Supermartingales in Prediction with Expert Advice

Alexey Chernov, Yuri Kalnishkan, Fedor Zhdanov, and Vladimir Vovk

Computer Learning Research Centre, Department of Computer Science
Royal Holloway, University of London, Egham, Surrey TW20 0EX, UK
{chernov,yura,fedor,vovk}@cs.rhul.ac.uk

Abstract. This paper compares two methods of prediction with expert advice, the Aggregating Algorithm and the Defensive Forecasting, in two different settings. The first setting is traditional, with a countable number of experts and a finite number of outcomes. Surprisingly, these two methods of fundamentally different origin lead to identical procedures. In the second setting the experts can give advice conditional on the learner's future decision. Both methods can be used in the new setting and give the same performance guarantees as in the traditional setting. However, whereas defensive forecasting can be applied directly, the AA requires substantial modifications.

1 Introduction

The framework of prediction with expert advice was introduced in the late 1980s. In contrast to statistical learning theory, the methods of prediction with expert advice work without making any statistical assumption about the source of data. The role of the assumptions is played by a "pool of experts" that the predictor competes with. For references and details, see the monograph [3].

Many methods for prediction with expert advice are known. This paper deals with two of them: the Aggregating Algorithm [15] and defensive forecasting [17]. The Aggregating Algorithm (the AA for short) is a member of the family of exponential-weights algorithms, and so implements a Bayesian-type aggregation; various optimality properties of the AA have been established [16]. Defensive forecasting is a recently developed technique that combines the ideas of game-theoretic probability [12] with Levin and Gács's ideas of neutral measure [7, 9] and Foster and Vohra's ideas of universal calibration [5].

The idea of defensive forecasting is that a prediction strategy is developed assuming that we are given probability forecasts satisfying a convenient law of probability. In game-theoretic probability, a law of probability is represented as a strategy for an imaginary opponent, Sceptic, whose capital tends to infinity (or becomes large) if the law is violated. Sceptic's capital is a supermartingale, and a well-known result (Lemma 3 of this paper) says that there is a forecasting strategy that prevents Sceptic's capital from growing, thereby forcing the law of probability. This paper gives a self-contained description of a simple version of the method of defensive forecasting; we have no need to talk about laws of probability (for us, they are synonymous with game-theoretic supermartingales).

The older versions of defensive forecasting (see, e.g., [17]) minimize Learner's actual loss with the help of the following trick: a probability forecasting system is constructed such that the actual losses (Learner's and experts') are close to the (one-step-ahead conditional) expected losses; at each step Learner minimizes the expected loss. (Therefore, the law of probability used is the conjunction of several laws of large numbers.) Defensive forecasting, as well as the AA, can be used for competitive online prediction against "pools of experts" consisting of all functions from a large functional class (see [18, 19]). However, the loss bounds proved so far are generally incomparable: for large classes (such as many Sobolev spaces), defensive forecasting is better, whereas for smaller classes (such as classes of analytical functions), the AA works better. Note that the optimality results for the AA are obtained for the case of free agents as experts, not functions from a given class, thus we need to evaluate the algorithms anew.

In this paper, the AA and defensive forecasting are compared in the simple case of a finite number of outcomes. Learner competes with a pool of experts Θ. Learner and the experts suffer some loss at each step. We are interested in the performance guarantees of the form

$$\forall\theta \in \Theta \quad L_N \le cL_N(\theta) + a(\theta)\,,$$

where L_N is the cumulative loss of Learner and $L_N(\theta)$ is the cumulative loss of expert θ over the first N steps, c is some constant and a depends on θ only. In this case, we prove the following fact: an exponent of the regret $(L_N - cL_N(\theta))$ is a supermartingale if and only if the AA guarantees for this c the bound above. A defensive forecasting algorithm exploiting this fact turns out to give the same predictions as the AA.

Then we consider a new setting for prediction with expert advice, where the experts are allowed to "second-guess", that is, to give "conditional" predictions that are functions of the future Learner's decision (cf. the notion of internal regret [6]). If the dependence is regular enough (the expert's loss is continuous in Learner's loss), the Defensive Forecasting algorithm works in the new setting virtually without changes and guarantees the same performance bound as in the traditional setting. The AA in its original form cannot work in the new setting, and we suggest a modified version of the AA for this case.

