \varepsilon] < \delta.$$

This proves that  $m^{UC}(\varepsilon, \delta)$ , and  $m^{PAC}(\varepsilon, \delta)$  are both polynomially bounded by  $Sd(\Psi)$ . If a solution is consistent, then we also know that there exists some  $s^* \in \mathbb{S}$  such that for all  $g \in \mathbb{G}|_T$ ,  $\phi_{g,s^*} = \mathbb{1} \in \Phi_g$ . The second part of the statement similarly follows from the upper bound of  $m^{PAC}$  for binary functions in Theorem 1.2.  $\square$

Let us next prove Corollary 2.5.

**Corollary 2.5.** *Given a STATISTICALSOLUTION problem  $\Psi = \langle \mathcal{X}, \mathcal{Y}, \mathbb{G}, \mathbb{S} \rangle$ :*

**Simultaneous Constraints:** *if multiple local loss functions  $\lambda_1, \dots, \lambda_k$  need to be simultaneously approximated within  $\varepsilon$ , i.e.  $\forall i \in [k] : |\hat{L}_i(g, s) - L_{i\mathcal{D}}(g, s)| < \varepsilon$ , then the sample complexity of finding a solution satisfying all of them is in  $\mathcal{O}(\max_{i \in [k]} \{m_i^{UC}(\varepsilon, \delta)\})$ .*

**Separable Conjunctions:** *if there are local constraints  $\lambda_1$  over  $\mathbb{S}_1$ , and  $\lambda_2$  over  $\mathbb{S}_2$ , where  $\mathbb{S}_\Psi = \mathbb{S}_1 \times \mathbb{S}_2$ , such that we need to bound their conjunction within  $\varepsilon$ , i.e.  $\Pr[\lambda_1(x, g, s_1) \wedge \lambda_2(x, g, s_2)]$ , then  $m^{UC}(\varepsilon, \delta)$  is in  $\mathcal{O}(\max_{i \in \{1,2\}} \{Sd(\Psi_i)\})$ .*

*Proof.* Part 1 (Simultaneous Constraints) is a direct corollary of Theorem 2.4; if  $m \geq \max_i \{m_i^{PAC}(\varepsilon, \delta)\}$ , then

$$\forall i : \Pr_{T \sim \mathcal{D}^m} [\Pr_{x \sim \mathcal{D}} [\lambda_i(x, g, s)] < \varepsilon] < \delta.$$

For part 2 (Separable Conjunctions), we are given a loss function  $\lambda : \mathcal{X} \times \mathbb{G} \times (\mathbb{S}_1 \times \mathbb{S}_2) \rightarrow \{0, 1\}$  with  $\lambda(x, g, (s_1, s_2)) = \lambda_1(x, g, s_1) \wedge \lambda_2(x, g, s_2)$ , such that for any  $g \in \mathbb{G}$ , we wish to find  $(s_1, s_2) \in \mathbb{S}_1 \times \mathbb{S}_2$  that bound the gap between the empirical and statistical loss.

For any  $s_2 \in \mathbb{S}_2$ , let us define  $\lambda_{1|s_2} : \mathcal{X} \times \mathbb{G} \times \mathbb{S}_1 \rightarrow \{0, 1\}$  with  $\lambda_{1|s_2}(x, g, s_1) = \lambda(x, g, (s_1, s_2)) = \lambda_1(x, g, s_1) \wedge \lambda_2(x, g, s_2)$ . Observe that  $\lambda_{1|s_2}(x, g, s_1)$  equals  $\lambda_1(x, g, s_1)$  whenever  $\lambda_2(x, g, s_2) = 1$ , and is otherwise 0 for every  $s_1$ . We note that a set shattered under  $\lambda_{1|s_2}$  cannot contain any point  $x'$  such that  $\lambda_2(x', g, s_2) = 0$  (points for which  $\lambda_2(x', g, s_2) = 0$  always evaluate to 0 under  $\lambda$  and do not admit Boolean functions  $b$  for which  $\lambda(x', g, (s_1, s_2)) = b(x') = 1$ ); thus, we know that a set shattered in  $\lambda_{1|s_2}$  is also shattered under  $\lambda_1$ , therefore the solution dimension corresponding to  $\lambda_{1|s_2}$  is bounded by the solution dimension corresponding to  $\lambda_1$ . We also observe that the formulas for empirical and statistical loss under  $\lambda_{1|s_2}$  and  $\lambda$  are equivalent; therefore, by Theorem 2.4, we know that if  $m \geq m_1^{UC}(\varepsilon, \delta)$

$$\forall g \in \mathbb{G}, s_1 \in \mathbb{S}_1, s_2 \in \mathbb{S}_2 : Pr[|\hat{L}_{1|s_2, T}(g, s) - L_{1|s_2, \mathcal{D}}(g, s)| < \varepsilon] > 1 - \delta \quad (2)$$

