Estimation of Priors with Applications to Preference Elicitation

Liu Yang
School of Computer Science
Carnegie Mellon University
liuy@cs.cmu.edu

Steve Hanneke
steve.hanneke@gmail.com

August 21, 2013

Abstract

We extend the work of [YHC13] on estimating prior distributions over VC classes to the case of real-valued functions in a VC subgraph class. We then apply this technique to the problem of maximizing customer satisfaction using a minimal number of value queries in an online preference elicitation scenario.

1 Introduction

Consider an online travel agency, where customers go to the site with some idea of what type of travel they are interested in; the site then poses a series of questions to each customer, and identifies a travel package that best suits their desires, budget, and dates. There are many options of travel packages, with options on location, site-seeing tours, hotel and room quality, etc. Because of this, serving the needs of an *arbitrary* customer might be a lengthy process, requiring many detailed questions. Fortunately, the stream of customers is typically not a worst-case sequence, and in particular obeys many statistical regularities: in particular, it is not too far from reality to think of the customers as being independent and identically distributed samples. With this assumption in mind, it becomes desirable to identify some of these statistical regularities so that we can pose the questions that are typically most relevant, and thereby more quickly identify the travel package that best suits the needs of the typical customer. One straightforward way to do this is to directly *estimate* the distribution of customer value functions, and optimize the questioning system to minimize the expected number of questions needed to find a suitable travel package.

One can model this problem in the style of Bayesian combinatorial auctions, in which each customer has a value function for each possible bundle of items. However, it is slightly different, in that we do not assume the distribution of customers is known, but rather are interested in estimating this distribution; the obtained estimate can then be used in combination with methods based on Bayesian decision theory. In contrast to the literature on Bayesian auctions (and subjectivist Bayesian decision theory in general), this technique is able to maintain general guarantees on performance that hold under an objective interpretation of the problem, rather than merely guarantees holding under an arbitrary assumed prior belief. This general idea is sometimes referred to as *Empirical Bayesian* decision theory in the machine learning and statistics literatures. The ideal result for an Empirical Bayesian algorithm is to be competitive with the corresponding Bayesian methods based on the *actual* distribution of the data (assuming the data are random, with an unknown distribution); that is, although the Empirical Bayesian methods only operate with a data-based estimate of the distribution, the aim is to perform nearly as well as methods based on the true (unobservable) distribution. In this work, we present results of this type, in the context of an abstraction of the aforementioned online travel agency problem, where the measure of performance is the expected number of questions to find a suitable package.

The technique we use here is rooted in the work of [YHC13] on *transfer learning* with a VC class. The component of that work of interest here is the estimation of prior distributions over VC classes. Essentially, there is a given class of functions, from which a sequence of functions is sampled i.i.d. according to an unknown distribution. We observe a number of values of each of these functions, evaluated at points chosen at random, and are then tasked with estimating the distribution of these functions. This is more challenging than the traditional problem of nonparametric density estimation, since we are not permitted direct access to these functions, but rather only a limited number of evaluations of the function (i.e., a number of $(x, f(x))$ pairs). The work of [YHC13] develops a technique for estimating the distribution of these functions, given that the functions are binary-valued, the class of functions has finite VC dimension, and the class of distributions is totally bounded. In this work, we extend this technique to classes of real-valued functions having finite pseudo-dimension, a natural generalization of VC dimension for real-valued functions [Hau92].

The specific application we are interested in here may be expressed abstractly as a kind of combinatorial auction with preference elicitation. Specifically, we suppose there is a collection of items on a menu, and each possible bundle of items has an associated fixed price. There is a stream of customers, each with a valuation function that provides a value for each possible bundle of items. The objective is to serve each customer a bundle of items that nearly-maximizes his or her surplus value (value minus price). However, we are not permitted direct observation of the customer valuation functions; rather, we may query for the value of any given bundle of items; this is referred to as a *value query* in the literature on preference elicitation in combinatorial auctions (see Chapter 14 of [CSS06], [ZBS03]). The objective is to achieve this near-maximal surplus guarantee, while making only a small number of queries per customer. We suppose the customer valuation function are sampled i.i.d. according to an unknown distribution over a known (but arbitrary) class of real-valued functions having finite pseudo-dimension. Reasoning that knowledge of this distribution should allow one to make a smaller number of value queries per customer, we are interested in estimating this unknown distribution, so that as we serve more and more customers, the number of queries per customer required to identify a near-optimal bundle should decrease. In this context, we in fact prove that in the limit, the expected number of queries per customer converges to the number required of a method having direct knowledge of the true distribution of valuation functions.

2 Notation

Let $\mathcal{B}$ denote a σ-algebra on $\mathcal{X} \times \mathbb{R}$, let $\mathcal{B}_{\mathcal{X}}$ denote the σ-algebra on $\mathcal{X}$. Also let $\rho(h, g) = \int |h - g| \mathrm{d}P_X$, where P_X is a marginal distribution over $\mathcal{X}$. Let $\mathcal{F}$ be a class of functions $\mathcal{X} \to \mathbb{R}$ with Borel σ-algebra $\mathcal{B}_{\mathcal{F}}$ induced by ρ. Let Θ be a set of parameters, and for each $\theta \in \Theta$, let π_θ denote a probability measure on $(\mathcal{F}, \mathcal{B}_{\mathcal{F}})$. We suppose $\{\pi_\theta : \theta \in \Theta\}$ is totally bounded in total variation distance, and that $\mathcal{F}$ is a uniformly bounded VC subgraph class with pseudodimension d. We also suppose ρ is a *metric* when restricted to $\mathcal{F}$.

Let $\{X_{ti}\}_{t,i \in \mathbb{N}}$ be i.i.d. P_X random variables. For each $\theta \in \Theta$, let $\{h^*_{t\theta}\}_{t \in \mathbb{N}}$ be i.i.d. π_θ random variables, independent from $\{X_{ti}\}_{t,i \in \mathbb{N}}$. For each $t \in \mathbb{N}$ and $\theta \in \Theta$, let $Y_{ti}(\theta) = h^*_{t\theta}(X_{ti})$ for $i \in \mathbb{N}$, and let $\mathcal{Z}_t(\theta) = \{(X_{t1}, Y_{t1}(\theta)), (X_{t2}, Y_{t2}(\theta)), \ldots\}$, $\mathbb{X}_t = \{X_{t1}, X_{t2}, \ldots\}$, and $\mathbb{Y}_t(\theta) = \{Y_{t1}(\theta), Y_{t2}(\theta), \ldots\}$; for each $k \in \mathbb{N}$, define $\mathcal{Z}_{tk}(\theta) = \{(X_{t1}, Y_{t1}(\theta)), \ldots, (X_{tk}, Y_{tk}(\theta))\}$, $\mathbb{X}_{tk} = \{X_{t1}, \ldots, X_{tk}\}$, and $\mathbb{Y}_{tk}(\theta) = \{Y_{t1}(\theta), \ldots, Y_{tk}(\theta)\}$.

For any probability measures μ, μ', we denote the total variation distance by

$$\|\mu - \mu'\| = \sup_A \mu(A) - \mu'(A),$$

where A ranges over measurable sets.

Lemma 1. *For any $\theta, \theta' \in \Theta$ and $t \in \mathbb{N}$,*

$$\|\pi_\theta - \pi_{\theta'}\| = \|\mathbb{P}_{\mathcal{Z}_t(\theta)} - \mathbb{P}_{\mathcal{Z}_t(\theta')}\|.$$

Proof. Fix $\theta, \theta' \in \Theta$, $t \in \mathbb{N}$. Let $\mathbb{X} = \{X_{t1}, X_{t2}, \ldots\}$, $\mathbb{Y}(\theta) = \{Y_{t1}(\theta), Y_{t2}(\theta), \ldots\}$, and for $k \in \mathbb{N}$ let $\mathbb{X}_k = \{X_{t1}, \ldots, X_{tk}\}$. and $\mathbb{Y}_k(\theta) = \{Y_{t1}(\theta), \ldots, Y_{tk}(\theta)\}$. For $h \in \mathcal{F}$, let $c_{\mathbb{X}}(h) = \{(X_{t1}, h(X_{t1})), (X_{t2}, h(X_{t2})), \ldots\}$.

