JMLR: Workshop and Conference Proceedings 19 (2011) 293–314 24th Annual Conference on Learning Theory

Complexity-Based Approach to Calibration with Checking Rules

Dean P. Foster
Department of Statistics
University of Pennsylvania

Alexander Rakhlin
Department of Statistics
University of Pennsylvania

Karthik Sridharan
TTI-Chicago

Ambuj Tewari
Department of Computer Science
University of Texas at Austin

Editor: Sham Kakade, Ulrike von Luxburg

Abstract

We consider the problem of forecasting a sequence of outcomes from an unknown source. The quality of the forecaster is measured by a family of checking rules. We prove upper bounds on the value of the associated game, thus certifying the existence of a calibrated strategy for the forecaster. We show that complexity of the family of checking rules can be captured by the notion of a sequential cover introduced in (Rakhlin et al., 2010a). Various natural assumptions on the class of checking rules are considered, including finiteness of Vapnik-Chervonenkis and Littlestone's dimensions.

1. Introduction

As many other papers on calibration, we start with the following motivating example: Consider a weatherman who predicts the probability of rain tomorrow and then observes the binary "rain/no rain" outcome. How can we measure the weatherman's performance? If we make no assumption on the way Nature selects outcomes, defining a notion of performance is a non-trivial matter. One approach, familiar to the learning community, is to prove regret bounds with respect to some class of strategies. However, in the absence of any assumptions on the sequence, the performance of the comparator will not be favorable, rendering the bounds meaningless. An alternative measure of performance is to ask that the forecaster satisfies certain properties with respect to the sequence. One such natural property is *calibration*. It posits that for all the days that the forecaster predicted a probability p of rain, the empirical frequency of rain was indeed close to p. It is not obvious, a priori, that there exists a forecasting strategy calibrated with respect to every p, no matter what sequences Nature presents. The question was raised in the Bayesian setting by Dawid (1982), followed by the negative result of Oakes (1985), who showed that no deterministic calibration strategy exists. The first positive result was shown by Foster and Vohra (1998), who provided a randomized calibration strategy.

Calibration is indeed the absolute minimum we should expect from a forecaster. Clearly a forecaster who makes a constant prediction of .6 on the binary sequence 11.0010010000111111... for π (which empirically is one half ones and believed by most to be half ones in the limit) should be fired at some point for a failure to be calibrated (Lehrer, 2004). However, forecasting the right overall frequency might not be enough. Indeed, consider a binary sequence "010101 . . ." of "rain/no rain" outcomes. A forecaster predicting 0.5 chance of rain is calibrated, yet such a lousy weatherman should be fired immediately! To cope with the obvious shortcoming of calibration, one may introduce more complex *checking rules* (Kalai et al., 1999; Sandroni et al., 2003; Cesa-Bianchi and Lugosi, 2006), such as "the forecaster should be calibrated on all even rounds." This additional rule clearly disallows a constant prediction of 0.5 since within the even rounds the empirical frequency is 1. While resolving the problem with the particular sequence "010101 . . .," the forecaster's performance might still appear unacceptable (by our standards) on other sequences. We refer to (Sandroni et al., 2003) for further discussion on checking rules.

How rich can we make the set of checking rules while being able to satisfy all of them at the same time? Of course, if checking rules are completely arbitrary, there is no hope, as the rule can be tailored to the particular sequence presented. It is then natural to ask the following questions: What is a sufficient restriction on the class of checking rules? What are the relevant measures of complexity of infinite classes of checking rules? What governs the rates of convergence in calibration? In addressing these matters, we come to questions of martingale convergence for function classes. In particular, this allows us to make a connection to the Vapnik-Chervonenkis theory which measures the complexity of the class using a combinatorial parameter. We can view the classical calibration results as a particular instance of checking rules with a finite VC dimension. To the best of our knowledge, the connection between calibration and statistical learning has not been previously observed.

Our results are based on tools recently developed in (Rakhlin et al., 2010b,a). These papers consider abstract repeated zero-sum games (subsuming Online Learning) and obtain upper bounds on the minimax value via the process of sequential symmetrization. Interestingly, these bounds are attained without explicitly talking about algorithms, and instead focusing on the inherent complexity of the problem. Analogously, in the present paper we prove convergence results which depend on the complexity of the class of checking rules without providing a computationally efficient algorithm (the inefficient algorithm can be recovered from the minimax formulation). We argue that an understanding of what is attainable in terms of satisfying checking rules is necessary before looking for an efficient implementation. Once the inherent complexity of calibration with checking rules is understood, algorithmic questions will arise. While there is an efficient algorithm for classical calibration with two actions (see Foster and Vohra (1998); Abernethy et al. (2011)), the question is still open for more complex classes of checking rules.

Classical decision theory typically divides problems into two pieces, probability and loss, and then combines these (via expectation) for making decisions. Calibrated forecasts allow this same division to be done in the setting of individual sequences: a probabilistic forecast can be made and then a loss function can be optimized as if these probabilities were in fact correct. These decisions can be made in a game theoretic setting, in which case calibrated forecasts can lead to equilibria in games (Foster and Vohra, 1997; Kakade and Foster, 2008). But unlike traditional decision theory which has viewed this division of decisions into probability and loss as having zero cost, there is a huge cost when using calibration in this way for individual sequences. Namely, the rates of convergence for a calibrated forecast have often been much poorer than the ones generated by optimizing the decisions directly, as is typically done in the experts literature. The cause of this rate difference is that calibration tries to optimize over details that the experts approach would ignore. We present alternative definitions of calibration that address this by focusing attention only on the parts of calibration that translate into difference at the decision-making level. We refer to (Young, 2004) for connections between calibration, decision making, and games.

Another motivation for studying checking rules comes from recent research at the intersection of game theory, learning, and economics, which often involves multiple agents acting in the world (Kakade et al., 2003). Being able to calibrate with respect to a class of checking rules can lead to good guarantees on the quality of actions taken by agents. For instance, one can consider multi-agent decision-making problems in large environments, where the agents only need to calibrate with respect to a small set of checking rules relevant to their decision making.

2. Notation

Let $\mathbb{E}_{x\sim p}$ denote expectation with respect to a random variable x with a distribution p. A Rademacher random variable is a symmetric ±1-valued random variable. The notation $x_{a:b}$ denotes the sequence $x_a, \ldots, x_b$. The indicator of an event A is denoted by $\mathbf{1}\{A\}$. The set $\{1, \ldots, T\}$ is denoted by $[T]$, while the $(k-1)$-dimensional probability simplex in $\mathbb{R}^k$ is denoted by Δ_k. Let E_k denote the k vertices of Δ_k. The set of all functions from $\mathcal{X}$ to $\mathcal{Y}$ is denoted by $\mathcal{Y}^{\mathcal{X}}$, and the t-fold product $\mathcal{X}\times\ldots\times\mathcal{X}$ is denoted by $\mathcal{X}^t$. Whenever a supremum (or infimum) is written in the form $\sup_a$ without a being quantified, it is assumed that a ranges over the set of all possible values which will be understood from the context.

Following (Rakhlin et al., 2010a), we define binary trees as follows. Consider a binary tree of uniform depth T where every interior node and every leaf is labeled with a value X chosen from some set $\mathcal{X}$. More precisely, given some set $\mathcal{X}$, an *$\mathcal{X}$-valued tree of depth T* is a sequence $(\mathbf{x}_1, \ldots, \mathbf{x}_T)$ of T mappings $\mathbf{x}_i : \{\pm1\}^{i-1} \mapsto \mathcal{X}$. Unless specified otherwise, $\epsilon = (\epsilon_1, \ldots, \epsilon_T) \in \{\pm1\}^T$ will define a path. For brevity, we will write $\mathbf{x}_t(\epsilon)$ instead of $\mathbf{x}_t(\epsilon_{1:t-1})$.

3. The Setting

In this paper we consider the k-outcome calibration game (in the weatherman example, $k = 2$). Each outcome is represented by an element of E_k, whereas the forecast is represented by a point in Δ_k. More precisely, the protocol can be viewed as the T-round game between player (learner) and the adversary (Nature):

FOR round $t = 1, \ldots, T$,

- the player chooses a mixed strategy $q_t \in \Delta(\Delta_k)$ (distribution on Δ_k)
- the adversary picks outcome $x_t \in E_k$

- the learner draws $f_t \in \Delta_k$ from q_t and observes outcome x_t

ENDFOR

Both opponents can base their next move on the history of actions observed so far. In particular, this makes the adversary *adaptive*. Throughout the paper, $z_t \in \mathcal{Z}$ is given by $z_t = ((f_1, x_1), \ldots, (f_{t-1}, x_{t-1}))$, the history of actions by both players at round t. Define the set of all possible histories by $\mathcal{Z} = \bigcup_{t=1}^{T} (\Delta_k \times E_k)^t$.

Definition 1 *A forecast-based checking rule is a binary-valued function $c : \mathcal{Z} \times \Delta_k \mapsto \{0, 1\}$.*

In other words, a checking rule depends on both the history and the current forecast. For simplicity, we only consider binary-valued checking rules; however, the results can be extended to real-valued functions and will appear in the full version of the paper.

Let ζ be a family of checking rules. The goal of the player is to minimize the performance metric

$$\mathbf{R}_T := \sup_{c \in \zeta} \left\| \frac{1}{T} \sum_{t=1}^{T} c(z_t, f_t) \cdot (f_t - x_t) \right\|$$

for some norm $\|\cdot\|$ on $\mathbb{R}^k$. While the ℓ_1 norm is typically used for calibration (Mannor and Stoltz, 2010), we can consider a general ℓ_p norm for $1 \leq p \leq \infty$. Informally, $\mathbf{R}_T$ says that the player needs to be calibrated (that is, average of forecasts close to the actual frequency) for any rule c that becomes active only on certain rounds. In the asymptotic sense, any rule that is not active infinitely often does not matter for the player.

