

A Geometric Proof of Calibration

Shie Mannor*

Gilles Stoltz†

Israel Institute of Technology (Technion)
Haifa, Israel

shie@ee.technion.ac.il

Ecole Normale Supérieure, CNRS
Paris, France

&

HEC Paris, CNRS,
Jouy-en-Josas, France

gilles.stoltz@ens.fr

December 17, 2009

Abstract

We provide yet another proof of the existence of calibrated forecasters; it has two merits. First, it is valid for an arbitrary finite number of outcomes. Second, it is short and simple and it follows from a direct application of Blackwell's approachability theorem to carefully chosen vector-valued payoff function and convex target set. Our proof captures the essence of existing proofs based on approachability (e.g., the proof by Foster [1999] in case of binary outcomes) and highlights the intrinsic connection between approachability and calibration.

JEL codes: C72, C73

Keywords: Calibration, Approachability, Convergence rates

*Partially supported by the ISF under contract 890015 and a Horev Fellowship.

† *Corresponding author.* Partially supported by the French "Agence Nationale pour la Recherche" under grant JCJC06-137444 "From applications to theory in learning and adaptive statistics" and by the PASCAL Network of Excellence under EC grant no. 506778.

1 Motivation

Already in 1999, Foster [1999] stated that:

Over the past few years many proofs of the existence of calibration have been discovered. Each of the following provides a different algorithm and proof of convergence: Foster and Vohra [1991, 1998], Hart [1995], Fudenberg and Levine [1999], Hart and Mas-Colell [2000]. Does the literature really need one more? Probably not.

Despite all, he argued, successfully, that his new proof of the existence of calibrated forecasters in the case of binary outcomes, based on Blackwell's approachability theorem (Blackwell [1956]), was shorter and more direct than most of the previous proofs.

In this paper, we consider the general case of finitely many outcomes and exhibit an even shorter (ten-line long) proof of the existence of calibrated forecasters based on approachability. We show therefore that calibration is a straightforward consequence of approachability. As we realized by browsing on the web, approachability and calibration are well-taught matters and we are pretty confident that this new proof will become a standard example in the list of direct applications of approachability (as is already, for instance, the existence of no-regret forecasters).

Foster [1999] mentions that his approachability-based proof of the existence of a calibrated forecaster was obtained by first considering (a modification of) an intuitive forecaster already stated in Foster and Vohra [1991] and then working out the proof of its guarantees. We proceed the other way round and start directly from the statement of Blackwell's approachability theorem for convex sets [Blackwell, 1956, Theorem 3] but, as a drawback, can only exhibit a forecaster which has to solve a linear program at each step. Taking a closer look at Foster [1999], one can see that we indeed capture the essence of his previous proof. His algorithm is a clever modification, in the case of binary outcomes, of the general approachability-based forecaster presented below; the former has a nice, explicit, and simple statement.

We now recall the informal definition and consequences of calibration. Consider a finite set of possible outcomes and suppose we obtain random forecasts about future events; these forecasts are each given by probability distributions over the outcomes. Now, such a sequence of forecasts is called calibrated whenever it is consistent in hindsight, that is, when, for all distributions \mathbf{p} , the actual empirical distribution of the outcomes on those rounds when the forecast was close to \mathbf{p} is also close to \mathbf{p} .

Having a calibrated forecasting scheme is beneficial in several ways. On the one hand, it allows some agent to choose the best responses to the predicted forecasts or to consider other risk measures which might be more valuable than greedily choosing the best action leading to highest reward. On the other hand, calibrated forecasting rules enable multiple agents to converge to a reasonable joint play in some situations. For instance, if all players use calibrated forecasts of other players' actions, then the empirical distribution of action profiles converges to the set of correlated equilibria; see Foster and Vohra [1997]. We refer to Sandroni et al. [2003] for further discussion on calibrated forecasting as well as its generalizations.