Therefore,  $m^{UC} \in O(\max\{m_i^{UC}\})$ .  $\square$

We note that the bound for conjuncts in Corollary 2.5 trivially generalizes to any number of conjunctions, i.e. if  $\lambda = \lambda_1 \wedge \dots \wedge \lambda_q$ , the dimension of  $\lambda$  is upper-bounded by the dimension of the domains corresponding to  $\lambda_1, \dots, \lambda_q$ .

**Corollary 2.6** (*Sd for Argmax*). *Let  $\Psi_{\max}$  be defined by  $\mathbb{G} = \{f : \mathcal{X} \rightarrow \mathcal{Y}\}$  and  $\mathbb{S} = \mathcal{X}$ , where  $\mathcal{Y}$  is endowed with a total order  $\succ$ , and  $\lambda(x, g, x^*) = \mathbb{1}[g(x) \succ g(x^*)]$ . Then,  $Sd(\Psi_{\max}) = 1$ .*

*Proof.* Let us assume that a set  $C = \{x_1, x_2\}$  is  $\mathbb{S}$ -shattered with  $g \in \mathbb{G}$ . This implies the existence of:  
i)  $x'$  such that  $\lambda(x_1, g, x') = 0 \implies g(x') \succeq g(x_1)$  and  $\lambda(x_2, g, x') = 1 \implies g(x_2) \succ g(x')$ . And,  
ii)  $x''$  such that  $\lambda(x_1, g, x'') = 1 \implies g(x_1) \succ g(x'')$  and  $\lambda(x_2, g, x'') = 0 \implies g(x'') \succeq g(x_2)$ .  
However, since  $\succeq$  is transitive, this leads to contradiction.  $\square$

Finally, we present the full proof for ERM solvers for non-realizable agnostic solution learning.

**Corollary 2.8.** *Given some STATISTICALSOLUTION, the sample complexity for Worst-Case and Bayesian agnostic PAC solving is in  $\mathcal{O}(Sd(\Psi))$ , and is achievable by an empirical risk minimizer.*

*Proof.* We first prove the result for worst-case agnostic learning. Let  $\mathcal{A}_m : (\mathcal{X} \times \mathcal{Y})^m \rightarrow \mathbb{S}$  be an ERM Solver that for any sample of  $m \geq m^{UC}(\varepsilon, \delta)$  points  $T = \langle (x_i, y_i) \rangle_{i=1}^m$ , outputs a solution  $\mathcal{A}_m(T) \in \mathbb{S}$  that minimizes  $\max_{g \in \mathbb{G}|_T} \hat{L}_T(g, \mathcal{A}_m(T))$ .

Let  $s^* \in \mathbb{S}$  be a solution that minimizes the worst-case statistical loss for any game  $g$  consistent with the sample, i.e.

$$s^* \in \operatorname{argmin}_{s \in \mathbb{S}} \max_{g \in \mathbb{G}|_T} L_{\mathcal{D}}(g, s).$$

By definition of  $\mathcal{A}_m(T)$ , we know that

$$\max_{g \in \mathbb{G}|_T} \hat{L}_T(g, \mathcal{A}_m(T)) \leq \max_{g \in \mathbb{G}|_T} \hat{L}_T(g, s^*).$$

For a sample of  $m \geq m^{UC}(\varepsilon/2, \delta/2)$  points drawn i.i.d. from  $\mathcal{D}$ , we know by Theorem 2.4, that for any  $g_0 \in \mathbb{G}$ , with probability  $\geq 1 - \delta/2$  we have:

$$|L_{\mathcal{D}}(g_0, \mathcal{A}_m(T)) - \hat{L}_T(g_0, \mathcal{A}_m(T))| < \frac{\varepsilon}{2},$$