For $h, g \in \mathcal{F}$, define $\rho_{\mathbb{X}}(h, g) = \lim\limits_{m \to \infty} \frac{1}{m} \sum_{i=1}^m |h(X_{ti}) - g(X_{ti})|$ (if the limit exists), and $\rho_{\mathbb{X}_k}(h, g) = \frac{1}{k} \sum_{i=1}^k |h(X_{ti}) - g(X_{ti})|$. Note that since $\mathcal{F}$ is a uniformly bounded VC subgraph class, so is the collection of functions $\{|h - g| : h, g \in \mathcal{F}\}$, so that the uniform strong law of large numbers implies that with probability one, $\forall h, g \in \mathcal{F}$, $\rho_{\mathbb{X}}(h, g)$ exists and has $\rho_{\mathbb{X}}(h, g) = \rho(h, g)$ [Vap82].

Consider any $\theta, \theta' \in \Theta$, and any $A \in \mathcal{B}_{\mathcal{F}}$. Then any $h \notin A$ has $\forall g \in A$, $\rho(h, g) > 0$ (by the metric assumption). Thus, if $\rho_{\mathbb{X}}(h, g) = \rho(h, g)$ for all $h, g \in \mathcal{F}$, then $\forall h \notin A$,

$$\forall g \in A, \rho_{\mathbb{X}}(h, g) = \rho(h, g) > 0 \implies \forall g \in A, c_{\mathbb{X}}(h) \neq c_{\mathbb{X}}(g) \implies c_{\mathbb{X}}(h) \notin c_{\mathbb{X}}(A).$$

This implies $c_{\mathbb{X}}^{-1}(c_{\mathbb{X}}(A)) = A$. Under these conditions,

$$\mathbb{P}_{\mathcal{Z}_t(\theta)|\mathbb{X}}(c_{\mathbb{X}}(A)) = \pi_\theta(c_{\mathbb{X}}^{-1}(c_{\mathbb{X}}(A))) = \pi_\theta(A),$$

and similarly for θ'.

Any measurable set C for the range of $\mathcal{Z}_t(\theta)$ can be expressed as $C = \{c_{\bar{x}}(h) : (h, \bar{x}) \in C'\}$ for some appropriate $C' \in \mathcal{B}_{\mathcal{F}} \otimes \mathcal{B}_{\mathcal{X}}^\infty$. Letting $C'_{\bar{x}} = \{h : (h, \bar{x}) \in C'\}$, we have

$$\mathbb{P}_{\mathcal{Z}_t(\theta)}(C) = \int \pi_\theta(c_{\bar{x}}^{-1}(c_{\bar{x}}(C'_{\bar{x}}))) \mathbb{P}_{\mathbb{X}}(\mathrm{d}\bar{x}) = \int \pi_\theta(C'_{\bar{x}}) \mathbb{P}_{\mathbb{X}}(\mathrm{d}\bar{x}) = \mathbb{P}_{(h^*_{t\theta}, \mathbb{X})}(C').$$

Likewise, this reasoning holds for θ'. Then

$$\begin{aligned}\|\mathbb{P}_{\mathcal{Z}_t(\theta)} - \mathbb{P}_{\mathcal{Z}_t(\theta')}\| &= \|\mathbb{P}_{(h^*_{t\theta},\mathbb{X})} - \mathbb{P}_{(h^*_{t\theta'},\mathbb{X})}\| \\ &= \sup_{C' \in \mathcal{B}_{\mathcal{F}} \otimes \mathcal{B}^\infty_{\mathcal{X}}} \left| \int (\pi_\theta(C'_{\bar{x}}) - \pi_{\theta'}(C'_{\bar{x}}))\mathbb{P}_{\mathbb{X}}(\mathrm{d}\bar{x}) \right| \\ &\leq \int \sup_{A \in \mathcal{B}_{\mathcal{F}}} |\pi_\theta(A) - \pi_{\theta'}(A)| \mathbb{P}_{\mathbb{X}}(\mathrm{d}\bar{x}) = \|\pi_\theta - \pi_{\theta'}\|.\end{aligned}$$

Since $h^*_{t\theta}$ and $\mathbb{X}$ are independent, for $A \in \mathcal{B}_{\mathcal{F}}$, $\pi_\theta(A) = \mathbb{P}_{h^*_{t\theta}}(A) = \mathbb{P}_{h^*_{t\theta}}(A)\mathbb{P}_{\mathbb{X}}(\mathcal{X}^\infty) = \mathbb{P}_{(h^*_{t\theta},\mathbb{X})}(A \times \mathcal{X}^\infty)$. Analogous reasoning holds for $h^*_{t\theta'}$. Thus, we have

$$\begin{aligned}\|\pi_\theta - \pi_{\theta'}\| &= \|\mathbb{P}_{(h^*_{t\theta},\mathbb{X})}(\cdot \times \mathcal{X}^\infty) - \mathbb{P}_{(h^*_{t\theta'},\mathbb{X})}(\cdot \times \mathcal{X}^\infty)\| \\ &\leq \|\mathbb{P}_{(h^*_{t\theta},\mathbb{X})} - \mathbb{P}_{(h^*_{t\theta'},\mathbb{X})}\| = \|\mathbb{P}_{\mathcal{Z}_t(\theta)} - \mathbb{P}_{\mathcal{Z}_t(\theta')}\|.\end{aligned}$$

Combining the above, we have $\|\mathbb{P}_{\mathcal{Z}_t(\theta)} - \mathbb{P}_{\mathcal{Z}_t(\theta')}\| = \|\pi_\theta - \pi_{\theta'}\|$. □

Lemma 2. *There exists a sequence $r_k = o(1)$ such that, $\forall t, k \in \mathbb{N}$, $\forall \theta, \theta' \in \Theta$,*

$$\|\mathbb{P}_{\mathcal{Z}_{tk}(\theta)} - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}\| \leq \|\pi_\theta - \pi_{\theta'}\| \leq \|\mathbb{P}_{\mathcal{Z}_{tk}(\theta)} - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}\| + r_k.$$

Proof. This proof follows identically to a proof of [YHC13], but is included here for completeness. Since $\mathbb{P}_{\mathcal{Z}_{tk}(\theta)}(A) = \mathbb{P}_{\mathcal{Z}_t(\theta)}(A \times (\mathcal{X} \times \mathbb{R})^\infty)$ for all measurable $A \subseteq (\mathcal{X} \times \mathbb{R})^k$, and similarly for θ', we have

$$\begin{aligned}\|\mathbb{P}_{\mathcal{Z}_{tk}(\theta)} - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}\| &= \sup_{A \in \mathcal{B}^k} \mathbb{P}_{\mathcal{Z}_{tk}(\theta)}(A) - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}(A) \\ &= \sup_{A \in \mathcal{B}^k} \mathbb{P}_{\mathcal{Z}_t(\theta)}(A \times (\mathcal{X} \times \mathbb{R})^\infty) - \mathbb{P}_{\mathcal{Z}_t(\theta')}(A \times (\mathcal{X} \times \mathbb{R})^\infty) \\ &\leq \sup_{A \in \mathcal{B}^\infty} \mathbb{P}_{\mathcal{Z}_t(\theta)}(A) - \mathbb{P}_{\mathcal{Z}_t(\theta')}(A) = \|\mathbb{P}_{\mathcal{Z}_t(\theta)} - \mathbb{P}_{\mathcal{Z}_t(\theta')}\|,\end{aligned}$$

which implies the left inequality when combined with Lemma 1.