Example 1 *For classical ϵ-calibration, choose $\zeta = \{c_p(z_t, f_t) = \mathbf{1}\{\|f_t - p\| \leq \epsilon \ : \ p \in \Delta_k\}\}$. In particular, ϵ-calibration captures the weather forecasting example discussed earlier. We refer to (Cesa-Bianchi and Lugosi, 2006; Mannor and Stoltz, 2010) for the details on the relationship between ϵ-calibration and well-calibration.*

Example 2 *Let $\mathcal{G}$ be an ϵ grid of the Δ_k. Define*

$$\zeta = \{c_A(z_t, f_t) = \mathbf{1}\{\|f_t - a\| \leq \epsilon \textit{ for some } a \in A\}\}_{A \in 2^{\mathcal{G}}}.$$

That is, c_A captures the set of forecasts for which f_t either over-forecasts or under-forecasts the correct probability of the outcome. This is a much richer set of rules than the previous example and is the implicit set used in the Brier quadratic calibration score used in (Foster and Vohra, 1998). As we will show later, the rate of convergence is much slower than for classical calibration.

Example 3 *Let $\hat{p}_{\theta,t}$ be the forecast made by a probabilistic model P_θ. Using $\zeta = \{c_{\theta,p}(z_t, f_t) = \mathbf{1}\{\|\hat{p}_{\theta,t} - p\| \leq \epsilon\}\}$ will test if the model P_θ is a much better fit to the data than the forecasting rule f_t. If complexity of the set of models $\{P_\theta\}$ is controlled, then theorems we will discuss later will guarantee existence of a rule that can do well against this family of tests. This connects to the testing of experts literature (Olszewski and Sandroni, 2009).*

Given the set ζ of checking rules, when is it possible to find a strategy for the forecaster such that $\mathbf{R}_T$ goes to zero as T increases? Instead of using, for instance, Blackwell's approachability to provide a calibration strategy with respect to the class ζ (as done in (Foster and Vohra, 1998; Sandroni et al., 2003)), we directly attack the value of the game. Given a $\theta > 0$, we define the value of the calibration game as

$$\mathcal{V}_T^\theta(\zeta) := \inf_{q_1} \sup_{x_1} \mathbb{E}_{f_1 \sim q_1} \dots \inf_{q_T} \sup_{x_T} \mathbb{E}_{f_T \sim q_T} \left[\mathbf{1} \left\{ \sup_{c \in \zeta} \left\| \frac{1}{T} \sum_{t=1}^{T} c(z_t, f_t) \cdot (f_t - x_t) \right\| > \theta \right\} \right],$$

where q_t's range over all distributions over Δ_k and x_t range over E_k. Note that the value can be interpreted as the probability of the performance metric $\mathbf{R}_T$ being larger than θ under the stochastic process arising from the successive infima, suprema, and expectations. An upper bound on $\mathcal{V}_T^\theta(\zeta)$ implies *existence* of a strategy for the learner such that the calibration metric $\mathbf{R}_T$ is smaller than θ with probability at least $1 - \mathcal{V}_T^\theta(\zeta)$. Or put more colloquially, our bound on $\mathcal{V}_T$ is an upper bound on the probability of the weatherman being fired for failure to be calibrated to accuracy θ. Alternatively, lower bounds on $\mathcal{V}_T^\theta(\zeta)$ imply impossibility results for the learner. Note that the definition of value of the game is for a fixed θ and number of rounds T. Thus, it is not obvious how to use the so-called "doubling trick" to get a player strategy that is Hannan consistent for the calibration game. The main difficulty is the dependence of the game (and hence the optimal player strategy) on θ. It is possible to define a game where the optimal player strategy will work *uniformly* over all θ (see (Rakhlin et al., 2010b)). Once this is done, we can proceed along similar lines as in Mannor and Stoltz (2010) to guarantee the existence of a Hannan consistent strategy for calibration with only an extra logarithmic factor on number of rounds played. However, for simplicty, we stick to the fixed θ, T definition above in this paper.

4. General Upper Bound on the Value $\mathcal{V}_T^\theta(\zeta)$

Let $\delta > 0$ and let C_δ be a minimal δ-cover of Δ_k in the norm $\|\cdot\|$. The size of the δ-cover can be bounded as

$$|C_\delta| \le (c_1/(2\delta))^{k-1} \quad . \tag{1}$$

where c_1 is some constant independent of k, but varying with the choice of the norm $\|\cdot\|$. This constant will appear throughout the paper. Further, for any $p_t \in \Delta_k$, let $p_t^\delta \in C_\delta$ be a point in C_δ such that $\|p_t - p_t^\delta\| \le \delta$. Slightly abusing the notation, define $z_t^\delta = \big((p_1^\delta, x_1), \dots, (p_{t-1}^\delta, x_{t-1})\big) \in \mathcal{Z}^\delta \subseteq \mathcal{Z}$ where $\mathcal{Z}^\delta := \bigcup_{t=1}^T (C_\delta \times E_k)^{t-1}$. (For the proofs of Lemmata 2–4, see Sec. 7 & Appendix.)

Lemma 2 *For any $\theta > 0$,*

$$\mathcal{V}_T^\theta(\zeta) \le \sup_{p_1} \mathbb{E}_{x_1 \sim p_1} \dots \sup_{p_T} \mathbb{E}_{x_T \sim p_T} \left[\mathbf{1} \left\{ \sup_{c \in \zeta} \left\| \frac{1}{T} \sum_{t=1}^{T} c(z_t^\delta, p_t^\delta) \cdot (p_t - x_t) \right\| > \theta/2 \right\} \right] \tag{2}$$

for any $\delta \le \theta/2$.

The interleaved suprema and expectations on the right-hand side of (2) can be written more succinctly as

$$\sup_{\mathbf{p}} \mathbb{E}\left[\mathbf{1}\left\{\sup_{c\in\zeta}\left\|\frac{1}{T}\sum_{t=1}^{T} c(z_t^\delta, p_t^\delta)\cdot(p_t - x_t)\right\| > \theta/2\right\}\right] \tag{3}$$

where $\mathbf{p}$ can be either thought of as a joint distribution over sequences $(x_1,\dots,x_T)$ or as a sequence of conditional distributions $\{p_t : E_k^{t-1} \to \Delta_k\}$. Using the notation of conditional distributions, the expectation in (3) can be expanded as $\mathbb{E}_{x_1\sim p_1}\mathbb{E}_{x_2\sim p_2(\cdot|x_1)}\mathbb{E}_{x_T\sim p_T(\cdot|x_{1:T-1})}$. Of course, expected value of an indicator is just the probability of the event. The goal is to relate (3) to the probability that the norm $\|\cdot\|$ of the average of a martingale difference sequence is large. The latter probability is exponentially small by a concentration of measure result which we present next.

Lemma 3 *For any $\mathbb{R}^k$-valued martingale difference sequence $\{d_t\}_{t=1}^T$ with $\|d_t\| \le 1$ a.s. for all $t \in [T]$, there exists a k-dependent constant c_k such that*

$$\mathbb{P}\left(\left\|\frac{1}{T}\sum_{t=1}^{T} d_t\right\| > \theta\right) \le 2\exp\left(-\frac{T\theta^2}{c_k}\right) .$$

In particular, $c_k = 8k$ for any ℓ_p norm with $1 \le p \le \infty$.

Armed with a concentration result for martingales, we apply the sequential symmetrization technique (see (Rakhlin et al., 2010b) for the high-probability version). In the lemma below, the supremum is over all binary E_k-valued trees $\mathbf{x}$ of depth T, as well as all binary C_δ-valued trees $\mathbf{p}^\delta$ of depth T. Given $\mathbf{x}, \mathbf{p}^\delta$, let the $\mathcal{Z}^\delta$-valued tree $\mathbf{z}^\delta$ be defined by

$$\mathbf{z}_t^\delta(\epsilon) = \Big((\mathbf{p}_1^\delta(\epsilon), \mathbf{x}_1(\epsilon)), \dots, (\mathbf{p}_{t-1}^\delta(\epsilon), \mathbf{x}_{t-1}(\epsilon))\Big)$$

for any $t \in [T]$. We also write $\mathbf{z}^{(\mathbf{x},\mathbf{p}^\delta)}$ instead of $\mathbf{z}^\delta$ to make the dependence on $\mathbf{x}, \mathbf{p}^\delta$ explicit.

Lemma 4 *For $T > \frac{16 c_k \log(4)}{\theta^2}$ and $\delta \le \theta/2$,*

$$\mathcal{V}_T^\theta(\zeta) \le 4 \sup_{\mathbf{x},\mathbf{p}^\delta} \ \mathbb{P}_\epsilon\left(\sup_{c\in\zeta}\left\|\frac{1}{T}\sum_{t=1}^{T}\epsilon_t\ c(\mathbf{z}_t^\delta(\epsilon), \mathbf{p}_\mathbf{t}^\delta(\epsilon))\ \mathbf{x}_t(\epsilon)\right\| > \theta/8\right) ,$$

where the probability is over an i.i.d. draw of Rademacher random variables $\epsilon_1, \dots, \epsilon_T$.

What has been achieved by this lemma? We were able to pass from the quantity in (3) which is defined with respect to a complicated stochastic process to a much simpler process. It is defined by fixing the worst-case trees (in the spaces of moves of the adversary and the player) and then generating the process by coin flips ϵ_t. The resulting quantity is a symmetrized one and can be seen as a sequential version of the classical Rademacher complexity. In particular, the symmetrized upper bound of Lemma 4 allows us to define appropriate covering numbers and thus analyze infinite classes of checking rules.

The definitions of a sequential *cover* and *covering number* below are from (Rakhlin et al., 2010a). Note that they differ from the corresponding classical "static" notions.

Definition 5 *Consider a binary-valued function class $\mathcal{G} \subseteq \{0,1\}^{\mathcal{Y}}$ over some set $\mathcal{Y}$. For any given $\mathcal{Y}$-valued tree $\mathbf{y}$ of depth T, a set V of binary-valued trees of depth T is called* a 0-cover of $\mathcal{G}$ on $\mathbf{y}$ *if*

$$\forall g \in \mathcal{G},\ \forall \epsilon \in \{\pm 1\}^T,\ \exists \mathbf{v} \in V \quad s.t. \qquad \forall t \in [T], \qquad g(\mathbf{y}_t(\epsilon)) = \mathbf{v}_t(\epsilon)\ . \tag{4}$$

The *covering number* at scale 0 of a class $\mathcal{G}$ (the 0-covering number) on a given tree $\mathbf{y}$ is defined as

$$N(\mathcal{G}, \mathbf{y}) = \min\{|V| : V \text{ is a 0-cover of } \mathcal{G} \text{ on } \mathbf{y}\}\,.$$

Also define the worst-case covering number for all depth-T trees as $N(\mathcal{G}, T) = \sup_{\mathbf{y}} N(\mathcal{G}, \mathbf{y})$.