2 Aggregating Algorithm

We begin with formulating the setting of the prediction with expert advice and the AA. Then we give a proof of the standard performance bound for the AA, which provides an insight to the "supermartingale" nature of the bound.

A *game of prediction* consists of three components: a non-empty finite set Ω of possible outcomes, a non-empty set Γ of possible decisions, and a function $\lambda : \Omega \times \Gamma \to [0, \infty]$ called the *loss function*. For technical convenience, we identify each decision $\gamma \in \Gamma$ with the function $\omega \mapsto \lambda(\omega, \gamma)$ (and also with a point in an $|\Omega|$-dimensional Euclidian space with pointwise operations). Let $\Lambda = \{ g \in [0,\infty]^\Omega \mid \exists \gamma \in \Gamma \, \forall \omega \in \Omega \; g(\omega) = \lambda(\omega,\gamma) \}$ be the set of *predictions*. From now on, a game is a pair (Ω,Λ), where $\Lambda \subseteq [0,\infty]^\Omega$.

The game of prediction with expert advice is played by Learner, Reality, and Experts; the set ("pool") of Experts is denoted by Θ. We will assume that Θ is (finite or) countable. In this paper, there is no loss of generality in assuming that Reality and all Experts are cooperative, since we are only interested in what can be achieved by Learner alone; therefore, we essentially consider a two-player game. The protocol of the game is the following.

PREDICTION WITH EXPERT ADVICE

$L_0 := 0$.
$L_0^\theta := 0$, $\theta \in \Theta$.
FOR $n = 1, 2, \dots$:
  Experts $\theta \in \Theta$ announce $\gamma_n^\theta \in \Lambda$.
  Learner announces $\gamma_n \in \Lambda$.
  Reality announces $\omega_n \in \Omega$.
  $L_n := L_{n-1} + \gamma_n(\omega_n)$.
  $L_n^\theta := L_{n-1}^\theta + \gamma_n^\theta(\omega_n)$.
END FOR.

The goal of Learner is to keep L_n less or at least not much greater than L_n^θ, at each step n and for all $\theta \in \Theta$.

To analyse the game, we need some additional notation. A point $g \in [0,\infty]^\Omega$ is called a *superprediction* if there is $\gamma \in \Lambda$ such that $\gamma(\omega) \le g(\omega)$ for all $\omega \in \Omega$. The last property will be denoted by $\gamma \le g$. Let Σ_Λ be the set of all superpredictions.

The *Aggregating Algorithm* is a strategy for Learner. It has four parameters: reals $c \ge 1$ and $\eta > 0$, a *distribution* P_0 on Θ (that is, $P_0(\theta) \in [0,1]$ for all $\theta \in \Theta$ and $\sum_\Theta P_0(\theta) = 1$), and a substitution function $\sigma : \Sigma_\Lambda \to \Lambda$ such that $\sigma(g) \le g$ for any $g \in \Sigma_\Lambda$.

At step N, the AA takes a point $g_N \in [0,\infty]^\Omega$ defined by the formula

$$g_N(\omega) = -\frac{c}{\eta} \ln \sum_{\theta \in \Theta} \frac{P_{N-1}(\theta)}{\sum_{\theta \in \Theta} P_{N-1}(\theta)} \exp(-\eta \gamma_N^\theta(\omega)) \,,$$

where

$$P_{N-1}(\theta) = P_0(\theta) \prod_{n=1}^{N-1} \exp(-\eta \gamma_n^\theta(\omega_n))$$

is an unnormalized distribution on Θ. Then, $\gamma_N = \sigma(g_N)$ is announced as the next prediction of Learner.