Next, we focus on the right inequality. Fix $\theta, \theta' \in \Theta$ and $\gamma > 0$, and let $B \in \mathcal{B}^\infty$ be such that

$$\|\pi_\theta - \pi_{\theta'}\| = \|\mathbb{P}_{\mathcal{Z}_t(\theta)} - \mathbb{P}_{\mathcal{Z}_t(\theta')}\| < \mathbb{P}_{\mathcal{Z}_t(\theta)}(B) - \mathbb{P}_{\mathcal{Z}_t(\theta')}(B) + \gamma.$$

Let $\mathcal{A} = \{A \times (\mathcal{X} \times \mathbb{R})^\infty : A \in \mathcal{B}^k, k \in \mathbb{N}\}$. Note that $\mathcal{A}$ is an algebra that generates $\mathcal{B}^\infty$. Thus, Carathéodory's extension theorem [Sch95] implies that there exist disjoint sets $\{A_i\}_{i \in \mathbb{N}}$ in $\mathcal{A}$ such that $B \subseteq \bigcup_{i \in \mathbb{N}} A_i$ and

$$\mathbb{P}_{\mathcal{Z}_t(\theta)}(B) - \mathbb{P}_{\mathcal{Z}_t(\theta')}(B) < \sum_{i \in \mathbb{N}} \mathbb{P}_{\mathcal{Z}_t(\theta)}(A_i) - \sum_{i \in \mathbb{N}} \mathbb{P}_{\mathcal{Z}_t(\theta')}(A_i) + \gamma.$$

Since these A_i sets are disjoint, each of these sums is bounded by a probability value, which implies that there exists some $n \in \mathbb{N}$ such that

$$\sum_{i \in \mathbb{N}} \mathbb{P}_{\mathcal{Z}_t(\theta)}(A_i) < \gamma + \sum_{i=1}^{n} \mathbb{P}_{\mathcal{Z}_t(\theta)}(A_i),$$

which implies

$$\sum_{i\in\mathbb{N}} \mathbb{P}_{\mathcal{Z}_t(\theta)}(A_i) - \sum_{i\in\mathbb{N}} \mathbb{P}_{\mathcal{Z}_t(\theta')}(A_i) < \gamma + \sum_{i=1}^{n} \mathbb{P}_{\mathcal{Z}_t(\theta)}(A_i) - \sum_{i=1}^{n} \mathbb{P}_{\mathcal{Z}_t(\theta')}(A_i)$$
$$= \gamma + \mathbb{P}_{\mathcal{Z}_t(\theta)}\left(\bigcup_{i=1}^{n} A_i\right) - \mathbb{P}_{\mathcal{Z}_t(\theta')}\left(\bigcup_{i=1}^{n} A_i\right).$$

As $\bigcup_{i=1}^{n} A_i \in \mathcal{A}$, there exists $m \in \mathbb{N}$ and measurable $B_m \in \mathcal{B}^m$ such that $\bigcup_{i=1}^{n} A_i = B_m \times (\mathcal{X} \times \mathbb{R})^{\infty}$, and therefore

$$\mathbb{P}_{\mathcal{Z}_t(\theta)}\left(\bigcup_{i=1}^{n} A_i\right) - \mathbb{P}_{\mathcal{Z}_t(\theta')}\left(\bigcup_{i=1}^{n} A_i\right) = \mathbb{P}_{\mathcal{Z}_{tm}(\theta)}(B_m) - \mathbb{P}_{\mathcal{Z}_{tm}(\theta')}(B_m)$$
$$\leq \|\mathbb{P}_{\mathcal{Z}_{tm}(\theta)} - \mathbb{P}_{\mathcal{Z}_{tm}(\theta')}\| \leq \lim_{k\to\infty} \|\mathbb{P}_{\mathcal{Z}_{tk}(\theta)} - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}\|.$$

Combining the above, we have $\|\pi_\theta - \pi_{\theta'}\| \leq \lim_{k\to\infty} \|\mathbb{P}_{\mathcal{Z}_{tk}(\theta)} - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}\| + 3\gamma$. By letting γ approach 0, we have

$$\|\pi_\theta - \pi_{\theta'}\| \leq \lim_{k\to\infty} \|\mathbb{P}_{\mathcal{Z}_{tk}(\theta)} - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}\|.$$

So there exists a sequence $r_k(\theta, \theta') = o(1)$ such that

$$\forall k \in \mathbb{N}, \|\pi_\theta - \pi_{\theta'}\| \leq \|\mathbb{P}_{\mathcal{Z}_{tk}(\theta)} - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}\| + r_k(\theta, \theta').$$

Now let $\gamma > 0$ and let Θ_γ be a minimal γ-cover of Θ. Define the quantity $r_k(\gamma) = \max_{\theta,\theta'\in\Theta_\gamma} r_k(\theta, \theta')$. Then for any $\theta, \theta' \in \Theta$, let $\theta_\gamma = \operatorname{argmin}_{\theta''\in\Theta_\gamma} \|\pi_\theta - \pi_{\theta''}\|$ and $\theta'_\gamma = \operatorname{argmin}_{\theta''\in\Theta_\gamma} \|\pi_{\theta'} - \pi_{\theta''}\|$. Then a triangle inequality implies that $\forall k \in \mathbb{N}$,

$$\begin{aligned}\|\pi_\theta - \pi_{\theta'}\| &\leq \|\pi_\theta - \pi_{\theta_\gamma}\| + \|\pi_{\theta_\gamma} - \pi_{\theta'_\gamma}\| + \|\pi_{\theta'_\gamma} - \pi_{\theta'}\| \\ &< 2\gamma + r_k(\theta_\gamma, \theta'_\gamma) + \|\mathbb{P}_{\mathcal{Z}_{tk}(\theta_\gamma)} - \mathbb{P}_{\mathcal{Z}_{tk}(\theta'_\gamma)}\| \\ &\leq 2\gamma + r_k(\gamma) + \|\mathbb{P}_{\mathcal{Z}_{tk}(\theta_\gamma)} - \mathbb{P}_{\mathcal{Z}_{tk}(\theta'_\gamma)}\|.\end{aligned}$$

Triangle inequalities and the left inequality from the lemma statement (already established) imply

$$\begin{aligned}&\|\mathbb{P}_{\mathcal{Z}_{tk}(\theta_\gamma)} - \mathbb{P}_{\mathcal{Z}_{tk}(\theta'_\gamma)}\| \\ &\leq \|\mathbb{P}_{\mathcal{Z}_{tk}(\theta_\gamma)} - \mathbb{P}_{\mathcal{Z}_{tk}(\theta)}\| + \|\mathbb{P}_{\mathcal{Z}_{tk}(\theta)} - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}\| + \|\mathbb{P}_{\mathcal{Z}_{tk}(\theta'_\gamma)} - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}\| \\ &\leq \|\pi_{\theta_\gamma} - \pi_\theta\| + \|\mathbb{P}_{\mathcal{Z}_{tk}(\theta)} - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}\| + \|\pi_{\theta'_\gamma} - \pi_{\theta'}\| \\ &< 2\gamma + \|\mathbb{P}_{\mathcal{Z}_{tk}(\theta)} - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}\|.\end{aligned}$$

So in total we have

$$\|\pi_\theta - \pi_{\theta'}\| \leq 4\gamma + r_k(\gamma) + \|\mathbb{P}_{\mathcal{Z}_{tk}(\theta)} - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}\|.$$