We point out that the order of quantifiers in (4) is crucial: For a given function g, the covering tree $\mathbf{v}$ can be chosen based on the path ϵ itself. It is thus not correct to think of the 0-cover as the number of distinct trees obtained by evaluating all functions from $\mathcal{G}$ on the given $\mathbf{y}$. Indeed, as described in (Rakhlin et al., 2010a), it is possible for an exponentially-large set of functions $\mathcal{G}$ to have a 0-cover of size 2, capturing the temporal structure of $\mathcal{G}$.

Definition 6 *Define the* minimal checking covering number *of ζ over depth T trees as*

$$\mathcal{N}_{\mathrm{ch}}(\zeta, T) = \sup_{\mathbf{x}, \mathbf{p}^\delta} N(\zeta, (\mathbf{z}^{(\mathbf{x}, \mathbf{p}^\delta)}, \mathbf{p}^\delta))$$

and the minimal checking cover *of ζ on $\mathbf{x}, \mathbf{p}^\delta$ as the set of size $N(\zeta, (\mathbf{z}^{(\mathbf{x}, \mathbf{p}^\delta)}, \mathbf{p}^\delta))$ that provides the cover. Here, abusing notation, $(\mathbf{z}^{(\mathbf{x}, \mathbf{p}^\delta)}, \mathbf{p}^\delta)$ is the $\mathcal{Z}^\delta \times C_\delta$-valued tree obtained by pairing the trees $\mathbf{z}^{(\mathbf{x}, \mathbf{p}^\delta)}$ and $\mathbf{p}^\delta$ together (and note that ζ is a class of binary functions on $\mathcal{Z}^\delta \times C_\delta$).*

Importantly, the minimal checking covering number is defined only over history trees $\mathbf{z}^{(\mathbf{x}, \mathbf{p}^\delta)}$ consistent with the chosen trees $\mathbf{x}, \mathbf{p}^\delta$. Clearly, we can upper bound the minimal checking covering number by the minimal cover $N(\zeta, T)$ over $\mathcal{Z}^\delta \times C_\delta$. It is immediate that $\mathcal{N}_{\mathrm{ch}}(\zeta, T) \le N(\zeta, T)$.

Theorem 7 *For $T > \frac{16 c_k \log(4)}{\theta^2}$ and $\delta \le \theta/2$,*

$$\mathcal{V}_T^\theta(\zeta) \le 8\ \mathcal{N}_{\mathrm{ch}}(\zeta, T) \exp\left(-\frac{T\theta^2}{64\ c_k}\right)$$

Proof [**Theorem 7**] Given any trees $\mathbf{x}, \mathbf{p}^\delta$, let the set of binary valued trees V be a (finite) minimal checking cover of ζ on $\mathbf{x}, \mathbf{p}^\delta$. For any $c \in \zeta$, let $\mathbf{v}[c, \epsilon] \in V$ be the member of the minimal checking cover that matches c on the tree $(\mathbf{x}, \mathbf{p}^\delta)$ over the path ϵ. Then we see that

$$\begin{aligned}\mathbb{P}_\epsilon\left(\sup_{c\in\zeta}\left\|\frac{1}{T}\sum_{t=1}^T \epsilon_t\ c(\mathbf{z}_t^\delta(\epsilon), \mathbf{p}_\mathbf{t}^\delta(\epsilon))\ \mathbf{x}_t(\epsilon)\right\| > \theta/8\right) &= \mathbb{P}_\epsilon\left(\sup_{c\in\zeta}\left\|\frac{1}{T}\sum_{t=1}^T \epsilon_t\ \mathbf{v}[c,\epsilon]_t(\epsilon)\ \mathbf{x}_t(\epsilon)\right\| > \theta/8\right)\\ &\le \mathbb{P}_\epsilon\left(\max_{\mathbf{v}\in V}\left\|\frac{1}{T}\sum_{t=1}^T \epsilon_t\ \mathbf{v}_t(\epsilon)\ \mathbf{x}_t(\epsilon)\right\| > \theta/8\right)\end{aligned}$$

Since $|V|$ is finite, by union bound we pass to the upper bound of

$$|V| \max_{\mathbf{v}\in V} \mathbb{P}_\epsilon \left(\left\| \frac{1}{T} \sum_{t=1}^{T} \epsilon_t \ \mathbf{v}_t(\epsilon) \ \mathbf{x}_t(\epsilon) \right\| > \theta/8 \right) \leq \mathcal{N}_{\mathrm{ch}}(\zeta, T) \max_{\mathbf{v}\in V} \mathbb{P}_\epsilon \left(\left\| \frac{1}{T} \sum_{t=1}^{T} \epsilon_t \ \mathbf{v}_t(\epsilon) \ \mathbf{x}_t(\epsilon) \right\| > \theta/8 \right)$$

We now appeal to Lemma 3. Note that $\mathbf{v}$ is binary-valued and $\mathbf{x}$ is E_k-valued, and, hence, $\|\mathbf{v}_t(\epsilon) \ \mathbf{x}_t(\epsilon)\| \leq 1$ for any t. Also, $\epsilon_t \ \mathbf{v}_t(\epsilon) \ \mathbf{x}_t(\epsilon)$ is a martingale difference sequence since $\mathbf{x}_t$ and $\mathbf{v}_t$ by definition only depend on $\epsilon_{1:t-1}$. Hence, for any $\mathbf{x}$ and $\mathbf{v}$,

$$\mathbb{P}_\epsilon \left(\left\| \frac{1}{T} \sum_{t=1}^{T} \epsilon_t \ \mathbf{v}_t(\epsilon) \ \mathbf{x}_t(\epsilon) \right\| > \theta/8 \right) \leq 2 \exp\left(-\frac{T\theta^2}{64 \ c_k} \right)$$

Combining with Lemma 4, we have that

$$\mathcal{V}_T^\theta(\zeta) \leq 4 \sup_{\mathbf{x},\mathbf{p}^\delta} \ \ \mathbb{P}_\epsilon \left(\sup_{c\in\zeta} \left\| \frac{1}{T} \sum_{t=1}^{T} \epsilon_t \ c(\mathbf{z}_t^\delta(\epsilon), \mathbf{p}_\mathbf{t}^\delta(\epsilon)) \ \mathbf{x}_t(\epsilon) \right\| > \theta/8 \right) \leq 8 \ \mathcal{N}_{\mathrm{ch}}(\zeta, T) \exp\left(-\frac{T\theta^2}{64 \ c_k} \right) \ .$$

∎

5. Families of Checking Rules

The main objective of this paper is to find general sufficient conditions on the set of checking rules that guarantee existence of a calibrated strategy. Theorem 7 guarantees decay of $\mathcal{V}_T^\theta(\zeta)$ if checking covering numbers of ζ can be controlled. In this section, we show control of these numbers under various assumptions on ζ, along with the resulting rates of convergence.

5.1. Finite Class of Checking Rules

The first straightforward consequence of Theorem 7 is that, for a finite class ζ,

$$\mathcal{V}_T^\theta(\zeta) \leq 8 \, |\zeta| \ \exp\left(-\frac{T\theta^2}{64 \ c_k} \right) \tag{5}$$

for $T > \frac{16 c_k \log(4)}{\theta^2}$. We can convert this statement into a probability of $\mathbf{R}_T$ being large. To this end, setting the right-hand side of (5) to η and solving for θ, we obtain

$$\theta = \sqrt{\frac{64 c_k \log(8|\zeta|/\eta)}{T}}.$$

For this value, the condition $T > \frac{16 c_k \log(4)}{\theta^2}$ is automatically satisfied. We can then state the result for finite ζ as follows: There exists a randomized strategy for the player such that

$$\mathbb{P}\left(\mathbf{R}_T \leq \sqrt{\frac{64 c_k \log(8|\zeta|/\eta)}{T}} \right) \geq 1 - \eta$$

for any $\eta > 0$, no matter how Nature chooses the outcomes.

As an example, consider the classic problem of digit identification, with the images of digits presented as "side information". A system that generates a prediction and gets scored against the true digit is then being effectively tested by a total of 10 checking rules.

5.2. History Invariant Checking Rules

A finite class of checking rules is, in some sense, too easy for the forecaster. Once we go to infinite classes, much of the difficulty arises from potentially complicated dependence of the rules on the history. Before attacking infinite classes of history-dependent rules, we consider the case of history-independence. The classical notion of calibration is an example of such a class of checking rules.

Formally, assume that ζ is a class of checking rules such that for all $c \in \zeta$, pair of histories $z, z' \in \mathcal{Z}$ and $p \in \Delta_k$:

$$c\,(z, p) = c\left(z', p\right)$$

Abusing notation, we can write each $c \in \zeta$ as a function $c : \Delta_k \mapsto \{0, 1\}$.

The next lemma recovers the rates obtained by Mannor and Stoltz (2010). For $k = 2$, the rate $T^{-1/3}$ has been also found previously by a variety of algorithms that reduced calibration on an ϵ-grid to the experts problem of no-internal regret with $O(1/\epsilon)$ experts.

Lemma 8 *For any class ζ of history invariant measurable checking rules, for any $\theta \in (0, 1]$ we have that*

$$\mathcal{V}_T^{\theta}(\zeta) \leq 8 \exp\left(-\frac{T\theta^2}{64\ c_k} + \left(\frac{c_1}{\theta}\right)^{k-1}\right)$$

for $T > \frac{16 c_k \log(4)}{\theta^2}$. This leads to

$$\mathbb{P}\left(\ \mathbf{R}_T\ \leq\ c_k'\ T^{-1/(k+1)} \sqrt{\log(8/\eta)}\ \right) \leq 1 - \eta$$

for an appropriate constant c_k'.