2 Setup and formal definition of calibration

We consider a finite set \mathcal{A} of outcomes, with cardinality denoted by A and denote by $\mathcal{P} = \Delta(\mathcal{A})$ the set of probability distributions over \mathcal{A} . We equip \mathcal{P} , which can be considered a subset of \mathbb{R}^A , with some¹ norm $\|\cdot\|$. In particular, the Dirac probability distribution on some outcome $a \in \mathcal{A}$ will be referred to as δ_a .

A forecaster plays a game against Nature. At each step, it outputs a probability distribution $P_t \in \mathcal{P}$ while Nature chooses simultaneously an outcome $a_t \in \mathcal{A}$. We make no assumption on Nature's strategy.

The goal of the forecaster is to ensure the following property, known as calibration: for all strategies of Nature,

$$\forall \varepsilon > 0, \quad \forall \mathbf{p} \in \mathcal{P}, \quad \lim_{T \rightarrow +\infty} \left\| \frac{1}{T} \sum_{t=1}^T \mathbb{I}_{\{\|P_t - \mathbf{p}\| \leq \varepsilon\}} (P_t - \delta_{a_t}) \right\| = 0 \quad \text{a.s.} \quad (1)$$

The a.s. statement accounts for randomized forecasters. (It was shown by Oakes [1985] and Dawid [1985] that randomization is essential for calibration.)

The literature (e.g., Foster and Vohra [1998], Foster [1999]) essentially considers a less ambitious goal, at least in a first step: ε -calibration. (We explain in Section 4.2 how to get a calibrated forecaster from some sequence of ε -calibrated forecasters with good properties.) Formally, given $\varepsilon > 0$, an ε -calibrated forecaster considers some finite covering of \mathcal{P} by N_ε balls of radius ε and abides by the following constraints. Denoting by $\mathbf{p}_1, \dots, \mathbf{p}_{N_\varepsilon}$ the centers of the balls in the covering (they form what will be referred to later on as an ε -grid), the forecaster chooses only forecasts $P_t \in \{\mathbf{p}_1, \dots, \mathbf{p}_{N_\varepsilon}\}$. We thus denote by K_t the index in $\{1, \dots, N_\varepsilon\}$ such that $P_t = \mathbf{p}_{K_t}$. The final condition to be satisfied is then that for all strategies of Nature,

$$\lim_{T \rightarrow +\infty} \sum_{k=1}^{N_\varepsilon} \left\| \frac{1}{T} \sum_{t=1}^T \mathbb{I}_{\{K_t=k\}} (\mathbf{p}_k - \delta_{a_t}) \right\| \leq \varepsilon \quad \text{a.s.} \quad (2)$$

When the considered norm is the ℓ^1 -norm $\|\cdot\|_1$, the sum appearing in this criterion is usually referred to as the ℓ^1 -calibration score (Foster [1999]). Another popular criterion is the Brier score (Foster and Vohra [1998]), which we consider in Section 4.3; it is bounded, up to a factor of 2, by the ℓ^1 -calibration score.

¹The precise nature of this norm, e.g., ℓ^1 , Euclidian ℓ^2 , or ℓ^∞ supremum norm, is irrelevant at this stage, since all norms are equivalent on finite-dimensional spaces.

3 A geometric construction of ε -calibrated forecasters

In this section we prove our main result regarding the existence of an ε -calibrated forecaster based on approachability theory. We start with recalling approachability theory, provide the main result, and then address the issue of computational complexity.

3.1 Statement of Blackwell's approachability theorem

Consider a vector-valued game between two players, with respective finite action sets \mathcal{I} and \mathcal{J} . We denote by d the dimension of the reward vector. The payoff function of the first player is given by a mapping $m : \mathcal{I} \times \mathcal{J} \rightarrow \mathbb{R}^d$, which is linearly extended to $\Delta(\mathcal{I}) \times \Delta(\mathcal{J})$, the set of product-distributions over $\mathcal{I} \times \mathcal{J}$.