and, with probability  $\geq 1 - \delta/2$ , for any  $g' \in \mathbb{G}$ ,

$$|\hat{L}_T(g_0, s^*) - L_{\mathcal{D}}(g', s^*)| < \frac{\varepsilon}{2}.$$

Putting it all together, we get that with probability  $\geq \delta$ ,

$$\begin{aligned} L_{\mathcal{D}}(g_0, \mathcal{A}_m(T)) &\leq \hat{L}_T(g_0, \mathcal{A}_m(T)) + \frac{\varepsilon}{2} \leq \max_{g \in \mathbb{G}|_T} \hat{L}_T(g, \mathcal{A}_m(T)) + \frac{\varepsilon}{2} \\ &\leq \max_{g \in \mathbb{G}|_T} \hat{L}_T(g, s^*) + \frac{\varepsilon}{2} \leq \max_{g \in \mathbb{G}|_T} L_{\mathcal{D}}(g, s^*) + \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \max_{g \in \mathbb{G}|_T} L_{\mathcal{D}}(g, s^*) + \varepsilon. \end{aligned}$$

Therefore, if  $m \geq m^{UC}(\varepsilon/2, \delta/2)$ , then

$$\Pr_{T \sim \mathcal{D}^m} [L_{\mathcal{D}}(g, \mathcal{A}_m(T)) \leq \min_{s \in \mathbb{S}} \max_{g' \in \mathbb{G}|_T} L_{\mathcal{D}}(g', s) + \varepsilon] \geq 1 - \delta.$$

For Bayesian agnostic learning, we present the result for distributions with a finite support over  $\mathbb{G}$ ; the case where the distribution has an infinite support over  $\mathbb{G}$  is similar. Let  $\tilde{\mathcal{D}}$  be some prior distribution with a finite support over  $\mathbb{G}$ . Then, by definition

$$\mathbb{E}_{g' \sim \tilde{\mathcal{D}}} [L_{\mathcal{D}}(g', \mathcal{A}(T)) | g' \in \mathbb{G}|_T] = \sum_{g \in \mathbb{G}|_T} L_{\mathcal{D}}(g, s) \Pr_{\tilde{\mathcal{D}}}(g).$$

When  $m \geq m^{UC}(\varepsilon, \delta)$ , by Theorem 2.4, for every  $g \in \mathbb{G}|_T$  and  $s \in \mathbb{S}$ ,  $\Pr[|L_{\mathcal{D}}(g, s) - \hat{L}_T(g, s)| \leq \varepsilon] > 1 - \delta$ . Combining these expressions, with probability mass  $\tilde{\mathcal{D}}$ , we get

$$\Pr[|\mathbb{E}_{g' \sim \tilde{\mathcal{D}}} [L_{\mathcal{D}}(g', s) | g' \in \mathbb{G}|_T] - \mathbb{E}_{g' \sim \tilde{\mathcal{D}}} [\hat{L}_T(g', s) | g' \in \mathbb{G}|_T]| \leq \varepsilon] > 1 - \delta.$$

Therefore, an ERM solver that minimizes  $\mathbb{E}_{g' \sim \tilde{\mathcal{D}}} [\hat{L}_T(g', s) | g' \in \mathbb{G}|_T]$ , also bounds the statistical loss within  $\varepsilon$ .  $\square$

## B The PAC Core in TU Cooperative Games

### B.1 Cooperative Games

In *transferable utility* (TU) cooperative games players' preferences are included by a function  $v : 2^N \rightarrow \mathbb{R}_+$  mapping every subset  $S \subseteq N$  to a value  $v(S) \in \mathbb{R}_+$ . We are interested in finding “good” payoff divisions for the game. These are simply vectors  $\vec{x} = (x_1, \dots, x_n) \in \mathbb{R}_+^n$  such that  $\sum_{i=1}^n x_i = v(N)$  (efficiency) and  $x_i \geq v(\{i\})$  for all  $i \in N$  (individual rationality). We say that a coalition  $S \subseteq N$  *blocks* a payoff division  $\vec{x}$  if  $\sum_{i \in S} x_i < v(S)$ ; that is, the coalition  $S$  can guarantee its members a strictly higher reward should they choose to break off from working with everyone else. The *core* is the (possibly empty) set of payoff divisions from which no coalition can deviate; in other words,  $\text{core}(N, v) = \{\vec{x} \in \mathbb{R}_+^n \mid \forall S \subseteq N : \sum_{i \in S} x_i \geq v(S); \sum_{i=1}^n x_i = v(N)\}$ .