Proof From Eq. (1), the total number of different labelings of set C_δ by ζ is bounded by $2^{(c_1/(2\delta))^{k-1}}$ (that is, the number of binary functions over set of size $|C_\delta|$). For $\delta = \theta/2$, we have that the size is bounded by $2^{(c_1/\theta)^{k-1}}$. By Theorem 7 we conclude that

$$\mathcal{V}_T^{\theta}(\zeta) \leq 8\ 2^{\left(\frac{c_1}{\theta}\right)^{k-1}} \exp\left(-\frac{T\theta^2}{64\ c_k}\right)\ .$$

Over-bounding, we obtain the first statement. Now, set $\theta = c_k' T^{-1/(k+1)} \sqrt{\log(8/\eta)}$ for some appropriate constant c_k'. For this value of θ, it holds that $\mathcal{V}_T^{\theta}(\zeta) \leq \eta$. We conclude that

$$\mathbb{P}\left(\ \mathbf{R}_T\ \leq\ c_k'\ T^{-1/(k+1)} \sqrt{\log(8/\eta)}\ \right) \leq 1 - \eta\ .$$

■

While the rate for all measurable history-invariant checking rules decays with k, we can get $\tilde{O}(\sqrt{T})$ rates as soon as we restrict the class of checking rules to have a finite combinatorial dimension. A finite combinatorial dimension limits the effective size of ζ as applied on C_δ. The first result we present holds for Vapnik-Chervonenkis classes.

Lemma 9 *For any class ζ of history invariant checking rules with VC dimension* $\mathrm{VCdim}(\zeta)$, *we have that*

$$\mathcal{V}_T^{\theta}(\zeta) \leq 8 \ \left(\frac{e\ c_1}{\theta}\right)^{(k-1)\ \mathrm{VCdim}(\zeta)} \exp\left(-\frac{T\theta^2}{64\ c_k}\right)$$

for $T > \frac{16 c_k \log(4)}{\theta^2}$. We therefore obtain

$$\mathbb{P}\left(\mathbf{R}_T \ \leq \ c'\sqrt{\frac{k\mathrm{VCdim}(\zeta)\cdot c_k \log(8/\eta)\log T}{T}} \ \right) \leq 1-\eta$$

for an appropriate constant c'_k.

Proof By the Vapnik-Chervonenkis-Sauer-Shelah lemma, the number of different labelings of the set C_δ by ζ is bounded by $(e\ |C_\delta|)^{\mathrm{VCdim}(\zeta)}$. Clearly, the size of the minimal 0-cover cannot be more than the number of different labelings on the set C_δ. Using $|C_\delta| \leq (c_1/(2\delta))^{k-1}$ with $\delta = \theta/2$ and Theorem 7 we conclude that

$$\mathcal{V}_T^{\theta}(\zeta) \leq 8 \ \left(\frac{e\ c_1}{\theta}\right)^{(k-1)\ \mathrm{VCdim}(\zeta)} \exp\left(-\frac{T\theta^2}{64\ c_k}\right)$$

which concludes the first statement. For the probability version, set

$$\theta = c'\sqrt{\frac{k\mathrm{VCdim}(\zeta)\cdot c_k \log(8/\eta)\log T}{T}}$$

For this setting, $\mathcal{V}_T^{\theta}(\zeta) \leq \eta$ for some appropriate k-independent constant c'. The second statement follows. ■

For the classical calibration problem, the VC dimension of the set of ℓ_1-balls is at most k^2 and the constant $c_k = 8k$ for the ℓ_1 norm (as shown in Lemma 3). Combining, we obtain the following corollary, which, to the best of our knowledge, does not appear in the literature.

Corollary 10 *For classical calibration with k actions and ℓ_1 norm, the rate of convergence is*

$$O\left(k^2\sqrt{\frac{\log(T)\log(1/\eta)}{T}}\right)$$

Next, we consider an alternative combinatorial parameter, called Littlestone's dimension (Littlestone, 1988; Ben-David et al., 2009). This dimension captures the sequential "richness" of the function class.

Definition 11 *An $\mathcal{X}$-valued tree $\mathbf{x}$ of depth d is* shattered *by a function class $\mathcal{F} \subseteq \{\pm 1\}^{\mathcal{X}}$ if for all $\epsilon \in \{\pm 1\}^d$, there exists $f \in \mathcal{F}$ such that $f(\mathbf{x}_t(\epsilon)) = \epsilon_t$ for all $t \in [d]$. The* Littlestone dimension $\mathrm{Ldim}(\mathcal{F}, \mathcal{X})$ *is the largest d such that $\mathcal{F}$ shatters some $\mathcal{X}$-valued tree of depth d.*

We use $\mathrm{Ldim}(\mathcal{F})$ for $\mathrm{Ldim}(\mathcal{F}, \mathcal{X})$ if the domain $\mathcal{X}$ is clear from context. As shown in (Rakhlin et al., 2010a), the Littlestone's dimension can be used to upper bound sequential covering numbers in a way similar to VC dimension upper bounding the classical covering numbers.

Lemma 12 *For any class ζ of history invariant checking rules with Littlestone's dimension* $\mathrm{Ldim}(\zeta)$,

$$\mathcal{V}_T^\theta(\zeta) \leq 8 \ (eT)^{\mathrm{Ldim}(\zeta)} \exp\left(-\frac{T\theta^2}{64 \ c_k}\right)$$

Proof Note that for any history invariant family of checking rules ζ, the definition of covering number here coincides with the definition of covering number in (Rakhlin et al., 2010a) for binary class of functions ζ on space C_δ. Therefore,

$$\mathcal{N}_{\mathrm{ch}}(\zeta, T) \leq (eT)^{\mathrm{Ldim}(\zeta, C_\delta)} \ .$$

The Littlestone's dimension on the set C_δ can be upper bounded by the Littlestone's dimension $\mathrm{Ldim}(\zeta)$ over the whole simplex Δ_k. Using Theorem 7 concludes the proof. ■

In the above lemma and in the rest of the paper, it will be assumed that T is large enough that $T > \frac{16c_k \log(4)}{\theta^2}$ so that we can appeal to Theorem 7.

5.3. Time Dependent Checking Rules

We now turn to richer classes of checking rules. Of particular interest are classes of history-invariant rules that have mild dependence on time. Our results have a flavor of "shifting experts" results in individual sequence prediction. Suppose the checking rules can be written as a family of functions $c : [T] \times \Delta_k \mapsto \{0, 1\}$ (i.e. the checking rule only depends on the length of the history and not the history itself). More specifically, given a family ζ of time invariant checking rules, we consider the family of time dependent checking rules ζ^n given by checking rules that are allowed to change at most $n \leq T$ times over the T rounds (checking rule for each round is chosen from ζ). Formally,

$$\begin{aligned} \zeta^n = \{c^n | \exists \ 1 = i_0 \leq \ldots \leq i_n \leq T \text{ and } c_1, \ldots, c_n \in \zeta \quad \text{s.t.} \\ \forall \ s \geq 0, \forall \ i_s \leq t \leq t' < i_{s+1}, \ c^n(t, \cdot) = c^n(t', \cdot) = c_s\big\} \end{aligned}$$

and i_{n+1} is assumed to be $T + 1$.

Lemma 13 *For any class ζ of history invariant measurable checking rules, we have that*

$$\mathcal{V}_T^\theta(\zeta^n) \leq 8 \exp\left(-\frac{T\theta^2}{64 \ c_k} + n\left(\frac{c_1}{\theta}\right)^{k-1} + n \log T\right)$$

Proof For any t, the total number of different labelings of set C_δ by ζ is bounded by $2^{(c_1/(2\delta))^{k-1}}$. To account for all the possibilities, we need to consider all possible ways of choosing n shifts out of T rounds, and then to choose a constant function for each interval out of the $2^{(c_1/(2\delta))^{k-1}}$ possibilities. Choosing $\delta = \theta/2$, the effective size of ζ on C_δ is bounded by $\binom{T}{n}\left(2^{(c_1/\theta)^{k-1}}\right)^n$. Hence by Theorem 7 we conclude that

$$\mathcal{V}_T^\theta(\zeta) \le 8\binom{T}{n} 2^{n\left(\frac{c_1}{\theta}\right)^{k-1}} \exp\left(-\frac{T\theta^2}{64\ c_k}\right)$$

which concludes the proof. ■

The corresponding statement in probability is analogous to that in Lemma 8 if n is constant. If n grows with T, a non-trivial rate in probability can still be shown as long as $n = o(T)$. Hence, there exists a calibration strategy for arbitrary sets of history-independent measurable checking rules which change $o(T)$ of times.

Lemma 14 *For any class ζ of history invariant checking rules with VC dimension* $\mathrm{VCdim}(\zeta)$,

$$\mathcal{V}_T^\theta(\zeta^n) \le 8\binom{T}{n}\left(\frac{e\ c_1}{\theta}\right)^{n(k-1)\mathrm{VCdim}(\zeta)} \exp\left(-\frac{T\theta^2}{64\ c_k}\right)$$

Proof For any $t \in [T]$ the number of different labelings of the set C_δ by ζ is bounded by $(e\ |C_\delta|)^{\mathrm{VCdim}(\zeta)}$. Hence the total possible number of different labelings of set C_δ by ζ in the T different rounds can be bounded by $\binom{T}{n}(e\ |C_\delta|)^{n\mathrm{VCdim}(\zeta)} \le \binom{T}{n}\left(\frac{e\ c_1}{\theta}\right)^{n(k-1)\ \mathrm{VCdim}(\zeta)}$. By Theorem 7 we conclude that

$$\mathcal{V}_T^\theta(\zeta) \le 8\binom{T}{n}\left(\frac{e\ c_1}{\theta}\right)^{n(k-1)\mathrm{VCdim}(\zeta)} \exp\left(-\frac{T\theta^2}{64\ c_k}\right)$$

which concludes the proof. ■

Similarly to Lemma 9, we obtain $\tilde{O}(\sqrt{T})$ rate of convergence for the class ζ^n constructed from a VC class of history-independent checking rules.

5.4. General Checking Rules

In this section we study checking rules that depend on history. We start with an assumption on the form of these rules: history is represented by some potentially smaller set. Such a smaller set can arise from a bound on the available memory, or from limited precision.