We denote by I_1, I_2, \dots and J_1, J_2, \dots the sequences of actions in \mathcal{I} and \mathcal{J} taken by each player (they are possibly given by randomized strategies). Let $C \subset \mathbb{R}^d$ be some set. By definition, C is approachable if there exists a strategy for the first player such that for all strategies of the second player,

$$\lim_{T \rightarrow \infty} \inf_{c \in C} \left\| c - \frac{1}{T} \sum_{t=1}^T m(I_t, J_t) \right\| = 0 \quad \text{a.s.}$$

That is, the first player has a strategy that ensures that the average of his vector-valued payoffs converges to the set C .

For closed convex sets C , there is a simple characterization of approachability that is a direct consequence of the minimax theorem.

Theorem 1 ([Blackwell, 1956, Theorem 3]). *A closed convex set $C \subset \mathbb{R}^d$ is approachable if and only if*

$$\forall \mathbf{q} \in \Delta(\mathcal{J}), \quad \exists \mathbf{p} \in \Delta(\mathcal{I}), \quad m(\mathbf{p}, \mathbf{q}) \in C .$$

3.2 Application to the existence of an ε -calibrated forecaster

As indicated above, we equip \mathcal{P} with some norm $\|\cdot\|$ and fix $\varepsilon > 0$; we then consider an associated ε -grid $\{\mathbf{p}_1, \dots, \mathbf{p}_{N_\varepsilon}\}$ in $\mathcal{P} = \Delta(\mathcal{A})$.

Theorem 2. *There exists an ε -calibrated forecaster which selects at every stage a distribution from this grid.*

Proof. We apply the results on approachability recalled above. To that end, we consider in our setting the action sets $\mathcal{I} = \{1, \dots, N_\varepsilon\}$ for the first player and $\mathcal{J} = \mathcal{A}$ for the second player.

We define the vector-valued payoff function as follows; it takes values in $\mathbb{R}^{AN_\varepsilon}$. For all $k \in \{1, \dots, N_\varepsilon\}$ and $a \in \mathcal{A}$,

$$m(k, a) = (\mathbf{0}, \dots, \mathbf{0}, \mathbf{p}_k - \delta_a, \mathbf{0}, \dots, \mathbf{0}) ,$$

which is a vector of N_ε elements of \mathbb{R}^A composed by $N_\varepsilon - 1$ occurrences of the zero element $\mathbf{0} \in \mathbb{R}^A$ and one non-zero element, located in the k -th position and given by the difference of probability distributions $\mathbf{p}_k - \delta_a$.

We now define the target set C as the following subset of the ε -ball around $(\underline{0}, \dots, \underline{0})$ for the underlying norm $\|\cdot\|$. We write (AN_ε) -dimensional vectors of $\mathbb{R}^{AN_\varepsilon}$ as N_ε -dimensional vectors with components in \mathbb{R}^A , i.e., for all $\underline{x} \in \mathbb{R}^{AN_\varepsilon}$,

$$\underline{x} = (\underline{x}_1, \dots, \underline{x}_{N_\varepsilon})$$

where $\underline{x}_k \in \mathbb{R}^A$ for all $k \in \{1, \dots, N_\varepsilon\}$. Then,

$$C = \left\{ \underline{x} \in \mathbb{R}^{AN_\varepsilon} : \sum_{k=1}^{N_\varepsilon} \|\underline{x}_k\| \leq \varepsilon \right\}.$$

Note that C is a closed convex set.

The condition (2) of ε -calibration can be rewritten as follows: the sequence of the

$$\bar{m}_T \stackrel{\text{def}}{=} \frac{1}{T} \sum_{t=1}^T m(K_t, a_t) = \left(\frac{1}{T} \sum_{t=1}^T \mathbb{I}_{\{K_t=1\}} (\mathbf{p}_1 - \delta_{a_t}), \dots, \frac{1}{T} \sum_{t=1}^T \mathbb{I}_{\{K_t=N_\varepsilon\}} (\mathbf{p}_{N_\varepsilon} - \delta_{a_t}) \right)$$

converges to the set C almost surely.