### B.2 The PAC Core for TU Cooperative Games

Balcan et al. [7] propose a learning-based approach to finding a PAC stable payoff division for TU cooperative games (see definitions in Section B.1). Given a distribution  $\mathcal{D}$  over  $2^N$ , a payoff division  $\vec{x}^*$   $\varepsilon$ -PAC stabilizes the game  $\langle N, v \rangle$  with respect to  $\mathcal{D}$  if

$$\Pr_{S \sim \mathcal{D}} \left[ v(S) < \sum_{i \in S} x_i^* \right] < \varepsilon.$$

In what follows, we provide a proof for the PAC stabilizability of TU cooperative games in the language of Theorem 2.4; direct proofs of this fact appear in [4, 11].

**Theorem B.1.** *The solution dimension of TU cooperative games is  $\mathcal{O}(n)$ .*

*Proof.* We show that any set of  $> n$  coalitions cannot be S-shattered as per Definition 2.3. Taking a set of coalitions  $\mathcal{S} = \{S_1, \dots, S_m\}$ , it is S-shattered if there is some TU cooperative game  $v : \mathcal{S} \rightarrow \mathbb{R}_+$  such that for all  $\mathcal{T} \subseteq \mathcal{S}$ , there exists some vector  $\vec{x}^*$  in  $\mathbb{R}^n$  such that for all  $T \in \mathcal{T}$ ,  $v(T) \geq x^*(T)$ , and for all  $S \in \mathcal{S} \setminus \mathcal{T}$ ,  $v(S) < x^*(S)$ . Let us bound the dimension  $m$  of  $\mathcal{S}$ . The problem is equivalent to shattering sets of vectors in the hypercube  $\{0, 1\}^n$  with linear classifiers, which is well-known to be impossible for sets of size  $> n$  [1]. We conclude that  $Sd$  for the PAC core of TU cooperative games is  $\leq n$ .  $\square$

We note that the solution computed in Balcan et al. [4] is only efficient (i.e. with  $\sum_{i=1}^n x_i = v(N)$ ) if the core of the cooperative game  $v$  is not empty. In the case where the game  $v$  has an empty core, the solution computed still satisfies the core constraints with high probability with respect to  $\mathcal{D}$ , but may not be efficient. However, the payoff outputted is using the *minimal subsidy required* in order to stabilize the game. In other words, the total payoff is no more than the *cost of stability* of the underlying game  $v$  [2]. Efficiency is an important requirement: without it, one can “cheat” and pay each player some arbitrarily high amount, guaranteeing that the underlying game is stable.

## C PAC Competitive Equilibria in Exchange Economies

In Section 1.4.2 we define Fisher markets; these are markets where goods are indivisible, and each player  $i \in N$  has a budget  $\beta_i$ . In what follows, we consider *exchange economies* [28, Chapters 6 and 9], which follow a somewhat different structure.

### C.1 Exchange Economies

In exchange economies we have a set  $G = \{g_1, \dots, g_k\}$  of  $k$  *divisible* goods, and player valuations are of the form  $v_i : [0, 1]^k \rightarrow \mathbb{R}_+$  for every  $i \in N$ ; bundle assignments are  $\pi : N \rightarrow [0, 1]^k$  (assigning a quantity  $q_j \leq 1$  of good  $g_j$  to player  $i$  can be thought of as player  $i$  receiving  $q_j$  percent of good  $g_j$ ).

In exchange economies with divisible goods, we assume that each player has an initial *endowment* of goods  $\vec{e}_i \in [0, 1]^k$ , denoting the (divisible) amount of each good that she possesses. It is no loss of generality to assume that  $\sum_{i=1}^n \vec{e}_{i,j} = 1$  for every good  $g_j$ ; in other words, the quantity  $e_{i,j}$  is the relative amount of good  $g_j$  that player  $i$  possesses. Given item prices, players *demand* certain item bundles. The *affordable set* is the set of all divisible goods whose total price is less than the worth of player  $i$ 's endowment under  $\vec{p}$ .

$$\mathcal{A}_i(\vec{p}) = \left\{ \vec{g} \in [0, 1]^k : \sum_{j=1}^k p_j g_j \leq \sum_{j=1}^k p_j e_{i,j} \right\}.$$

An outcome  $\langle \pi, \vec{p} \rangle$  is a competitive equilibrium if  $\pi(i) \in \mathcal{A}_i(\vec{p})$ , and  $\forall \vec{g} \in \mathcal{A}(\vec{p}), v_i(\pi(i)) \geq v_i(\vec{g})$ .