Formally, assume that for some set $\mathcal{Y}$ there exists a mapping $\phi : \mathcal{Z}^\delta \mapsto \mathcal{Y}$ and a class of binary functions $\mathcal{G} \subseteq \{0,1\}^{\mathcal{Y}\times\Delta_k}$ with the following property: For any $c \in \zeta$ there exists $g \in \mathcal{G}$ such that

$$c(z,p) = g(\phi(z),p) \quad \text{ for any } z \in \mathcal{Z} \text{ and } p \in \Delta_k\ .$$

Clearly, if we set $\mathcal{Y} = \mathcal{Z}^\delta$ and ϕ the identity mapping, $\mathcal{G}$ and ζ coincide.

Lemma 15 *For any set $\mathcal{Y}$ and class of binary functions $\mathcal{G}$ satisfying the above mentioned assumption with mapping ϕ, we have that*

$$\mathcal{V}_T^\theta \le 8\,(eT)^{\mathrm{Ldim}(\mathcal{G})}\ \exp\left(-\frac{T\theta^2}{64 c_k}\right)$$

Proof Note that

$$\mathcal{N}_{\mathrm{ch}}(\zeta,T) = \sup_{\mathbf{x},\mathbf{p}^\delta} N(\zeta,(\mathbf{z}^{(\mathbf{x},\mathbf{p}^\delta)},\mathbf{p}^\delta)) = \sup_{\mathbf{x},\mathbf{p}^\delta} N(\mathcal{G},(\phi(\mathbf{z}^{(\mathbf{x},\mathbf{p}^\delta)}),\mathbf{p}^\delta)) \le \sup_{\mathbf{y},\mathbf{p}^\delta} N(\mathcal{G},(\mathbf{y},\mathbf{p}^\delta)) \le (eT)^{\mathrm{Ldim}(\mathcal{G})} \ .$$

Using this with Theorem 7 we conclude the proof. ■

Corollary 16 *For any class of checking rules ζ,*

$$\mathcal{V}_T^\theta \le 8\,(eT)^{\mathrm{Ldim}(\zeta,\mathcal{Z}^\delta\times C_\delta)} \ \exp\left(-\frac{T\theta^2}{64c_k}\right)$$

Proof Use previous lemma with $\mathcal{G}=\zeta$, $\mathcal{Y}=\mathcal{Z}^\delta$ and ϕ the identity mapping. ■

5.5. Checking Rules With Limited History Lookback

We now consider a family of checking rules that only depend on at most m of the most recent pairs of actions played by the two players. We call such a class of rules an *m-look back family*. Specifically, for $0 \le m \le T-1$, define $\mathcal{Y} = \bigcup_{t=0}^m (C_\delta \times E_k)^t \subset \mathcal{Z}^\delta$, $\mathcal{G}=\zeta$ and $\phi : \mathcal{Z}^\delta \mapsto \mathcal{Y}$ is given by:

$$\phi(z) = \begin{cases} z & \text{if } z \in \mathcal{Y} \\ (z_{t-m-1},\ldots,z_t) & \text{if } z \in (C_\delta \times E_k)^t \text{ for some } m < t \le T \end{cases}$$

The first bound we can get here directly is the one implied by Lemma 15 for the $\mathcal{G}$ and $\mathcal{Y}$ mentioned above.

Lemma 17 *For any m-look back family of checking rules ζ,*

$$\mathcal{V}_T^\theta \le 8 \cdot 2^{m\ k^m \left(\frac{c_1}{\theta}\right)^{km}} \ \exp\left(-\frac{T\theta^2}{64c_k}\right)$$

Proof Note that

$$|\mathcal{Y}| = \sum_{t=0}^m \left|(C_\delta \times E_k)^t\right| \le \sum_{t=0}^m (|C_\delta|\cdot k)^t \le \sum_{t=0}^m \left(\left(\frac{c_1}{2\delta}\right)^{(k-1)} \cdot k\right)^t \le m\ k^m \left(\frac{c_1}{2\delta}\right)^{(k-1)m}$$

So for $\delta = \theta/2$ we have $|\mathcal{Y}| \le mk^m\left(\frac{c_1}{\theta}\right)^{(k-1)m}$. This implies that the total number of different possible binary labelings of elements of the set $\mathcal{Y}\times C_\delta$ (and hence $\mathcal{N}_{\mathrm{ch}}(\zeta,T)$) is bounded by

$$\mathcal{N}_{\mathrm{ch}}(\zeta,T) \le 2^{m\ k^m \left(\frac{c_1}{\theta}\right)^{km}}$$

Hence using Theorem 7 we conclude the theorem statement. ■

Note that the above bound gives polynomial convergence for any $m \le \frac{\log\ T}{1+\epsilon}$ for any $\epsilon > 0$. That is, there exists a forecasting strategy that can calibrate against any family of measurable checking rules which have dependence on a logarithmic (in T) number of past forecasts and outcomes.

Lemma 18 *For any m-look back family of checking rules ζ, if VC dimension of the class as applied on input space $\mathcal{Y} \times C_\delta$ is given by* $\mathrm{VCdim}(\zeta, \mathcal{Y} \times C_\delta)$ *then,*

$$\mathcal{V}_T^\theta \leq 2 \left(e \ m \ k^m \left(\frac{c_1}{\theta} \right)^{km} \right)^{\mathrm{VCdim}(\zeta, \mathcal{Y} \times C_\delta)} \exp\left(-\frac{T\theta^2}{64 c_k} \right)$$

Proof By VC lemma the number of different labelings of the set $\mathcal{Y} \times C_\delta$ by the class ζ is bounded by $(e|\mathcal{Y} \times C_\delta|)^{\mathrm{VCdim}(\zeta, \mathcal{Y} \times C_\delta)}$. However

$$|\mathcal{Y} \times C_\delta| \leq m \ k^m \left(\frac{c_1}{\theta} \right)^{km}$$

Hence

$$\mathcal{N}(\zeta, T) \leq \left(e \ m \ k^m \left(\frac{c_1}{\theta} \right)^{km} \right)^{\mathrm{VCdim}(\zeta, \mathcal{Y} \times C_\delta)}$$

We conclude the proof by appealing to Theorem 7. ■

The above bound guarantees existence of a calibration strategy whenever $m = o(T)$. That is, as long as the checking rule with bounded VC only looks back up to $o(T)$ steps in history, the forecaster has a successful strategy.

5.6. Checking Rules with Bounded Computation

Whenever the number of arithmetic operations required to compute each function in a class is bounded by some constant, the VC dimension of the class can be bounded from above Goldberg and Jerrum (1995). Specifically result in Goldberg and Jerrum (1995) states that for binary function class ζ over domain $\mathcal{X} \subset \mathbb{R}^n$ defined by algorithms of description length bounded by ℓ and which run in time U using only the operations of conditional jumps and $+$, $-$, $\times$ and $/$ (in constant time), the VC dimension of the function class is bounded by $O(\ell U)$. Using this with Lemma 18 we make the following observation.

For m-look back family of checking rules ζ defined by algorithms with description length bounded by ℓ and runtime bounded by U, applying Lemma 18, the value of the game is bounded by

$$\mathcal{V}_T^\theta \leq 2 \left(e \ m \ k \left(\frac{c_1}{\theta} \right)^{k} \right)^{O(m\ell U)} \exp\left(-\frac{T\theta^2}{64 c_k} \right)$$

Hence we can gaurantee calibration against set of all checking rules defined by algorithms of description length bounded by ℓ and whose run times are bounded by U as long as $m\ell U = o(T)$.

6. Lower Bounds

In this section we show that the $\sqrt{T}$ rate for classical calibration cannot be improved. While the argument is not difficult, we could not find it in the literature.

Lemma 19 *For two actions, the rate for the classical calibration game is lower bounded for any $\theta > 0$ as*

$$\mathcal{V}_T^\theta \geq \mathbb{P}\left(\frac{1}{T}\sum_{t=1}^T x_t \geq 2\theta\right)$$

where $x_1, \ldots, x_T$ are independent Rademacher random variables.

Proof Note that for $k = 2$, the vector notation for the outcomes is no longer necessary. Indeed, the difference of any two vectors in the simplex is $|(a, 1-a) - (b, 1-b)| = 2|a-b|$, and thus the value of the game can be written as

$$\mathcal{V}_T^\theta(\zeta) := \inf_{q_1}\sup_{x_1}\mathbb{E}_{f_1\sim q_1}\ldots\inf_{q_T}\sup_{x_T}\mathbb{E}_{f_T\sim q_T}\left[\mathbf{1}\left\{\sup_{c\in\zeta}\left|\frac{1}{T}\sum_{t=1}^T c(z_t, f_t)\cdot(f_t - x_t)\right| > \theta\right\}\right]$$

where q_t is a distribution over $[0,1]$, $f_t \in [0,1]$, and $x_t \in \{0,1\}$. In fact, the mathematical exposition is easier if q_t is a distribution on $[-1,1]$, $f_t \in [-1,1]$, and $x_t \in \{-1,1\}$. The problem is not changed, as one can easily translate between the two formulations. We consider a particular ζ consisting of two rules: $c_1(z_t, f_t) = \mathbf{1}\{f_t \geq 0\}$ and $c_2(z_t, f_t) = \mathbf{1}\{f_t < 0\}$. Note that we can equivalently write these rules as being 1/4-close to the centers 1/4 and 3/4. Hence, this is genuinely a classical ϵ-calibration problem with $\epsilon = 1/4$. We can then write the value of the game as

$$\inf_{q_1}\sup_{x_1}\mathbb{E}_{f_1\sim q_1}\ldots\inf_{q_T}\sup_{x_T}\mathbb{E}_{f_T\sim q_T}$$
$$\mathbf{1}\left\{\max\left\{\left|\frac{1}{T}\sum_{t=1}^T(x_t - f_t)\mathbf{1}\{f_t \geq 0\}\right|, \left|\frac{1}{T}\sum_{t=1}^T(x_t - f_t)\mathbf{1}\{f_t < 0\}\right|\right\} > \theta\right\}$$