The existence of an ε -calibrated forecaster is thus equivalent to the approachability of C , which we now prove by showing that the characterization provided by Theorem 1 is satisfied. Let $\mathbf{q} \in \Delta(\mathcal{J}) = \mathcal{P}$. By construction, there exists $k \in \{1, \dots, N_\varepsilon\}$ such that $\|\mathbf{p}_k - \mathbf{q}\| \leq \varepsilon$ and thus

$$m(k, \mathbf{q}) \in C.$$

(Here, the distribution \mathbf{p} of the approachability theorem can be taken as the Dirac distribution δ_k .) \square

3.3 Computation of the exhibited ε -calibrated forecaster

The proof of the approachability theorem gives rise to an implicit strategy, as indicated in Blackwell [1956]. We denote here by Π_C the projection in ℓ^2 -norm onto C .

At each round $t \geq 2$ and with the notations above, the forecaster should pick his action K_t at random according to a distribution $\psi_t = (\psi_{t,1}, \dots, \psi_{t,N_\varepsilon})$ on $\{1, \dots, N_\varepsilon\}$ such that

$$\forall a \in \mathcal{A}, \quad \left(\bar{m}_{t-1} - \Pi_C(\bar{m}_{t-1}) \right) \cdot \left(m(\psi, a) - \Pi_C(\bar{m}_{t-1}) \right) \leq 0, \quad (3)$$

where \cdot denotes the inner product in $\mathbb{R}^{AN_\varepsilon}$. The proof of Theorem 1 (see Blackwell [1956]) shows that such a distribution ψ_t indeed exists; the question is how to efficiently compute it. To do so, we first need to compute the projection $\Pi_C(\bar{m}_{t-1})$ of \bar{m}_{t-1} .

We address the two computational issues separately. We first indicate how to find the projection efficiently and then explain how to find the distribution ψ_t based on the knowledge of this projection.

3.3.1 Projecting onto C

We need to find the closest point in C to \bar{m}_{t-1} . Since C is convex and the ℓ^2 -norm is convex, we have to deal with a minimization problem of a convex function over a convex set. Since

answering the question whether a given point is in C or not can be done in time linear in AN_ε , the projection problem can be solved (approximately) in time polynomial in AN_ε .

Now, for the special case where the underlying norm is the ℓ^1 -norm $\|\cdot\|_1$, we can do much better. For $i \in \{1, \dots, AN_\varepsilon\}$, we denote by $s_{i,t-1} \in \{-1, 1\}$ the sign of the i -th component $\bar{m}_{i,t-1}$ of the vector \bar{m}_{t-1} . (The value of the sign function at x is arbitrary at $x = 0$, equal to -1 when $x < 0$ and to 1 when $x > 0$.) Then, $\Pi_C(\bar{m}_{t-1})$ is the solution of the following optimization problem, where the unknown is $\underline{y} = (y_1, \dots, y_{AN_\varepsilon})$:

$$\begin{aligned} \min \quad & \|\underline{y} - \bar{m}_{t-1}\|_2^2 \\ \text{such that} \quad & \begin{cases} \sum_{i=1}^{AN_\varepsilon} y_i s_{i,t-1} \leq \varepsilon \\ y_i s_{i,t-1} \geq 0, \quad \forall i \in \{1, \dots, AN_\varepsilon\}. \end{cases} \end{aligned}$$

It can be easily shown (e.g., by an immediate adaptation of [Palomar, 2005, Lemma 1]) that the optimal solution is unique; it is given by $\underline{y}(\mu^*)$ where for all $\mu \geq 0$,

$$\underline{y}(\mu) = s_{i,t-1} (s_{i,t-1} \bar{m}_{i,t-1} - \mu)^+$$

and μ^* is chosen as the minimum nonnegative value such that $\sum_i y_i(\mu) s_{i,t-1} \leq \varepsilon$. (Note that if $\mu^* > 0$ then $\sum_i y_i(\mu^*) s_{i,t-1} = \varepsilon$.)