Since this holds for all $\gamma > 0$, defining $r_k = \inf_{\gamma>0}(4\gamma + r_k(\gamma))$, we have the right inequality of the lemma statement. Furthermore, since each $r_k(\theta, \theta') = o(1)$, and $|\Theta_\gamma| < \infty$, we have $r_k(\gamma) = o(1)$ for each $\gamma > 0$, and thus we also have $r_k = o(1)$. $\square$

Lemma 3. $\forall t, k \in \mathbb{N}$, *there exists a monotone function* $M_k(x) = o(1)$ *such that,* $\forall \theta, \theta' \in \Theta$,

$$\|\mathbb{P}_{\mathcal{Z}_{tk}(\theta)} - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}\| \leq M_k\left(\|\mathbb{P}_{\mathcal{Z}_{td}(\theta)} - \mathbb{P}_{\mathcal{Z}_{td}(\theta')}\|\right).$$

Proof. Fix any $t \in \mathbb{N}$, and let $\mathbb{X} = \{X_{t1}, X_{t2}, \ldots\}$ and $\mathbb{Y}(\theta) = \{Y_{t1}(\theta), Y_{t2}(\theta), \ldots\}$, and for $k \in \mathbb{N}$ let $\mathbb{X}_k = \{X_{t1}, \ldots, X_{tk}\}$ and $\mathbb{Y}_k(\theta) = \{Y_{t1}(\theta), \ldots, Y_{tk}(\theta)\}$.

If $k \leq d$, then $\mathbb{P}_{\mathcal{Z}_{tk}(\theta)}(\cdot) = \mathbb{P}_{\mathcal{Z}_{td}(\theta)}(\cdot \times (\mathcal{X} \times \{-1, +1\})^{d-k})$, so that

$$\|\mathbb{P}_{\mathcal{Z}_{tk}(\theta)} - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}\| \leq \|\mathbb{P}_{\mathcal{Z}_{td}(\theta)} - \mathbb{P}_{\mathcal{Z}_{td}(\theta')}\|,$$

and therefore the result trivially holds.

Now suppose $k > d$. Fix any $\gamma > 0$, and let $B_{\theta,\theta'} \subseteq (\mathcal{X} \times \mathbb{R})^k$ be a measurable set such that

$$\begin{aligned}\mathbb{P}_{\mathcal{Z}_{tk}(\theta)}(B_{\theta,\theta'}) - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}(B_{\theta,\theta'}) &\leq \|\mathbb{P}_{\mathcal{Z}_{tk}(\theta)} - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}\| \\ &\leq \mathbb{P}_{\mathcal{Z}_{tk}(\theta)}(B_{\theta,\theta'}) - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}(B_{\theta,\theta'}) + \gamma.\end{aligned}$$

By Carathéodory's extension theorem, there exists a disjoint sequence of sets $\{B_i\}_{i=1}^{\infty}$ such that

$$\mathbb{P}_{\mathcal{Z}_{tk}(\theta)}(B_{\theta,\theta'}) - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}(B_{\theta,\theta'}) < \gamma + \sum_{i=1}^{\infty} \mathbb{P}_{\mathcal{Z}_{tk}(\theta)}(B_i) - \sum_{i=1}^{\infty} \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}(B_i),$$

and such that each $B_i(\theta, \theta')$ is representable as follows; for some $\ell_i(\theta, \theta') \in \mathbb{N}$, and sets $C_{ij} = (A_{ij1} \times (-\infty, t_{ij1}]) \times \cdots \times (A_{ijk} \times (-\infty, t_{ijk}])$, for $j \leq \ell_i(\theta, \theta')$, where each $A_{ijp} \in \mathcal{B}_{\mathcal{X}}$, the set $B_i(\theta, \theta')$ is representable as $\bigcup_{s \in S_i} \bigcap_{j=1}^{\ell_i(\theta,\theta')} D_{ijs}$, where $S_i \subseteq \{0, \ldots, 2^{\ell_i(\theta,\theta')} - 1\}$, each $D_{ijs} \in \{C_{ij}, C_{ij}^c\}$, and $s \neq s' \Rightarrow \bigcap_{j=1}^{\ell_i(\theta,\theta')} D_{ijs} \cap \bigcap_{j=1}^{\ell_i(\theta,\theta')} D_{ijs'} = \emptyset$. Since the $B_i(\theta, \theta')$ are disjoint, the above sums are bounded, so that there exists $m_k(\theta, \theta', \gamma) \in \mathbb{N}$ such that every $m \geq m_k(\theta, \theta', \gamma)$ has

$$\begin{aligned}&\mathbb{P}_{\mathcal{Z}_{tk}(\theta)}(B_{\theta,\theta'}) - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}(B_{\theta,\theta'}) \\ &\qquad < 2\gamma + \sum_{i=1}^{m} \mathbb{P}_{\mathcal{Z}_{tk}(\theta)}(B_i(\theta, \theta')) - \sum_{i=1}^{m} \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}(B_i(\theta, \theta')),\end{aligned}$$

Now define $\tilde{M}_k(\gamma) = \max_{\theta,\theta' \in \Theta_\gamma} m_k(\theta, \theta', \gamma)$. Then for any $\theta, \theta' \in \Theta$, let $\theta_\gamma, \theta'_\gamma \in \Theta_\gamma$ be such that $\|\pi_\theta - \pi_{\theta_\gamma}\| < \gamma$ and $\|\pi_{\theta'} - \pi_{\theta'_\gamma}\| < \gamma$, which implies $\|\mathbb{P}_{\mathcal{Z}_{tk}(\theta)} - \mathbb{P}_{\mathcal{Z}_{tk}(\theta_\gamma)}\| < \gamma$ and $\|\mathbb{P}_{\mathcal{Z}_{tk}(\theta')} - \mathbb{P}_{\mathcal{Z}_{tk}(\theta'_\gamma)}\| < \gamma$ by Lemma 2. Then

$$\begin{aligned}\|\mathbb{P}_{\mathcal{Z}_{tk}(\theta)} - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}\| &< \|\mathbb{P}_{\mathcal{Z}_{tk}(\theta_\gamma)} - \mathbb{P}_{\mathcal{Z}_{tk}(\theta'_\gamma)}\| + 2\gamma \\ &\leq \mathbb{P}_{\mathcal{Z}_{tk}(\theta_\gamma)}(B_{\theta_\gamma,\theta'_\gamma}) - \mathbb{P}_{\mathcal{Z}_{tk}(\theta'_\gamma)}(B_{\theta_\gamma,\theta'_\gamma}) + 3\gamma \\ &\leq \sum_{i=1}^{\tilde{M}_k(\gamma)} \mathbb{P}_{\mathcal{Z}_{tk}(\theta_\gamma)}(B_i(\theta_\gamma, \theta'_\gamma)) - \mathbb{P}_{\mathcal{Z}_{tk}(\theta'_\gamma)}(B_i(\theta_\gamma, \theta'_\gamma)) + 5\gamma.\end{aligned}$$