Let $\mathrm{sign}(b)$ denote the sign of $b \in \mathbb{R}$, and $\mathrm{sign}(0) = 1$. Let us write

$$A(f_{1:T}, x_{1:T}) := \frac{1}{T}\sum_{t=1}^T(x_t - f_t)\mathbf{1}\{f_t \geq 0\} \quad \text{and} \quad B(f_{1:T}, x_{1:T}) := \frac{1}{T}\sum_{t=1}^T(x_t - f_t)\mathbf{1}\{f_t < 0\}.$$

The suprema over x_t's can equivalently be written as suprema over all distributions on $\{-1,1\}$. The lower bound is then achieved by choosing x_t to be i.i.d. Rademacher random variables. The lower bound on the value of the game can thus be written as

$$\begin{aligned}
\mathcal{V}_T^\theta &\geq \inf_{q_1}\mathbb{E}_{f_1\sim q_1}\mathbb{E}_{x_1}\ldots\inf_{q_T}\mathbb{E}_{f_T\sim q_T}\mathbb{E}_{x_T}\left[\mathbf{1}\{\max\{|A(f_{1:T}, x_{1:T})|, |B(f_{1:T}, x_{1:T})|\} > \theta\}\right] \\
&= \inf_{f_1}\mathbb{E}_{x_1}\ldots\inf_{f_T}\mathbb{E}_{x_T}\left[\mathbf{1}\{\max\{|A(f_{1:T}, x_{1:T})|, |B(f_{1:T}, x_{1:T})|\} > \theta\}\right] \\
&= \inf_{f_1}\sup_{a_1\in\{\pm 1\}}\mathbb{E}_{x_1}\ldots\inf_{f_T}\sup_{a_T\in\{\pm 1\}}\mathbb{E}_{x_T}\left[\mathbf{1}\left\{\max\left\{|A(f_{1:T}, \{a_t x_t\}_{t=1}^T)|, |B(f_{1:T}, \{a_t x_t\}_{t=1}^T)|\right\} > \theta\right\}\right]
\end{aligned}$$

The last equality holds because x_t have the same distribution as $a_t x_t$. Now, choosing $a_t = \mathrm{sign}(f_t)$, we get

$$\begin{aligned}
\mathcal{V}_T^\theta &\geq \inf_{f_1}\mathbb{E}_{x_1}\ldots\inf_{f_T}\mathbb{E}_{x_T}\left[\mathbf{1}\left\{\max\left\{|A(f_{1:T}, \{\mathrm{sign}(f_t)x_t\}_{t=1}^T)|, |B(f_{1:T}, \{\mathrm{sign}(f_t)x_t\}_{t=1}^T)|\right\} > \theta\right\}\right] \\
&= \inf_{f_1}\mathbb{E}_{x_1}\ldots\inf_{f_T}\mathbb{E}_{x_T}\left[\mathbf{1}\{\max\{|A(f_{1:T}, x_{1:T})|, |B(f_{1:T}, -x_{1:T})|\} > \theta\}\right].
\end{aligned}$$

Observe that

$$\begin{aligned} A(f_{1:T}, x_{1:T}) - B(f_{1:T}, -x_{1:T}) &= \frac{1}{T}\sum_{t=1}^{T}(x_t - f_t)\mathbf{1}\{f_t \geq 0\} - \frac{1}{T}\sum_{t=1}^{T}(-x_t - f_t)\mathbf{1}\{f_t < 0\} \\ &= \frac{1}{T}\sum_{t=1}^{T} x_t - \frac{1}{T}\sum_{t=1}^{T} f_t\mathbf{1}\{f_t \geq 0\} + \frac{1}{T}\sum_{t=1}^{T} f_t\mathbf{1}\{f_t < 0\} \\ &\leq \frac{1}{T}\sum_{t=1}^{T} x_t \ . \end{aligned}$$

Hence,

$$\mathbf{1}\{\max\{|A(f_{1:T}, x_{1:T})|, |B(f_{1:T}, -x_{1:T})|\} > \theta\} > \mathbf{1}\left\{\frac{1}{T}\sum_{t=1}^{T} x_t < -2\theta\right\} \ .$$

We conclude

$$\mathcal{V}_T^\theta \geq \mathbb{P}\left(\frac{1}{T}\sum_{t=1}^{T} x_t < -2\theta\right) \ .$$

■

The lower bound of Lemma 19 can be immediately extended to $k > 2$ actions and history-invariant checking rules that change $O(k)$ times. This can be done by dividing T rounds into $\lfloor k/2 \rfloor$ equal-length periods and then constructing the lower bound for each period based on two actions.

7. Proofs

Proof [**Lemma 2**] The first step is replacing the suprema over x_t with suprema over distributions p_t on E_k. The second step is exchanging each infimum and supremum by appealing to the minimax theorem.

$$\begin{aligned} \mathcal{V}_T^\theta(\zeta) &= \inf_{q_1}\sup_{p_1} \mathop{\mathbb{E}}_{\substack{f_1\sim q_1 \\ x_1\sim p_1}} \ldots \inf_{q_T}\sup_{p_T} \mathop{\mathbb{E}}_{\substack{f_T\sim q_T \\ x_T\sim p_T}} \left[\mathbf{1}\left\{\sup_{c\in\zeta}\left\|\frac{1}{T}\sum_{t=1}^{T} c(z_t, f_t)\cdot(f_t - x_t)\right\| > \theta\right\}\right] \\ &= \sup_{p_1}\inf_{q_1} \mathop{\mathbb{E}}_{\substack{f_1\sim q_1 \\ x_1\sim p_1}} \ldots \sup_{p_T}\inf_{q_T} \mathop{\mathbb{E}}_{\substack{f_T\sim q_T \\ x_T\sim p_T}} \left[\mathbf{1}\left\{\sup_{c\in\zeta}\left\|\frac{1}{T}\sum_{t=1}^{T} c(z_t, f_t)\cdot(f_t - x_t)\right\| > \theta\right\}\right] \\ &= \sup_{p_1}\inf_{f_1\in\Delta_k} \mathbb{E}_{x_1\sim p_1} \ldots \sup_{p_T}\inf_{f_T\in\Delta_k} \mathbb{E}_{x_T\sim p_T} \left[\mathbf{1}\left\{\sup_{c\in\zeta}\left\|\frac{1}{T}\sum_{t=1}^{T} c(z_t, f_t)\cdot(f_t - x_t)\right\| > \theta\right\}\right] \end{aligned}$$

Now since $C_\delta \subset \Delta_k$ we have

$$\begin{aligned}
\mathcal{V}_T^\theta(\zeta) &= \sup_{p_1} \inf_{f_1 \in \Delta_k} \mathbb{E}_{x_1 \sim p_1} \dots \sup_{p_T} \inf_{f_T \in \Delta_k} \mathbb{E}_{x_T \sim p_T} \left[\mathbf{1}\left\{ \sup_{c \in \zeta} \left\| \frac{1}{T} \sum_{t=1}^T c(z_t, f_t) \cdot (f_t - x_t) \right\| > \theta \right\} \right] \\
&\le \sup_{p_1} \inf_{f_1 \in C_\delta} \mathbb{E}_{x_1 \sim p_1} \dots \sup_{p_T} \inf_{f_T \in C_\delta} \mathbb{E}_{x_T \sim p_T} \left[\mathbf{1}\left\{ \sup_{c \in \zeta} \left\| \frac{1}{T} \sum_{t=1}^T c(z_t, f_t) \cdot (f_t - x_t) \right\| > \theta \right\} \right] \\
&\le \sup_{p_1} \mathbb{E}_{x_1 \sim p_1} \dots \sup_{p_T} \mathbb{E}_{x_T \sim p_T} \left[\mathbf{1}\left\{ \sup_{c \in \zeta} \left\| \frac{1}{T} \sum_{t=1}^T c(z_t^\delta, p_t^\delta) \cdot (p_t^\delta - x_t) \right\| > \theta \right\} \right] \qquad (6)
\end{aligned}$$

where the last inequality is obtained by replacing each $\inf_{f_t \in C_\delta}$ by the (possibly) sub-optimal choice of p_t^δ, thus only increasing the value.

By triangle inequality

$$\left\| \frac{1}{T} \sum_{t=1}^T c(z_t, p_t^\delta) \cdot (p_t^\delta - x_t) \right\| \le \left\| \frac{1}{T} \sum_{t=1}^T c(z_t^\delta, p_t^\delta) \cdot (p_t^\delta - p_t) \right\| + \left\| \frac{1}{T} \sum_{t=1}^T c(z_t^\delta, p_t^\delta) \cdot (p_t - x_t) \right\|$$

and the first term above is further bounded above by

$$\frac{1}{T} \sum_{t=1}^T \left\| c(z_t^\delta, p_t^\delta) \cdot (p_t^\delta - p_t) \right\| \le \frac{1}{T} \sum_{t=1}^T \left\| p_t^\delta - p_t \right\| \le \delta \ .$$

Using this in Equation 6, we get

$$\begin{aligned}
\mathcal{V}_T^\theta(\zeta) &\le \sup_{p_1} \mathbb{E}_{x_1 \sim p_1} \dots \sup_{p_T} \mathbb{E}_{x_T \sim p_T} \left[\mathbf{1}\left\{ \delta + \sup_{c \in \zeta} \left\| \frac{1}{T} \sum_{t=1}^T c(z_t^\delta, p_t^\delta) \cdot (p_t - x_t) \right\| > \theta \right\} \right] \\
&\le \mathbf{1}\left\{ \delta > \theta/2 \right\} + \sup_{p_1} \mathbb{E}_{x_1 \sim p_1} \dots \sup_{p_T} \mathbb{E}_{x_T \sim p_T} \left[\mathbf{1}\left\{ \sup_{c \in \zeta} \left\| \frac{1}{T} \sum_{t=1}^T c(z_t^\delta, p_t^\delta) \cdot (p_t - x_t) \right\| > \theta/2 \right\} \right]
\end{aligned}$$

Choosing $\delta \le \theta/2$ concludes the proof. ∎

Proof [Lemma 3]

The result is a straightforward consequence of concentration results for 2-smooth functions of an average of a martingale difference sequence due to Pinelis (1994). We also refer to (Rakhlin et al., 2010b) for a short but detailed proof. The result states that, for a 2-smooth norm (in particular, $\|\cdot\|_2$),

$$\mathbb{P}\left(\left\| \frac{1}{T} \sum_{t=1}^T d_t \right\|_2 \ge \epsilon \right) \le 2 \exp\left(-\frac{\epsilon^2 T}{8B^2} \right)$$

if $\|d_t\|_2 \le B$ almost surely for all t. It remains to pass from our norm $\|\cdot\|$ to the ℓ_2 norm. Here, we make this transition explicit for any ℓ_p norm ($1 \le p \le \infty$), but it can also be done for any appropriately normalized norm on $\mathbb{R}^k$.