It follows, as a conclusion, that projecting onto C can be done in linear time in AN_ε when the underlying norm is the ℓ^1 -norm $\|\cdot\|_1$.

3.4 Finding the optimal distribution ψ_t in (3)

The question that has to be resolved is therefore how to find ψ_t that satisfies condition (3). Since we know that such a ψ_t exists, it suffices, for instance, to compute an element of

$$\operatorname{argmin}_{\psi} \max_{a \in \mathcal{A}} \left(\bar{m}_{t-1} - \Pi_C(\bar{m}_{t-1}) \right) \cdot m(\psi_t, a) = \operatorname{argmin}_{\psi} \max_{a \in \mathcal{A}} \sum_{k=1}^{N_\varepsilon} \psi_k \gamma_{k,a,t-1}$$

where we denoted $\gamma_{k,a,t-1} = \left(\bar{m}_{t-1} - \Pi_C(\bar{m}_{t-1}) \right) \cdot m(k, a)$.

This can be done efficiently by linear programming leading to a polynomial complexity in N_ε and A .

However, if instead of solving the minimax problem exactly we are satisfied with solving it approximately, i.e., allowing a small violation $\delta > 0$ in each of the A constraints given by (3), we can use the multiplicative weights algorithm as explained in Freund and Schapire [1999]; see also [Cesa-Bianchi and Lugosi, 2006, Section 7.2]. The complexity of such a solution would be

$$O\left(\frac{AN_\varepsilon}{\delta^2} \ln N_\varepsilon\right),$$

since $(\ln N_\varepsilon)/\delta^2$ steps of complexity AN_ε each have to be performed.

The proof of Blackwell’s approachability theorem shows that in this case the sequence of the average payoff vectors \bar{m}_t converges rather to the $\sqrt{\delta}$ -expansion (in ℓ^2 -norm) of C ; it is easy² to see that the latter is included in the δ -expansion (in ℓ^1 -norm) of C .

Putting all things together, we obtain that if we consider the ℓ^1 -norm $\|\cdot\|_1$ as the underlying norm (in particular, to define C), we can find a 2ε -calibrated forecaster whose complexity is of the order of $AN_\varepsilon \varepsilon^{-2} \log N_\varepsilon$ at each step. Since N_ε behaves like $\varepsilon^{-(A-1)}$ we have that the dependence of the complexity per stage behaves like $\varepsilon^{-(A+1)}$ (ignoring multiplicative and logarithmic factors). This implies a polynomial dependence in ε but an exponential dependence in A .

Remark 1. It is worth noting that when choosing a solution ψ_t , it is not possible to replace ψ_t with its mean or with an element of $\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_{N_\varepsilon}$ that is close to its mean. The reason is that this would give rise to a deterministic rule, which, as we mentioned in Section 2, cannot be calibrated. The fact that we have to randomize rather than take the mean is due to our construction of the vector-valued game; therein, playing a mixed action ψ_t over the \mathbf{p}_i ’s leads to a very different vector-valued reward than playing the (element \mathbf{p}_k closest to the) mean of the mixed action. This is because different indices of the (AN_ε) -dimensional space are involved.

4 Rates of convergence and construction of a calibrated forecaster

In this section we provide rates of convergence and discuss the construction of a calibrated (rather than ε -calibrated) forecaster. We finally compare our results to some existing calibrated forecasters in the literature.