Again, since the $B_i(\theta_\gamma, \theta'_\gamma)$ are disjoint, this equals

$$
\begin{aligned}
&5\gamma + \mathbb{P}_{\mathcal{Z}_{tk}(\theta_\gamma)}\left(\bigcup_{i=1}^{\tilde{M}_k(\gamma)} B_i(\theta_\gamma, \theta'_\gamma)\right) - \mathbb{P}_{\mathcal{Z}_{tk}(\theta'_\gamma)}\left(\bigcup_{i=1}^{\tilde{M}_k(\gamma)} B_i(\theta_\gamma, \theta'_\gamma)\right) \\
&\leq 7\gamma + \mathbb{P}_{\mathcal{Z}_{tk}(\theta)}\left(\bigcup_{i=1}^{\tilde{M}_k(\gamma)} B_i(\theta_\gamma, \theta'_\gamma)\right) - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}\left(\bigcup_{i=1}^{\tilde{M}_k(\gamma)} B_i(\theta_\gamma, \theta'_\gamma)\right) \\
&= 7\gamma + \sum_{i=1}^{\tilde{M}_k(\gamma)} \mathbb{P}_{\mathcal{Z}_{tk}(\theta)}(B_i(\theta_\gamma, \theta'_\gamma)) - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}(B_i(\theta_\gamma, \theta'_\gamma)) \\
&\leq 7\gamma + \tilde{M}_k(\gamma) \max_{i \leq \tilde{M}_k(\gamma)} \left|\mathbb{P}_{\mathcal{Z}_{tk}(\theta)}(B_i(\theta_\gamma, \theta'_\gamma)) - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}(B_i(\theta_\gamma, \theta'_\gamma))\right|.
\end{aligned}
$$

Thus, if we can show that each $\left|\mathbb{P}_{\mathcal{Z}_{tk}(\theta)}(B_i(\theta_\gamma, \theta'_\gamma)) - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}(B_i(\theta_\gamma, \theta'_\gamma))\right|$ is bounded by a $o(1)$ function of $\|\mathbb{P}_{\mathcal{Z}_{td}(\theta)} - \mathbb{P}_{\mathcal{Z}_{td}(\theta')}\|$, then the result will follow by substituting this relaxation into the above expression and defining M_k by minimizing the resulting expression over $\gamma > 0$.

Toward this end, let C_{ij} be as above from the definition of $B_i(\theta_\gamma, \theta'_\gamma)$, and note that $I_{B_i(\theta_\gamma, \theta'_\gamma)}$ is representable as a function of the $I_{C_{ij}}$ indicators, so that

$$
\begin{aligned}
&\left|\mathbb{P}_{\mathcal{Z}_{tk}(\theta)}(B_i(\theta_\gamma, \theta'_\gamma)) - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}(B_i(\theta_\gamma, \theta'_\gamma))\right| \\
&= \|\mathbb{P}_{I_{B_i(\theta_\gamma, \theta'_\gamma)}(\mathcal{Z}_{tk}(\theta))} - \mathbb{P}_{I_{B_i(\theta_\gamma, \theta'_\gamma)}(\mathcal{Z}_{tk}(\theta'))}\| \\
&\leq \|\mathbb{P}_{(I_{C_{i1}}(\mathcal{Z}_{tk}(\theta)), \ldots, I_{C_{i\ell_i(\theta_\gamma, \theta'_\gamma)}}(\mathcal{Z}_{tk}(\theta)))} - \mathbb{P}_{(I_{C_{i1}}(\mathcal{Z}_{tk}(\theta')), \ldots, I_{C_{i\ell_i(\theta_\gamma, \theta'_\gamma)}}(\mathcal{Z}_{tk}(\theta')))}\| \\
&\leq 2^{\ell_i(\theta_\gamma, \theta'_\gamma)} \max_{J \subseteq \{1, \ldots, \ell_i(\theta_\gamma, \theta'_\gamma)\}} \mathbb{E}\left[\left(\prod_{j \in J} I_{C_{ij}}(\mathcal{Z}_{tk}(\theta))\right) \prod_{j \notin J}\left(1 - I_{C_{ij}}(\mathcal{Z}_{tk}(\theta))\right)\right. \\
&\left. - \left(\prod_{j \in J} I_{C_{ij}}(\mathcal{Z}_{tk}(\theta'))\right) \prod_{j \notin J}\left(1 - I_{C_{ij}}(\mathcal{Z}_{tk}(\theta'))\right)\right] \\
&\leq 2^{\ell_i(\theta_\gamma, \theta'_\gamma)} \sum_{J \subseteq \{1, \ldots, 2^{\ell_i(\theta_\gamma, \theta'_\gamma)}\}} \left|\mathbb{E}\left[\prod_{j \in J} I_{C_{ij}}(\mathcal{Z}_{tk}(\theta)) - \prod_{j \in J} I_{C_{ij}}(\mathcal{Z}_{tk}(\theta'))\right]\right| \\
&\leq 4^{\ell_i(\theta_\gamma, \theta'_\gamma)} \max_{J \subseteq \{1, \ldots, 2^{\ell_i(\theta_\gamma, \theta'_\gamma)}\}} \left|\mathbb{E}\left[\prod_{j \in J} I_{C_{ij}}(\mathcal{Z}_{tk}(\theta)) - \prod_{j \in J} I_{C_{ij}}(\mathcal{Z}_{tk}(\theta'))\right]\right| \\
&= 4^{\ell_i(\theta_\gamma, \theta'_\gamma)} \max_{J \subseteq \{1, \ldots, 2^{\ell_i(\theta_\gamma, \theta'_\gamma)}\}} \left|\mathbb{P}_{\mathcal{Z}_{tk}(\theta)}\left(\bigcap_{j \in J} C_{ij}\right) - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}\left(\bigcap_{j \in J} C_{ij}\right)\right|.
\end{aligned}
$$

Note that $\bigcap_{j \in J} C_{ij}$ can be expressed as some $(A_1 \times (-\infty, t_1]) \times \cdots \times (A_k \times (-\infty, t_k])$, where each $A_p \in \mathcal{B}_\mathcal{X}$ and $t_p \in \mathbb{R}$, so that, letting $\hat{\ell} = \max_{\theta, \theta' \in \Theta_\gamma} \max_{i \leq \tilde{M}_k(\gamma)} \ell_i(\theta, \theta')$ and $\mathcal{C}_k = \{(A_1 \times (-\infty, t_1]) \times \cdots \times (A_k \times (-\infty, t_k]) : \forall j \leq k, A_j \in \mathcal{B}_\mathcal{X}, t_k \in \mathbb{R}\}$, this last expression is at most

$$4^{\hat{\ell}} \sup_{C \in \mathcal{C}_k} \left| \mathbb{P}_{\mathcal{Z}_{tk}(\theta)}(C) - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}(C) \right|.$$

Next note that for any $C = (A_1 \times (-\infty, t_1]) \times \cdots \times (A_k \times (-\infty, t_k]) \in \mathcal{C}_k$, letting $C_1 = A_1 \times \cdots \times A_k$ and $C_2 = (-\infty, t_1] \times \cdots \times (-\infty, t_k]$,

$$\begin{aligned} \mathbb{P}_{\mathcal{Z}_{tk}(\theta)}(C) - \mathbb{P}_{\mathcal{Z}_{tk}(\theta')}(C) &= \mathbb{E}\left[\left(\mathbb{P}_{\mathbb{Y}_{tk}(\theta)|\mathbb{X}_{tk}}(C_2) - \mathbb{P}_{\mathbb{Y}_{tk}(\theta')|\mathbb{X}_{tk}}(C_2) \right) I_{C_1}(\mathbb{X}_{tk}) \right] \\ &\leq \mathbb{E}\left[\left| \mathbb{P}_{\mathbb{Y}_{tk}(\theta)|\mathbb{X}_{tk}}(C_2) - \mathbb{P}_{\mathbb{Y}_{tk}(\theta')|\mathbb{X}_{tk}}(C_2) \right| \right]. \end{aligned}$$