For $p \leq 2$, $\|\cdot\|_2 \leq \|\cdot\|_p$ and thus the condition $\|d_t\|_p \leq 1$ implies $\|d_t\|_2 \leq 1$. Further, $\|\cdot\|_p \leq \sqrt{k}\|\cdot\|_2$ and so $\|\cdot\|_p \geq \epsilon$ implies $\|\cdot\|_2 \geq \epsilon/\sqrt{k}$. Thus, $c_k = 8k$. Now, for the case $p \geq 2$, $\|\cdot\|_2 \leq \sqrt{k}\|\cdot\|_p$ and thus we set $B = \sqrt{k}$, leading to the value $c_k = 8k$.

■

Acknowledgments

A. Rakhlin gratefully acknowledges the support of NSF under grant CAREER DMS-0954737 and Dean's Research Fund.

References

J. Abernethy, P. L. Bartlett, and E. Hazan. Blackwell approachability and no-regret learning are equivalent. In *Proceedings of the 24th Annual Conference on Learning Theory*, 2011.

S. Ben-David, D. Pal, and S. Shalev-Shwartz. Agnostic online learning. In *Proceedings of the 22th Annual Conference on Learning Theory*, 2009.

N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games.* Cambridge University Press, 2006.

A. P. Dawid. The well-calibrated Bayesian. *Journal of the American Statistical Association*, 77(379):605–610, 1982.

V. de la Peña and E. Giné. *Decoupling: From Dependence to Independence.* Springer, 1998.

D. P. Foster and R. V. Vohra. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21(1-2):40–55, October 1997.

D. P. Foster and R. V. Vohra. Asymptotic calibration. *Biometrika*, 85(2):379, 1998.

P.W. Goldberg and M.R. Jerrum. Bounding the Vapnik-Chervonenkis dimension of concept classes parameterized by real numbers. *Machine Learning*, 18(2):131–148, 1995. ISSN 0885-6125.

S. Kakade, M. Kearns, J. Langford, and L. Ortiz. Correlated equilibria in graphical games. In *Proceedings of the 4th ACM Conference on Electronic Commerce*, pages 42–47. ACM, 2003. ISBN 158113679X.

S. M. Kakade and D. P. Foster. Deterministic calibration and nash equilibrium. *Journal of Computer and System Sciences*, 74(1):115 – 130, 2008. ISSN 0022-0000. Learning Theory 2004.

E. Kalai, E. Lehrer, and R. Smorodinsky. Calibrated Forecasting and Merging. *Games and Economic Behavior*, 29(1-2):151–169, 1999. ISSN 0899-8256.

E. Lehrer. The game of normal numbers. *Mathematics of Operations Research*, 29(2): 259–265, 2004. ISSN 0364-765X.

N. Littlestone. Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm. *Machine Learning*, 2(4):285–318, 04 1988.

S. Mannor and G. Stoltz. A geometric proof of calibration. *Mathematics of Operations Research*, 35(4):721–727, 2010. doi: 10.1287/moor.1100.0465.

D. Oakes. Self-calibrating priors do not exist. *Journal of the American Statistical Association*, 80(390):339–339, 1985. ISSN 0162-1459.

W. Olszewski and A. Sandroni. A nonmanipulable test. *Annals of Statistics*, 37(2):1013–1039, 2009.

I. Pinelis. Optimum bounds for the distributions of martingales in Banach spaces. *The Annals of Probability*, 22(4):1679–1706, 1994.

A. Rakhlin, K. Sridharan, and A. Tewari. Online learning: Random averages, combinatorial parameters, and learnability. In *NIPS*, 2010a. Full version available at arXiv:1006.1138.

A. Rakhlin, K. Sridharan, and A. Tewari. Online learning: Beyond regret. *ArXiv preprint arXiv:1011.3168v2*, 2010b.

A. Sandroni, R. Smorodinsky, and R.V. Vohra. Calibration with many checking rules. *Mathematics of Operations Research*, 28(1):141–153, 2003. ISSN 0364-765X.

H.P. Young. *Strategic learning and its limits.* Oxford University Press, USA, 2004. ISBN 0199269181.

Appendix

Proof [**Lemma 4**] Fix a $\mathbf{p}$. If we condition on $x_1, \ldots, x_T$, the sequence of $p_1, \ldots, p_T$ is well-defined, and we can consider a *tangent* sequence $x'_t \sim p_t$. This sequence is independent (see (de la Peña and Giné, 1998; Rakhlin et al., 2010a)). Note also that for any t, $c(z_t^\delta, p_t^\delta)$ is constant given $x_1, \ldots, x_T$. Then for any fixed $c \in \zeta$,

$$\mathbb{E}_{x'_1 \sim p_1, \ldots, x'_T \sim p_T}\left[\mathbf{1}\left\{\left\|\frac{1}{T}\sum_{t=1}^T c(z_t^\delta, p_t^\delta)\cdot(p_t - x'_t)\right\| > \theta/4\right\}\middle| x_1, \ldots, x_T\right]$$
$$= \mathbb{P}\left(\left\|\frac{1}{T}\sum_{t=1}^T c(z_t^\delta, p_t^\delta)\cdot(p_t - x'_t)\right\| > \theta/4 \;\middle|\; x_1, \ldots, x_T\right) \le 2\exp\left(-\frac{T\theta^2}{16c_k}\right) \le \frac{1}{2}$$

where the last inequality is by our assumption that $T > \frac{16c_k \log(4)}{\theta^2}$. Hence we can conclude that for any fixed $c \in \zeta$,

$$\mathbb{P}\left(\left\|\frac{1}{T}\sum_{t=1}^T c(z_t^\delta, p_t^\delta)\cdot(p_t - x'_t)\right\| \le \theta/4 \;\middle|\; x_1, \ldots, x_T\right) \ge \frac{1}{2}$$

Now since we are conditioning on $x_1, \ldots, x_T$ we can pick $c^* \in \zeta$ as :

$$c^* = \operatorname*{argmax}_{c \in \zeta} \left\|\frac{1}{T}\sum_{t=1}^T c(z_t^\delta, p_t^\delta)\cdot(p_t - x_t)\right\|$$

and so

$$\mathbb{P}\left(\left\|\frac{1}{T}\sum_{t=1}^T c^*(z_t^\delta, p_t^\delta)\cdot(p_t - x'_t)\right\| \le \theta/4 \;\middle|\; x_1, \ldots, x_T\right) \ge \frac{1}{2} \tag{7}$$

Since the Inequality (7) holds for any $x_1, \ldots, x_T$ we assert that

$$\frac{1}{2} \le \mathbb{P}\left(\left\|\frac{1}{T}\sum_{t=1}^T c^*(z_t^\delta, p_t^\delta)\cdot(p_t - x'_t)\right\| \le \theta/4 \;\middle|\; \sup_{c\in\zeta}\left\|\frac{1}{T}\sum_{t=1}^T c(z_t^\delta, p_t^\delta)\cdot(p_t - x_t)\right\| > \theta/2\right)$$

Hence we can conclude that for any distribution,

$$\frac{1}{2}\mathbb{P}\left(\sup_{c\in\zeta}\left\|\frac{1}{T}\sum_{t=1}^T c(z_t^\delta, p_t^\delta)\cdot(p_t - x_t)\right\| > \theta/2\right)$$
$$\le \mathbb{P}\left(\left\|\frac{1}{T}\sum_{t=1}^T c^*(z_t^\delta, p_t^\delta)\cdot(p_t - x'_t)\right\| \le \theta/4 \;\middle|\; \sup_{c\in\zeta}\left\|\frac{1}{T}\sum_{t=1}^T c(z_t^\delta, p_t^\delta)\cdot(p_t - x_t)\right\| > \theta/2\right)$$
$$\times\, \mathbb{P}\left(\sup_{c\in\zeta}\left\|\frac{1}{T}\sum_{t=1}^T c(z_t^\delta, p_t^\delta)\cdot(p_t - x_t)\right\| > \theta/2\right)$$
$$= \mathbb{P}\left(\sup_{c\in\zeta}\left\|\frac{1}{T}\sum_{t=1}^T c(z_t^\delta, p_t^\delta)\cdot(p_t - x_t)\right\| > \theta/2\ ,\ \left\|\frac{1}{T}\sum_{t=1}^T c^*(z_t^\delta, p_t^\delta)\cdot(p_t - x'_t)\right\| \le \theta/4\right)$$
$$\le \mathbb{P}\left(\sup_{c\in\zeta}\left\|\frac{1}{T}\sum_{t=1}^T c(z_t^\delta, p_t^\delta)\cdot(x_t - x'_t)\right\| > \theta/4\right)$$

Note that the probability is both with respect to the stochastic process $x_1, \ldots, x_T$ and the tangent sequence $x'_1, \ldots, x'_T$. Furthermore, the above inequality holds for any $\mathbf{p}$. Thus,

$$\begin{aligned}
&\frac{1}{2} \sup_{\mathbf{p}} \mathbb{E}_{x_1,\ldots,x_T} \left[\mathbf{1}\left\{ \sup_{c\in\zeta} \left\| \frac{1}{T} \sum_{t=1}^{T} c(z_t^\delta, p_t^\delta)\cdot(p_t - x_t) \right\| > \theta/2 \right\} \right] \\
&\le \sup_{\mathbf{p}} \mathbb{E}_{x_1,\ldots,x_T} \mathbb{E}_{x'_1,\ldots,x'_T} \left[\mathbf{1}\left\{ \sup_{c\in\zeta} \left\| \frac{1}{T} \sum_{t=1}^{T} c(z_t^\delta, p_t^\delta)\cdot(x_t - x'_t) \right\| > \theta/4 \right\} \right]
\end{aligned}$$