4.1 Rates of convergence

Approachability theory provides uniform convergence rates of sequence of empirical payoff vectors to the target set, see [Cesa-Bianchi and Lugosi, 2006, Exercise 7.23]. Formally, denoting by $\|\cdot\|_2$ the Euclidian norm in $\mathbb{R}^{AN_\varepsilon}$, there exists some absolute constant γ (independent of A and N_ε) such that for all strategies of Nature and for all T , with probability $1 - \delta$,

$$\|\bar{m}_T - \Pi_C(\bar{m}_T)\|_2 \leq \gamma \sqrt{\frac{\ln(1/\delta)}{T}}.$$

Here, it is crucial to state the convergence rates based on the Euclidian norm because of some underlying martingale convergence argument in Hilbert spaces proved by Chen and White [1996]. The reason why the convergence rate is independent here of A and N_ε is that the payoff vectors $m(k, a)$ all have an Euclidian norm bounded by an absolute constant, e.g., 2; this happens because most of their components are 0.

However, the set C can be defined by a different norm $\|\cdot\|$. Below, we will define it based on the ℓ^1 -norm, for instance. The stated uniform convergence rate can be used since, via a triangle inequality and an application of the Cauchy-Schwarz inequality,

$$\|\bar{m}_T\|_1 \leq \|\Pi_C(\bar{m}_T)\|_1 + \|\bar{m}_T - \Pi_C(\bar{m}_T)\|_1 \leq \varepsilon + \sqrt{AN_\varepsilon} \|\bar{m}_T - \Pi_C(\bar{m}_T)\|_2.$$

²It suffices to note that for all vectors Δ of a finite-dimensional space, one has $\|\Delta\|_\infty \leq \|\Delta\|_2$, so that the inequality $\|\Delta\|_2 \leq \sqrt{\|\Delta\|_\infty \|\Delta\|_1}$ yields $\sqrt{\|\Delta\|_2} \leq \|\Delta\|_1$.

N_ε is of the order of $\varepsilon^{-(A-1)}$; we let γ' be an absolute constant such that $N_\varepsilon \leq \gamma' \varepsilon^{-(A-1)}$ for all $\varepsilon \leq 1$ (say). We therefore have proved that given $0 < \varepsilon \leq 1$, the forecaster defined in the previous section is such that for all strategies of Nature and for all T , with probability $1 - \delta$,

$$\|\bar{m}_T\|_1 = \sum_{k=1}^{N_\varepsilon} \left\| \frac{1}{T} \sum_{t=1}^T \mathbb{I}_{\{K_t=k\}} (\mathbf{p}_k - \delta_{a_t}) \right\|_1 \leq \varepsilon + \gamma \gamma' \sqrt{A} \sqrt{\frac{\ln(1/\delta)}{\varepsilon^{(A-1)} T}} \stackrel{\text{def}}{=} U_{\varepsilon, T, \delta}.$$

4.2 Construction of a calibrated forecaster

We use a standard approach which is commonly known as the “doubling trick,” see, e.g., Cesa-Bianchi and Lugosi [2006]. It consists in defining a meta-forecaster that proceeds in regimes; regime r (where $r \geq 1$) lasts T_r rounds and resorts for the forecasts to an ε_r -calibrated forecaster, for some $\varepsilon_r > 0$ to be defined by the analysis. We now show that for appropriate values of the T_r and ε_r , the resulting meta-forecaster is calibrated in the sense of (1), and even uniformly calibrated in the following sense, where \mathcal{B} denotes the Borel sigma-algebra of \mathcal{P} :

$$\lim_{T \rightarrow +\infty} \sup_{B \in \mathcal{B}} \left\| \frac{1}{T} \sum_{t=1}^T \mathbb{I}_{\{P_t \in B\}} (P_t - \delta_{a_t}) \right\| = 0 \quad \text{a.s.} \quad (4)$$

Of course, uniform calibration (4) implies calibration (1) via the choices for B given by ε -balls around probability distributions \mathbf{p} .

For concreteness, we focus below on the ℓ^1 -calibration score.