For $p \in \{1, \ldots, k\}$, let $C_{2p} = (-\infty, t_p]$. Then note that, by definition of d, for any given $x = (x_1, \ldots, x_k)$, the class $\mathcal{H}_x = \{x_p \mapsto I_{C_{2p}}(h(x_p)) : h \in \mathcal{F}\}$ is a VC class over $\{x_1, \ldots, x_k\}$ with VC dimension at most d. Furthermore, we have

$$\begin{aligned} &\left| \mathbb{P}_{\mathbb{Y}_{tk}(\theta)|\mathbb{X}_{tk}}(C_2) - \mathbb{P}_{\mathbb{Y}_{tk}(\theta')|\mathbb{X}_{tk}}(C_2) \right| \\ &\quad = \Big| \mathbb{P}_{(I_{C_{21}}(h^*_{t\theta}(X_{t1})), \ldots, I_{C_{2k}}(h^*_{t\theta}(X_{tk})))|\mathbb{X}_{tk}}(\{(1, \ldots, 1)\}) \\ &\qquad - \mathbb{P}_{(I_{C_{21}}(h^*_{t\theta'}(X_{t1})), \ldots, I_{C_{2k}}(h^*_{t\theta'}(X_{tk})))|\mathbb{X}_{tk}}(\{(1, \ldots, 1)\}) \Big|. \end{aligned}$$

Therefore, the results of [YHC13] (in the proof of their Lemma 3) imply that

$$\begin{aligned} &\left| \mathbb{P}_{\mathbb{Y}_{tk}(\theta)|\mathbb{X}_{tk}}(C_2) - \mathbb{P}_{\mathbb{Y}_{tk}(\theta')|\mathbb{X}_{tk}}(C_2) \right| \\ &\leq 2^k \max_{y \in \{0,1\}^d} \max_{D \in \{1, \ldots, k\}^d} \Big| \mathbb{P}_{\{I_{C_{2j}}(h^*_{t\theta}(X_{tj}))\}_{j \in D} | \{X_{tj}\}_{j \in D}}(\{y\}) \\ &\qquad - \mathbb{P}_{\{I_{C_{2j}}(h^*_{t\theta'}(X_{tj}))\}_{j \in D} | \{X_{tj}\}_{j \in D}}(\{y\}) \Big|. \end{aligned}$$

Thus, we have

$$\begin{aligned} &\mathbb{E}\left[\left| \mathbb{P}_{\mathbb{Y}_{tk}(\theta)|\mathbb{X}_{tk}}(C_2) - \mathbb{P}_{\mathbb{Y}_{tk}(\theta')|\mathbb{X}_{tk}}(C_2) \right| \right] \\ &\leq 2^k \mathbb{E}\Bigg[\max_{y \in \{0,1\}^d} \max_{D \in \{1, \ldots, k\}^d} \Big| \mathbb{P}_{\{I_{C_{2j}}(h^*_{t\theta}(X_{tj}))\}_{j \in D} | \{X_{tj}\}_{j \in D}}(\{y\}) \\ &\qquad - \mathbb{P}_{\{I_{C_{2j}}(h^*_{t\theta'}(X_{tj}))\}_{j \in D} | \{X_{tj}\}_{j \in D}}(\{y\}) \Big| \Bigg] \\ &\leq 2^k \sum_{y \in \{0,1\}^d} \sum_{D \in \{1, \ldots, k\}^d} \mathbb{E}\Bigg[\Big| \mathbb{P}_{\{I_{C_{2j}}(h^*_{t\theta}(X_{tj}))\}_{j \in D} | \{X_{tj}\}_{j \in D}}(\{y\}) \\ &\qquad - \mathbb{P}_{\{I_{C_{2j}}(h^*_{t\theta'}(X_{tj}))\}_{j \in D} | \{X_{tj}\}_{j \in D}}(\{y\}) \Big| \Bigg] \\ &\leq 2^{d+k} k^d \max_{y \in \{0,1\}^d} \max_{D \in \{1, \ldots, k\}^d} \mathbb{E}\Bigg[\Big| \mathbb{P}_{\{I_{C_{2j}}(h^*_{t\theta}(X_{tj}))\}_{j \in D} | \{X_{tj}\}_{j \in D}}(\{y\}) \\ &\qquad - \mathbb{P}_{\{I_{C_{2j}}(h^*_{t\theta'}(X_{tj}))\}_{j \in D} | \{X_{tj}\}_{j \in D}}(\{y\}) \Big| \Bigg]. \end{aligned}$$

Exchangeability implies this is at most

$$
\begin{aligned}
&2^{d+k}k^d \max_{y\in\{0,1\}^d} \sup_{t_1,\ldots,t_d\in\mathbb{R}} \mathbb{E}\Bigg[\bigg|\mathbb{P}_{\{I_{(-\infty,t_j]}(h^*_{t\theta}(X_{tj}))\}_{j=1}^d|\mathbb{X}_{td}}(\{y\}) \\
&\qquad\qquad - \mathbb{P}_{\{I_{(-\infty,t_j]}(h^*_{t\theta'}(X_{tj}))\}_{j=1}^d|\mathbb{X}_{td}}(\{y\})\bigg|\Bigg] \\
&= 2^{d+k}k^d \max_{y\in\{0,1\}^d} \sup_{t_1,\ldots,t_d\in\mathbb{R}} \mathbb{E}\Bigg[\bigg|\mathbb{P}_{\{I_{(-\infty,t_j]}(Y_{tj}(\theta))\}_{j=1}^d|\mathbb{X}_{td}}(\{y\}) \\
&\qquad\qquad - \mathbb{P}_{\{I_{(-\infty,t_j]}(Y_{tj}(\theta'))\}_{j=1}^d|\mathbb{X}_{td}}(\{y\})\bigg|\Bigg].
\end{aligned}
$$

[YHC13] argue that for all $y \in \{0,1\}^d$ and $t_1,\ldots,t_d \in \mathbb{R}$,

$$
\begin{aligned}
&\mathbb{E}\left[\left|\mathbb{P}_{\{I_{(-\infty,t_j]}(Y_{tj}(\theta))\}_{j=1}^d|\mathbb{X}_{td}}(\{y\}) - \mathbb{P}_{\{I_{(-\infty,t_j]}(Y_{tj}(\theta'))\}_{j=1}^d|\mathbb{X}_{td}}(\{y\})\right|\right] \\
&\qquad \le 4\sqrt{\|\mathbb{P}_{\{I_{(-\infty,t_j]}(Y_{tj}(\theta))\}_{j=1}^d,\mathbb{X}_{td}} - \mathbb{P}_{\{I_{(-\infty,t_j]}(Y_{tj}(\theta'))\}_{j=1}^d,\mathbb{X}_{td}}\|}.
\end{aligned}
$$

Noting that

$$
\|\mathbb{P}_{\{I_{(-\infty,t_j]}(Y_{tj}(\theta))\}_{j=1}^d,\mathbb{X}_{td}} - \mathbb{P}_{\{I_{(-\infty,t_j]}(Y_{tj}(\theta'))\}_{j=1}^d,\mathbb{X}_{td}}\| \le \|\mathbb{P}_{\mathcal{Z}_{td}(\theta)} - \mathbb{P}_{\mathcal{Z}_{td}(\theta')}\|
$$

completes the proof. □

We can use the above lemmas to design an estimator of $\pi_{\theta_\star}$. Specifically, we have the following result.