### C.2 PAC Market Equilibria in Exchange Economies

We assume that player preferences are convex. We show that for any sample of fractional bundles  $T = \{\vec{b}_1, \dots, \vec{b}_m\}$ , there exists a solution  $\langle \pi^*, \vec{p}^* \rangle$  consistent with  $T$  with non-positive excess assignment (but potentially leaving some goods unassigned). We assume that none of the goods are undesirable, i.e. for every good there exists at least one player that assigns a positive value to some quantity of that good.

**Theorem C.1.** *Suppose we are given an exchange economy for divisible goods with convex preferences and without undesirable goods. We observe  $m$  sampled bundles  $T = \{\vec{b}_1, \dots, \vec{b}_m\}$  and player valuations over the bundles, along with player endowments  $\vec{e}_1, \dots, \vec{e}_n$ . There exists a solution  $\langle \pi^*, \vec{p}^* \rangle$  such that every player  $i$  is assigned a bundle they can afford given their endowment, which is consistent (against any possible valuation functions that could have generated the observed values).*

*Proof.* Without loss of generality, let us work with the reduced space of only observed goods. Let  $\mathcal{U}$  denote the underlying space of convex preferences from which we draw player preferences over assignments. Let  $\mathcal{U}_{i,T}$  denote the space of all valuation functions  $u$  that satisfy the observed values, i.e.  $u(\vec{b}_j) = v_i(\vec{b}_j)$  for every  $j \in [m]$ . Since prices only need to satisfy the affordability criterion for every player, i.e.  $\vec{p}^* \cdot \pi(i) \leq \vec{p}^* \cdot \vec{e}_i$ , we can normalize and assume that prices belong to the simplex  $\Delta_{n-1}$ . Also, observe that the absence of undesirable goods implies that in any consistent solution the price of any observed good cannot be 0.

Now let us define the demand set function as  $D : \mathcal{U} \times \Delta_{n-1} \times [0, 1]^n \rightarrow 2^{[0, 1]^n}$ , such that

$$D(u, \vec{p}, \vec{e}_i) = \left\{ \vec{b} \in [0, 1]^n : u(\vec{b}) \geq u(\vec{b}_j) \forall \vec{b}_j \in T; \text{ and } \vec{p} \cdot \vec{b} \leq \vec{p} \cdot \vec{e}_i \right\}.$$

We observe that under convex preferences (i.e. quasi-concave utility functions), for every  $u, \vec{p}$  and  $\vec{e}$ ,  $D(u, \vec{p}, \vec{e})$  is a convex and compact body; in addition,  $D$  is continuous in  $\vec{p}$ . Define, for every player  $i \in N$ ,  $D_i^T(\vec{p}) = \bigcap_{u \in \mathcal{U}_{i,T}} D(u, \vec{p}, \vec{e}_i)$ :  $D_i^T(\vec{p})$  is the set of all possible bundles that player  $i$  might demand under the price vector  $\vec{p}$ , under all possible utility functions that agree with the sample  $T$ . The intersection  $D_i^T$  is convex and compact, as well as continuous in  $\vec{p}$ . Also observe that  $D_i^T(\vec{p})$  is always non-empty, since there is at least one bundle among the observed samples and the empty bundle which belongs to each of the  $D(u, \vec{p}, \vec{e}_i)$ .

Let  $f : [0, 1]^{k \times n} \rightarrow [0, n]^k$  be the *excess demand function*:  $f(\pi) = \sum_{i \in N} \pi(i) - \vec{1}$  (where  $\vec{1} = (1, 1, \dots, 1) \in [0, 1]^n$ ). Let  $\mathbf{z} : \Delta_{n-1} \rightarrow 2^{[0, k]^n}$ , be the function  $\mathbf{z}(\vec{p}) = \{f(\pi) : \pi \in \prod_{i \in K} D_i^T(\vec{p})\}$ . The function  $f$  is linear, therefore  $\mathbf{z}(\vec{p})$  is convex and compact, and  $\mathbf{z}$  is continuous. Using  $\mathbf{z}$ , we define a function  $\mathbf{g} : \Delta_{n-1} \rightarrow 2^{\Delta_{n-1}}$  such that  $r$ -th component is given by

$$\mathbf{g}(\vec{p}) = \left\{ \vec{g} : \text{where } g_r(\vec{p}) = \frac{p_r + \max\{0, z_r\}}{1 + \sum_{s=1}^n \max\{0, z_s\}}, \text{ for some } \vec{z} \in \mathbf{z}(\vec{p}) \right\}$$