Regimes are indexed by $r = 1, 2, \dots$ and the index of the regime corresponding to round T is referred to as R_T . The set of the rounds within regime $r \leq R_T - 1$ is called \mathcal{T}_r ; rounds in regime R_T with index less than T are gathered in the set \mathcal{T}_{R_T} (we commit here an abuse of notations). We denote by $\mathbf{p}_{k,r}$, where $k \in \{1, \dots, N_{\varepsilon_r}\}$, the finite ε_r -grid considered in the r -th regime. By the triangle inequality satisfied by $\|\cdot\|$, we first decompose the quantity of interest according to the regimes and to the played points of the grids,

$$\left\| \sum_{t=1}^T \mathbb{I}_{\{P_t \in B\}} (P_t - \delta_{a_t}) \right\|_1 \leq \sum_{r=1}^{R_T} \sum_{k=1}^{N_{\varepsilon_r}} \mathbb{I}_{\{\mathbf{p}_{k,r} \in B\}} \left\| \sum_{t \in \mathcal{T}_r} \mathbb{I}_{\{K_t=k\}} (\mathbf{p}_{k,r} - \delta_{a_t}) \right\|_1.$$

We now substitute the uniform bound obtained in the previous section and get that with probability $1 - (\delta_{1,T} + \dots + \delta_{R_T,T}) \geq 1 - 1/T^2$,

$$\sup_{B \in \mathcal{B}} \left\| \frac{1}{T} \sum_{t=1}^T \mathbb{I}_{\{P_t \in B\}} (P_t - \delta_{a_t}) \right\|_1 \leq \frac{1}{T} \sum_{r=1}^{R_T} T_r U_{\varepsilon_r, T_r, \delta_{r,T}}$$

where we defined $\delta_{r,T} = 1/(2^r T^2)$.

An application of the Borel-Cantelli lemma and Cesaro's lemma shows that for suitable choices of a sequence ε_r decreasing towards 0 and an increasing sequence T_r such that $\varepsilon_r^{A-1} T_r$ tends to infinity fast enough, one then gets the desired convergence (4). For instance, if $T_r = 2^r$, and ε_r is chosen such that

$$\varepsilon_r \quad \text{and} \quad \sqrt{\frac{1}{\varepsilon_r^{(A-1)} T_r}}$$

are of the same order of magnitude, e.g., $\varepsilon_r = 2^{-r/(A+1)}$, then

$$\limsup_{T \rightarrow \infty} \frac{T^{1/(A+1)}}{\sqrt{\ln T}} \sup_{B \in \mathcal{B}} \left\| \frac{1}{T} \sum_{t=1}^T \mathbb{I}_{\{P_t \in B\}} (P_t - \delta_{a_t}) \right\|_1 \leq \Gamma_A \quad \text{a.s. ,}$$

where the constant Γ_A depends only on A . To the best of our knowledge, this is the first rates results for calibration for an alphabet of size A larger than 2.

Remark 2. Our approach can be easily extended to calibration of countably infinite alphabets; in that case both ε and A would change with the regime. We omit the technical details here.

4.3 Comparison to previous forecasters

4.3.1 ℓ^1 -calibration score

Foster [1999] first considered the ℓ^1 -calibration score in the context of the prediction of binary outcomes only, i.e., when $A = 2$. The ε -calibrated forecaster he explicitly exhibited has a computational complexity of the order of $1/\varepsilon$. He did not work out the convergence rates but since his procedure is mostly a clever twist on our general procedure, they should be similar to the ones we proved in Section 4.1.

4.3.2 Brier score

What follows is extracted from Foster and Vohra [1998]; see also [Cesa-Bianchi and Lugosi, 2006, Section 4.5].