Theorem 1. *There exists an estimator* $\hat{\theta}_{T\theta_\star} = \hat{\theta}_T(\mathcal{Z}_{1d}(\theta_\star),\ldots,\mathcal{Z}_{Td}(\theta_\star))$*, and functions* $R : \mathbb{N}_0 \times (0,1] \to [0,\infty)$ *and* $\delta : \mathbb{N}_0 \times (0,1] \to [0,1]$ *such that, for any* $\alpha > 0$*,* $\lim_{T\to\infty} R(T,\alpha) = \lim_{T\to\infty} \delta(T,\alpha) = 0$ *and for any* $T \in \mathbb{N}_0$ *and* $\theta_\star \in \Theta$*,*

$$
\mathbb{P}\left(\|\pi_{\hat{\theta}_{T\theta_\star}} - \pi_{\theta_\star}\| > R(T,\alpha)\right) \le \delta(T,\alpha) \le \alpha.
$$

Proof. The estimator $\hat{\theta}_{T\theta_\star}$ we will use is precisely the minimum-distance skeleton estimate of $\mathbb{P}_{\mathcal{Z}_{td}(\theta_\star)}$ [Yat85, DL01]. [Yat85] proved that if $N(\varepsilon)$ is the ε-covering number of $\{\mathbb{P}_{\mathcal{Z}_{td}(\theta_\star)} : \theta \in \Theta\}$, then taking this $\hat{\theta}_{T\theta_\star}$ estimator, then for some $T_\varepsilon = O((1/\varepsilon^2)\log N(\varepsilon/4))$, any $T \ge T_\varepsilon$ has

$$
\mathbb{E}\left[\|\mathbb{P}_{\mathcal{Z}_{td}(\hat{\theta}_{T\theta_\star})} - \mathbb{P}_{\mathcal{Z}_{td}(\theta_\star)}\|\right] < \varepsilon.
$$

Thus, taking $G_T = \inf\{\varepsilon > 0 : T \ge T_\varepsilon\}$, we have

$$
\mathbb{E}\left[\|\mathbb{P}_{\mathcal{Z}_{td}(\hat{\theta}_{T\theta_\star})} - \mathbb{P}_{\mathcal{Z}_{td}(\theta_\star)}\|\right] \le G_T = o(1).
$$

Letting $R'(T,\alpha)$ be any positive sequence with $G_T \ll R'(T,\alpha) \ll 1$ and $R'(T,\alpha) \geq G_T/\alpha$, and letting $\delta(T,\alpha) = G_T/R'(T,\alpha) = o(1)$, Markov's inequality implies

$$\mathbb{P}\left(\|\mathbb{P}_{\mathcal{Z}_{td}(\hat{\theta}_{T\theta_\star})} - \mathbb{P}_{\mathcal{Z}_{td}(\theta_\star)}\| > R'(T,\alpha)\right) \leq \delta(T,\alpha) \leq \alpha. \tag{1}$$

Letting $R(T,\alpha) = \min_k \left(M_k\left(R'(T,\alpha)\right) + r_k\right)$, since $R'(T,\alpha) = o(1)$ and $r_k = o(1)$, we have $R(T,\alpha) = o(1)$. Furthermore, composing (1) with Lemmas 1, 2, and 3, we have

$$\mathbb{P}\left(\|\pi_{\hat{\theta}_{T\theta_\star}} - \pi_{\theta_\star}\| > R(T,\alpha)\right) \leq \delta(T,\alpha) \leq \alpha.$$

$\square$

Remark: Although the above result makes use of the minimum-distance skeleton estimator, which is typically not computationally efficient, it is often possible to achieve this same result (for certain families of distributions) using a simpler estimator, such as the maximum likelihood estimator. All we require is that the risk of the estimator converges to 0 at a known rate that is independent of $\theta_\star$. For instance, see [vdG00] for conditions on the family of distributions sufficient for this to be true of the maximum likelihood estimator.

3 Maximizing Customer Satisfaction in Combinatorial Auctions

We can use Theorem 1 in the context of various applications. For instance, consider the following application to the problem of serving a sequence of customers so as to maximize their satisfaction.

Suppose there is a menu of n items $[n] = \{1,\ldots,n\}$, and each bundle $B \subseteq [n]$ has an associated price $p(B) \geq 0$. Suppose also there is a sequence of customers, each with a valuation function $v_t : 2^{[n]} \to \mathbb{R}$. We suppose these v_t functions are i.i.d. samples. We can then calculate the satisfaction function for each customer as $s_t(x)$, where $x \in \{0,1\}^n$, and $s_t(x) = v_t(B_x) - p(B_x)$, where $B_x \subseteq [n]$ contains element $i \in [n]$ iff $x_i = 1$.

Now suppose we are able to ask each customer a number of questions before serving up a bundle $B_{\hat{x}_t}$ to that customer. More specifically, we are able to ask for the value $s_t(x)$ for any $x \in \{0,1\}^n$. This is referred to as a *value query* in the literature on preference elicitation in combinatorial auctions (see Chapter 14 of [CSS06], [ZBS03]). We are interested in asking as few questions as possible, while satisfying the guarantee that $\mathbb{E}[s_t(\hat{x}_t) - \max_x s_t(x)] \leq \varepsilon$.

Now suppose, for every π and ε, we have a method $A(\pi,\varepsilon)$ such that, given that π is the actual distribution of the s_t functions, $A(\pi,\varepsilon)$ guarantees that the $\hat{x}_t$ value it selects has $\mathbb{E}[\max_x s_t(x) - s_t(\hat{x}_t)] \leq \varepsilon$; also let $\hat{N}_t(\pi,\varepsilon)$ denote the actual (random) number of queries the method $A(\pi,\varepsilon)$ would ask for the s_t function, and let $Q(\pi,\varepsilon) = \mathbb{E}[\hat{N}_t(\pi,\varepsilon)]$. We suppose the method never queries any $s_t(x)$ value twice for a given t, so that its number of queries for any given t is bounded.

Also suppose $\mathcal{F}$ is a VC subgraph class of functions mapping $\mathcal{X} = \{0,1\}^n$ into $[-1,1]$ with pseudodimension d, and that $\{\pi_\theta : \theta \in \Theta\}$ is a known totally bounded family of distributions over $\mathcal{F}$ such that the s_t functions have distribution $\pi_{\theta_\star}$ for some unknown $\theta_\star \in \Theta$. For any $\theta \in \Theta$ and $\gamma > 0$, let $\mathrm{B}(\theta,\gamma) = \{\theta' \in \Theta : \|\pi_\theta - \pi_{\theta'}\| \leq \gamma\}$.

Suppose, in addition to A, we have another method $A'(\varepsilon)$ that is not π-dependent, but still provides the ε-correctness guarantee, and makes a bounded number of queries (e.g., in the worst case, we could consider querying all 2^n points, but in most cases there are more clever π-independent methods that use far fewer queries, such as $O(1/\varepsilon^2)$). Consider the following method; the quantities $\hat{\theta}_{T\theta_\star}$, $R(T,\alpha)$, and $\delta(T,\alpha)$ from Theorem 1 are here considered with respect P_X taken as the uniform distribution on $\{0,1\}^n$.

Algorithm 1 An algorithm for sequentially maximizing expected customer satisfaction.

for $t = 1, 2, \ldots, T$ **do**
  Pick points $X_{t1}, X_{t2}, \ldots, X_{td}$ uniformly at random from $\{0,1\}^n$
  if $R(t-1, \varepsilon/2) > \varepsilon/8$ **then**
    Run $A'(\varepsilon)$
    Take $\hat{x}_t$ as the returned value
  else
    Let $\check{\theta}_{t\theta_\star} \in \mathrm{B}\left(\hat{\theta}_{(t-1)\theta_\star}, R(t-1,\varepsilon/2)\right)$ be such that
    $Q(\pi_{\check{\theta}_{t\theta_\star}}, \varepsilon/4) \leq \min_{\theta \in \mathrm{B}\left(\hat{\theta}_{(t-1)\theta_\star}, R(t-1,\varepsilon/2)\right)} Q(\pi_\theta, \varepsilon/4) + 1/t$
    Run $A(\pi_{\check{\theta}_{t\theta_\star}}, \varepsilon/4)$ and let $\hat{x}_t$ be its return value
  end if
end for

The following theorem indicates that this method is correct, and furthermore that the long-run average number of queries is not much worse than that of a method that has direct knowledge of $\pi_{\theta_\star}$.