By applying Kakutani's fixed-point theorem over  $\mathbf{g}$ , we get the existence of some  $\vec{p}^*$  such that  $\vec{p}^* \in \mathbf{g}(\vec{p}^*)$ . This implies the existence of some  $\pi \in \prod_{i \in K} D_i^T(\vec{p}^*)$ , such that  $f(\pi) \in \mathbf{z}(\vec{p}^*)$  satisfies

$$p_r^* = \frac{p_r^* + \max(0, f_r(\vec{p}^*))}{1 + \sum_{s=1}^n \max(0, f_s(\vec{p}^*))}$$

Let  $r^*$  some non-positive component of  $f(\pi)$ , as argued above; then  $p_{r^*}^* = p_{r^*}^*(1 + \sum_{s=1}^n \max(0, f_s(\vec{p}^*)))$ ; this implies that for all  $r$ :  $\max(0, f_r(\vec{p}^*)) = 0$ . Therefore, there exists some allocation  $\pi^*$  with non-positive excess demand at  $\vec{p}^*$ , such that  $\langle \pi^*, \vec{p}^* \rangle$  is consistent with the observed bundles against all possible  $u \in \mathcal{U}|_{i,T}$  for every player  $i$ .  $\square$

## D PAC Condorcet Winners

Recall that a Condorcet winner is a candidate  $c^*$  that beats every other candidate  $c$  in a pairwise election, i.e. for every other candidate  $c$ , a majority of voters prefer  $c^*$  to  $c$ . We note that if there are only two candidates, then a Condorcet winner trivially exists (barring the case when the votes are tied); however, when there are three or more candidates, a Condorcet winner is not guaranteed to exist. This case is known in the literature as Condorcet cycles (or Condorcet paradoxes). For the sake of completeness, we provide a simple example of a voting profile where no candidate is a Condorcet winner.

**Example D.1.** Consider a setting with three candidates  $a, b, c$  and three voters, 1, 2 and 3 whose preferences over  $a, b, c$  are as follows:

$$\begin{aligned} 1 : a \succ_1 b \succ_1 c \\ 2 : b \succ_2 c \succ_2 a \\ 3 : c \succ_3 a \succ_3 b \end{aligned}$$

1 and 3 prefer  $a$  to  $b$ ; 1 and 2 prefer  $b$  to  $c$ ; 2 and 3 prefer  $c$  to  $a$ . Therefore, there are no Condorcet winners.

Next, let us provide a complete proof of Theorem 3.6.

**Theorem 3.6.** *Let  $k$  be the largest number of candidates, such that for some tournament graph in  $\mathbb{H}$ , every pair among them is part of some 3-cycle. Then  $Sd(\Psi_{Cond}(\mathbb{H})) \leq \log_2(k+2)$ .*

*Proof.* Let us assume that there is a set of size  $d$ ,  $C_0 \subseteq C$ , that is shattered. Then by definition of shattering, there exists  $2^d$  different candidates corresponding to every subset of  $C_0$ , such that for  $f \in 2^{C_0}$  there exists a candidate  $c_f \in C$ , such that  $c_1 \in C_0$  beats  $c_f$  in the tournament graph if and only if  $f(c_1) = 1$ .

Let us focus on  $2^d - 2$  of these candidates, corresponding to all non-trivial functions  $f \in 2^{C_0}$  (let us, for now, ignore the functions that assign a constant value (of 0 or 1) to all candidates in  $C_0$ ). Then, for every pair of these functions  $f_1$  and  $f_2$ , there exists a candidate  $c_1 \in C_0$  such that  $f_1(c_1) = 1$  and  $f_2(c_1) = 0$ , and a candidate  $c_2$  such that  $f_1(c_2) = 0$  and  $f_2(c_2) = 1$ . This implies the existence of a directed path of length at most 2 from  $c_{f_1}$  to  $c_{f_2}$ , and vice versa. Since, in a tournament graph, either the edge  $c_{f_1} \rightarrow c_{f_2}$  or  $c_{f_2} \rightarrow c_{f_1}$  exists, we know that  $c_{f_1}$  and  $c_{f_2}$  are members of some 3-cycle. Since this is true for all such  $c_f$ 's, we know that  $2^d - 2$  is less than or equal to largest number of candidates, such that for some tournament graph in  $\mathbb{H}$ , every pair amongst them is part of some 3-cycle.  $\square$