Given an ε -grid over the simplex \mathcal{P} , we define, for all $k \in \{1, \dots, N_\varepsilon\}$, the empirical distribution of the outcomes chosen by Nature at those rounds t when the forecaster used \mathbf{p}_k ,

$$\rho_T(k) = \begin{cases} \mathbf{p}_k & \text{if } \sum_{t=1}^T \mathbb{I}_{\{K_t=k\}} = 0, \\ \sum_{t=1}^T \mathbb{I}_{\{K_t=k\}} \frac{1}{\sum_{s=1}^T \mathbb{I}_{\{K_s=k\}}} \delta_{a_t} & \text{if } \sum_{t=1}^T \mathbb{I}_{\{K_t=k\}} > 0. \end{cases}$$

The Brier score is defined as the following cumulative criterion,

$$\sum_{k=1}^{N_\varepsilon} \|\rho_T(k) - \mathbf{p}_k\|_2^2 \left(\frac{1}{T} \sum_{t=1}^T \mathbb{I}_{\{K_t=k\}} \right).$$

Since for two probability distributions \mathbf{p} and \mathbf{q} of \mathcal{P} , one always has

$$\|\mathbf{p} - \mathbf{q}\|_2^2 \leq 2 \|\mathbf{p} - \mathbf{q}\|_1,$$

the Brier score can be seen to be upper bounded by twice the ℓ^1 -calibration score; it is thus a weaker criterion.

Section 4.5 in Cesa-Bianchi and Lugosi [2006] shows however that forecasters with Brier scores asymptotically smaller than ε can be the keystones to construct calibrated forecasters, in a way similar to the construction exhibited in Section 4.2.

In the case $A = 2$, these forecasters essentially bound the Brier score, with probability at least $1 - \delta$, by something of the order of

$$\varepsilon + \frac{1}{\varepsilon} \sqrt{\frac{\ln(1/\varepsilon) + \ln(1/\delta)}{T}},$$

which is worse than the rate we could exhibit in Section 4.1 for the ℓ^1 -calibration score.

In addition, the computational complexity of the underlying procedure (based on the minimization of internal regret) is of the order of $1/\varepsilon^2$ per stage and thus is similar to the complexity $1/\varepsilon^{A+1} = 1/\varepsilon^2$ we derived in Section 3.3 for our new procedure.

The general case of $A \geq 3$ is briefly mentioned in Cesa-Bianchi and Lugosi [2006] (Section 4.5) who indicate that the case of $A = 2$ can be extended to $A \geq 3$ without further details. As far as we can say, the computational complexity of such an extension per step would be of the order of $1/\varepsilon^{2(A-1)}$ versus $1/\varepsilon^{(A+1)}$ for the approachability-based procedure we suggested above. The convergence rates, for a straightforward extension, seem to be quite slow.

References

- D. Blackwell. An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6:1–8, 1956.
- N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, Cambridge, United Kingdom, 2006.
- X. Chen and H. White. Laws of large numbers for Hilbert space-valued mixingales with applications. *Econometric Theory*, 12:284–304, 1996.
- A.P. Dawid. The impossibility of inductive inference. *Journal of the American Statistical Association*, 80:340–341, 1985.
- D. Foster. A proof of calibration via Blackwell’s approachability theorem. *Games and Economic Behavior*, 29:73–78, 1999.
- D. Foster and R. Vohra. Asymptotic calibration. Technical report, Graduate School of Business, University of Chicago, 1991.
- D. Foster and R. Vohra. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21:40–55, 1997.
- D. Foster and R. Vohra. Asymptotic calibration. *Biometrika*, 85:379–390, 1998.
- Y. Freund and R. Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29:79–103, 1999.
- D. Fudenberg and D. Levine. An easier way to calibrate. *Games and Economic Behavior*, 29:131–137, 1999.
- S. Hart. Personal communication to Dean P. Foster, 1995.
- S. Hart and A. Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68:1127–1150, 2000.
- D. Oakes. Self-calibrating priors do not exist. *Journal of the American Statistical Association*, 80:339, 1985.
- D.P. Palomar. Convex primal decomposition for multicarrier linear MIMO transceivers. *IEEE Transactions on Signal Processing*, 53(12):4661–4674, 2005.
- A. Sandroni, R. Smorodinsky, and R. Vohra. Calibration with many checking rules. *Mathematics of Operations Research*, 28:141–153, 2003.