Theorem 2. *For the above method,* $\forall t \leq T, \mathbb{E}[\max_x s_t(x) - s_t(\hat{x}_t)] \leq \varepsilon$. *Furthermore, if $S_T(\varepsilon)$ is the total number of queries made by the method, then*

$$\limsup_{T\to\infty} \frac{\mathbb{E}[S_T(\varepsilon)]}{T} \leq Q(\pi_{\theta_\star}, \varepsilon/4) + d.$$

Proof. By Theorem 1, for any $t \leq T$, if $R(t-1,\varepsilon/2) \leq \varepsilon/8$, then with probability at least $1 - \varepsilon/2$, $\|\pi_{\theta_\star} - \pi_{\hat{\theta}_{(t-1)\theta_\star}}\| \leq R(t-1,\varepsilon/2)$, so that a triangle inequality implies $\|\pi_{\theta_\star} - \pi_{\check{\theta}_{t\theta_\star}}\| \leq 2R(t-1,\varepsilon/2) \leq \varepsilon/4$. Thus,

$$\begin{aligned} &\mathbb{E}\left[\max_x s_t(x) - s_t(\hat{x}_t)\right] \\ &\quad \leq \mathbb{E}\left[\mathbb{E}\left[\max_x s_t(x) - s_t(\hat{x}_t) \middle| \check{\theta}_{t\theta_\star}\right] \mathbb{1}\left[\|\pi_{\check{\theta}_{t\theta_\star}} - \pi_{\theta_\star}\| \leq \varepsilon/2\right]\right] + \varepsilon/2. \quad (2) \end{aligned}$$

For $\theta \in \Theta$, let $\hat{x}_{t\theta}$ denote the point x that would be returned by $A(\pi_{\check{\theta}_{t\theta_\star}}, \varepsilon/4)$ when queries are answered by some $s_{t\theta} \sim \pi_\theta$ instead of s_t (and supposing $s_t = s_{t\theta_\star}$). If

$\|\pi_{\check{\theta}_{t\theta_\star}} - \pi_{\theta_\star}\| \leq \varepsilon/4$, then

$$\mathbb{E}\left[\max_x s_t(x) - s_t(\hat{x}_t)\middle|\check{\theta}_{t\theta_\star}\right] = \mathbb{E}\left[\max_x s_{t\theta_\star}(x) - s_{t\theta_\star}(\hat{x}_t)\middle|\check{\theta}_{t\theta_\star}\right]$$
$$\leq \mathbb{E}\left[\max_x s_{t\check{\theta}_{t\theta_\star}}(x) - s_{t\check{\theta}_{t\theta_\star}}(\hat{x}_{t\check{\theta}_{t\theta_\star}})\middle|\check{\theta}_{t\theta_\star}\right] + \|\pi_{\check{\theta}_{t\theta_\star}} - \pi_{\theta_\star}\| \leq \varepsilon/4 + \varepsilon/4 = \varepsilon/2.$$

Plugging into (2), we have

$$\mathbb{E}\left[\max_x s_t(x) - s_t(\hat{x}_t)\right] \leq \varepsilon.$$

For the result on $S_T(\varepsilon)$, first note that $R(t-1, \varepsilon/2) > \varepsilon/8$ only finitely many times (due to $R(t, \alpha) = o(1)$), so that we can ignore those values of t in the asymptotic calculation (as the number of queries is always bounded), and rely on the correctness guarantee of A' for correctness. For the remaining t values, let N_t denote the number of queries made by $A(\pi_{\check{\theta}_{t\theta_\star}}, \varepsilon/4)$. then

$$\limsup_{T\to\infty} \frac{\mathbb{E}[S_T(\varepsilon)]}{T} \leq d + \limsup_{T\to\infty} \sum_{t=1}^{T} \mathbb{E}\left[N_t\right]/T.$$

Since

$$\lim_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} \mathbb{E}\left[N_t \mathbb{1}[\|\pi_{\hat{\theta}_{(t-1)\theta_\star}} - \pi_{\theta_\star}\| > R(t-1, \varepsilon/2)]\right]$$
$$\leq \lim_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} 2^n \mathbb{P}\left(\|\pi_{\hat{\theta}_{(t-1)\theta_\star}} - \pi_{\theta_\star}\| > R(t-1, \varepsilon/2)\right)$$
$$\leq 2^n \lim_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} \delta(t-1, \varepsilon/2) = 0,$$

we have

$$\limsup_{T\to\infty} \sum_{t=1}^{T} \mathbb{E}\left[N_t\right]/T = \limsup_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} \mathbb{E}\left[N_t \mathbb{1}[\|\pi_{\hat{\theta}_{(t-1)\theta_\star}} - \pi_{\theta_\star}\| \leq R(t-1, \varepsilon/2)]\right].$$

For any $t \leq T$, let $N_t(\check{\theta}_{t\theta_\star})$ denote the number of queries $A(\pi_{\check{\theta}_{t\theta_\star}}, \varepsilon/4)$ would make if queries were answered with $s_{t\check{\theta}_{t\theta_\star}}$ instead of s_t. On the event $\|\pi_{\hat{\theta}_{(t-1)\theta_\star}} - \pi_{\theta_\star}\| \leq R(t-1, \varepsilon/2)$, we have

$$\mathbb{E}\left[N_t\middle|\check{\theta}_{t\theta_\star}\right] \leq \mathbb{E}\left[N_t(\check{\theta}_{t\theta_\star})\middle|\check{\theta}_{t\theta_\star}\right] + 2R(t-1, \varepsilon/2)$$
$$= Q(\pi_{\check{\theta}_{t\theta_\star}}, \varepsilon/4) + 2R(t-1, \varepsilon/2) \leq Q(\pi_{\theta_\star}, \varepsilon/4) + 2R(t-1, \varepsilon/2) + 1/t.$$

Therefore,

$$\limsup_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} \mathbb{E}\left[N_t \mathbb{1}[\|\pi_{\hat{\theta}_{(t-1)\theta_\star}} - \pi_{\theta_\star}\| \leq R(t-1, \varepsilon/2)]\right]$$
$$\leq Q(\pi_{\theta_\star}, \varepsilon/4) + \limsup_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} 2R(t-1, \varepsilon/2) + 1/t = Q(\pi_{\theta_\star}, \varepsilon/4).$$

$\square$

Note that in many cases, this result will even continue to hold with an infinite number of goods ($n = \infty$), since the general results of the previous section have no dependence on the cardinality of the space $\mathcal{X}$.

References

[CSS06] P. Cramton, Y. Shoham, and R. Steinberg. *Combinatorial Auctions*. The MIT Press, 2006.

[DL01] L. Devroye and G. Lugosi. *Combinatorial Methods in Density Estimation*. Springer, New York, NY, USA, 2001.

[Hau92] D. Haussler. Decision theoretic generalizations of the PAC model for neural net and other learning applications. *Information and Computation*, 100:78–150, 1992.

[Sch95] M. J. Schervish. *Theory of Statistics*. Springer, New York, NY, USA, 1995.

[Vap82] V. Vapnik. *Estimation of Dependencies Based on Empirical Data*. Springer-Verlag, New York, 1982.

[vdG00] S. van de Geer. *Empirical Processes in M-Estimation*. Cambridge University Press, 2000.

[Yat85] Y. G. Yatracos. Rates of convergence of minimum distance estimators and Kolmogorov's entropy. *The Annals of Statistics*, 13:768–774, 1985.

[YHC13] L. Yang, S. Hanneke, and J. Carbonell. A theory of transfer learning with applications to active learning. *Machine Learning*, 90(2):161–189, 2013.

[ZBS03] M. Zinkevich, A. Blum, and T. Sandholm. On polynomial-time preference elicitation with value queries. In *Proceedings of the* 4^{th} *ACM Conference on Electronic Commerce*, pages 175–185, 2003.