Moving back to the expanded notation of (2) and using Lemma 2,

$$\begin{aligned}
\frac{1}{2}\mathcal{V}_T^\theta &\le \sup_{p_1} \mathbb{E}_{x_1\sim p_1} \ldots \sup_{p_T} \mathbb{E}_{x_T\sim p_T} \mathbb{E}_{x'_1\sim p_1,\ldots,x'_T\sim p_T} \left[\mathbf{1}\left\{ \sup_{c\in\zeta} \left\| \frac{1}{T} \sum_{t=1}^{T} c(z_t^\delta, p_t^\delta)\cdot(x_t - x'_t) \right\| > \theta/4 \right\} \right] \\
&\le \sup_{p_1} \mathbb{E}_{x_1,x'_1\sim p_1} \ldots \sup_{p_T} \mathbb{E}_{x_T,x'_T\sim p_T} \left[\mathbf{1}\left\{ \sup_{c\in\zeta} \left\| \frac{1}{T} \sum_{t=1}^{T} c(z_t^\delta, p_t^\delta)\cdot(x_t - x'_t) \right\| > \theta/4 \right\} \right]
\end{aligned}$$

Next, we upper bound the above expression by introducing suprema over p_t^δ (we are slightly abusing the notation, as these variables will no longer depend on p_t):

$$\begin{aligned}
&\sup_{p_1} \sup_{p_1^\delta\in C_\delta} \mathbb{E}_{x_1,x'_1\sim p_1} \ldots \sup_{p_T} \sup_{p_T^\delta\in C_\delta} \mathbb{E}_{x_T,x'_T\sim p_T} \left[\mathbf{1}\left\{ \sup_{c\in\zeta} \left\| \frac{1}{T} \sum_{t=1}^{T} c(z_t^\delta, p_t^\delta)\cdot(x_t - x'_t) \right\| > \theta/4 \right\} \right] \\
&= \sup_{p_1} \sup_{p_1^\delta\in C_\delta} \mathop{\mathbb{E}}_{x_1,x'_1\sim p_1} \ldots \sup_{p_T} \sup_{p_T^\delta\in C_\delta} \mathop{\mathbb{E}}_{x_T,x'_T\sim p_T} \mathop{\mathbb{E}}_{\epsilon_T} \\
&\qquad \left[\mathbf{1}\left\{ \sup_{c\in\zeta} \left\| \frac{1}{T} \sum_{t=1}^{T-1} c(z_t^\delta, p_t^\delta)\cdot(x_t - x'_t) + \epsilon_T\ c(z_T^\delta, p_T^\delta)\cdot(x_T - x'_T) \right\| > \theta/4 \right\} \right]
\end{aligned}$$

The last step is justified because x_T and x'_T have the same distribution p_t when conditioned on $x_1, \ldots, x_{T-1}$, and thus we can introduce the Rademacher random variable ϵ_T. Next, we pass to the supremum over (x_T, x'_T):

$$\begin{aligned}
&\sup_{p_1} \sup_{p_1^\delta\in C_\delta} \mathop{\mathbb{E}}_{x_1,x'_1\sim p_1} \ldots \sup_{x_T,x'_T\in E_k} \sup_{p_T^\delta\in C_\delta} \mathop{\mathbb{E}}_{\epsilon_T} \\
&\qquad \left[\mathbf{1}\left\{ \sup_{c\in\zeta} \left\| \frac{1}{T} \sum_{t=1}^{T-1} c(z_t^\delta, p_t^\delta)\cdot(x_t - x'_t) + \epsilon_T\ c(z_T^\delta, p_T^\delta)\cdot(x_T - x'_T) \right\| > \theta/4 \right\} \right] \\
&= \sup_{p_1} \sup_{p_1^\delta\in C_\delta} \mathop{\mathbb{E}}_{x_1,x'_1\sim p_1} \ldots \sup_{p_{T-1}} \sup_{p_{T-1}^\delta\in C_\delta} \mathop{\mathbb{E}}_{x_{T-1},x'_{T-1}\sim p_{T-1}} \mathop{\mathbb{E}}_{\epsilon_{T-1}} \sup_{x_T,x'_T\in E_k} \sup_{p_T^\delta\in C_\delta} \mathop{\mathbb{E}}_{\epsilon_T} \\
&\qquad \left[\mathbf{1}\left\{ \sup_{c\in\zeta} \left\| \frac{1}{T} \sum_{t=1}^{T-2} c(z_t^\delta, p_t^\delta)\cdot(x_t - x'_t) + \sum_{j=T-1}^{T} \epsilon_j\ c(z_j^\delta, p_j^\delta)\cdot(x_j - x'_j) \right\| > \theta/4 \right\} \right] \\
&\le \sup_{p_1} \sup_{p_1^\delta\in C_\delta} \mathop{\mathbb{E}}_{x_1,x'_1\sim p_1} \ldots \sup_{x_{T-1},x'_{T-1}\in E_k} \sup_{p_{T-1}^\delta\in C_\delta} \mathop{\mathbb{E}}_{\epsilon_{T-1}} \sup_{x_T,x'_T\in E_k} \sup_{p_T^\delta\in C_\delta} \mathop{\mathbb{E}}_{\epsilon_T} \\
&\qquad \left[\mathbf{1}\left\{ \sup_{c\in\zeta} \left\| \frac{1}{T} \sum_{t=1}^{T-2} c(z_t^\delta, p_t^\delta)\cdot(x_t - x'_t) + \sum_{j=T-1}^{T} \epsilon_j\ c(z_j^\delta, p_j^\delta)\cdot(x_j - x'_j) \right\| > \theta/4 \right\} \right]
\end{aligned}$$

Continuing similarly all the way to the first term, we obtain an upper bound

$$\sup_{x_1,x_1'\in E_k}\sup_{p_1^\delta\in C_\delta}\mathbb{E}_{\epsilon_1}\ldots\sup_{x_T,x_T'\in E_k}\sup_{p_T^\delta\in C_\delta}\mathbb{E}_{\epsilon_T}\left[\mathbf{1}\left\{\sup_{c\in\zeta}\left\|\frac{1}{T}\sum_{t=1}^T\epsilon_t\ c(z_t^\delta,p_t^\delta)\cdot(x_t-x_t')\right\|>\theta/4\right\}\right]$$

We now pass to the tree notation. The above quantity is equal to

$$\sup_{\mathbf{x},\mathbf{x}',\mathbf{p}^\delta}\quad\mathbb{E}_\epsilon\left[\mathbf{1}\left\{\sup_{c\in\zeta}\left\|\frac{1}{T}\sum_{t=1}^T\epsilon_t\ c(\mathbf{z}_t^\delta(\epsilon),\mathbf{p}_\mathbf{t}^\delta(\epsilon))\cdot(\mathbf{x}_t(\epsilon)-\mathbf{x}_t'(\epsilon))\right\|>\theta/4\right\}\right]$$

where $\mathbf{x},\mathbf{x}'$ are E_k-valued trees of depth T, $\mathbf{p}^\delta$ is a C_δ-valued tree of depth T, and the $\mathcal{Z}$-valued history tree is defined for by

$$\mathbf{z}_t^\delta(\epsilon):=\Big((\mathbf{p}_1^\delta(\epsilon),\mathbf{x}_1(\epsilon)),\ldots,(\mathbf{p}_{t-1}^\delta(\epsilon),\mathbf{x}_{t-1}(\epsilon))\Big)\,.$$

Here, $\epsilon=(\epsilon_1,\ldots,\epsilon_T)\in\{\pm1\}^T$ denotes a path. The last quantity is upper bounded by

$$\begin{aligned}
&\sup_{\mathbf{x},\mathbf{x}',\mathbf{p}^\delta}\quad\mathbb{E}_\epsilon\left[\mathbf{1}\left\{\sup_{c\in\zeta}\left\|\frac{1}{T}\sum_{t=1}^T\epsilon_t\ c(\mathbf{z}_t^\delta(\epsilon),\mathbf{p}_\mathbf{t}^\delta(\epsilon))\mathbf{x}_t(\epsilon)\right\|+\left\|\frac{1}{T}\sum_{t=1}^T\epsilon_t\ c(\mathbf{z}_t^\delta(\epsilon),\mathbf{p}_\mathbf{t}^\delta(\epsilon))\mathbf{x}_t'(\epsilon)\right\|>\theta/4\right\}\right]\\
&\le\sup_{\mathbf{x},\mathbf{x}',\mathbf{p}^\delta}\quad\mathbb{E}_\epsilon\left\{\mathbf{1}\left\{\sup_{c\in\zeta}\left\|\frac{1}{T}\sum_{t=1}^T\epsilon_t\ c(\mathbf{z}_t^\delta(\epsilon),\mathbf{p}_\mathbf{t}^\delta(\epsilon))\mathbf{x}_t(\epsilon)\right\|>\theta/8\right\}\right.\\
&\qquad\qquad\left.+\mathbf{1}\left\{\sup_{c\in\zeta}\left\|\frac{1}{T}\sum_{t=1}^T\epsilon_t\ c(\mathbf{z}_t^\delta(\epsilon),\mathbf{p}_\mathbf{t}^\delta(\epsilon))\mathbf{x}_t'(\epsilon)\right\|>\theta/8\right\}\right\}\\
&\le2\sup_{\mathbf{x},\mathbf{p}^\delta}\quad\mathbb{E}_\epsilon\left[\mathbf{1}\left\{\sup_{c\in\zeta}\left\|\frac{1}{T}\sum_{t=1}^T\epsilon_t\ c(\mathbf{z}_t^\delta(\epsilon),\mathbf{p}_\mathbf{t}^\delta(\epsilon))\ \mathbf{x}_t(\epsilon)\right\|>\theta/8\right\}\right]\\
&=2\sup_{\mathbf{x},\mathbf{p}^\delta}\quad\mathbb{P}_\epsilon\left(\sup_{c\in\zeta}\left\|\frac{1}{T}\sum_{t=1}^T\epsilon_t\ c(\mathbf{z}_t^\delta(\epsilon),\mathbf{p}_\mathbf{t}^\delta(\epsilon))\ \mathbf{x}_t(\epsilon)\right\|>\theta/8\right)
\end{aligned}$$

■