The step of AA is correct if and only if g_N is a superprediction. The necessary and sufficient condition for this is

$$\exists \gamma_N \in \Lambda \, \forall \omega \quad \gamma_N(\omega_N) \le -\frac{c}{\eta} \ln \sum_{\theta \in \Theta} \frac{P_{N-1}(\theta)}{\sum_{\theta \in \Theta} P_{N-1}(\theta)} \exp(-\eta \gamma_N^\theta(\omega)) \,. \qquad (1)$$

We say that the AA is *realizable* for certain c and η if the condition (1) is true regardless of $\gamma_N^\theta \in \Lambda$ and P_{N-1} (that is, regardless of P_0, the history, and the opponents' moves). In other words, for any finite set $G \subseteq \Lambda$ and for any distribution ρ on G, it holds that

$$\exists \gamma \in \Lambda \; \forall \omega \quad \exp\left(-\frac{\eta}{c}\gamma(\omega)\right) \geq \sum_{g \in G} \rho(g) \exp(-\eta g(\omega)) \,. \tag{2}$$

A detailed survey of the AA, its properties, attainable bounds and respective conditions on c and η for different games can be found in [16]. Remark only that if the AA is realizable for $c = 1$ and some η, the game is called η-*mixable* (and mixable if it is η-mixable for some η), and this case is of special interest.

Theorem 1 (Vovk, 1990). *If the AA is realizable for c and η, then the AA with parameters c, η, P_0, and σ guarantees that at each step n for all experts θ*

$$L_n \leq cL_n^\theta + \frac{c}{\eta} \ln \frac{1}{P_0(\theta)} \,.$$

The theorem was proved in [15]. Here we reproduce the proof emphasizing the points we need in the sequel.

Proof. We need to deduce the performance bound from the condition (1). To this end, we will rewrite (1) and get a semi-invariant of AA—a value that does not grow.

First, note that if we replace $\exists \gamma_N \in \Lambda$ by $\exists \gamma_N \in \Sigma_\Lambda$ in (1), we get an equivalent statement Indeed, $\Lambda \subseteq \Sigma_\Lambda$, thus one direction is trivial. The other direction holds by definition of a superprediction.

Second, note P_{N-1} occur in (1) only as a ratio of $P_{N-1}(\theta)$ to their sum, so we can multiply all $P_{N-1}(\theta)$ by a constant (an expression without θ). Let us define Q_{N-1} by the formula $P_0(\theta)Q_{N-1}(\theta) = P_{N-1}(\theta)\prod_{n=1}^{N-1} \exp\left(\frac{\eta}{c}\gamma_n(\omega_n)\right)$, that is,

$$Q_{N-1}(\theta) = \exp\left(\eta \sum_{n=1}^{N-1} \left(\frac{\gamma_n(\omega_n)}{c} - \gamma_n^\theta(\omega_n)\right)\right) \,,$$

and replace $P_{N-1}(\theta)$ by $P_0(\theta)Q_{N-1}(\theta)$ in (1). The inequality transforms to

$$\sum_{\theta \in \Theta} P_0(\theta)Q_{N-1}(\theta) \geq \sum_{\theta \in \Theta} P_0(\theta)Q_{N-1}(\theta) \exp(-\eta\gamma_N^\theta(\omega)) \exp(\frac{\eta}{c}\gamma_N(\omega)) \,.$$

Finally, we get that (1) is equivalent to the following condition:

$$\exists \gamma_N \in \Sigma_\Lambda \; \forall \omega \quad \sum_{\theta \in \Theta} P_0(\theta)Q_N(\theta) \leq \sum_{\theta \in \Theta} P_0(\theta)Q_{N-1}(\theta) \tag{3}$$

(for ω_N in Q_N we substitute ω).

In other words, the AA (if realizable for c and η) guarantees that after each step n the value $\sum_{\theta \in \Theta} P_0(\theta)Q_n(\theta)$ does not increase whatever ω_n is chosen by

Reality. Since $\sum_{\theta\in\Theta} P_0(\theta)Q_0(\theta) = \sum_{\theta\in\Theta} P_0(\theta) = 1$, we get $\sum_{\theta\in\Theta} P_0(\theta)Q_N(\theta) \le 1$ and $Q_N(\theta) \le 1/P_0(\theta)$ for each step N. To complete the proof it remains to note that

$$Q_N(\theta) = \exp\left(\eta\left(\frac{L_N}{c} - L_N^\theta\right)\right) . \qquad \square$$

For $c = 1$, the value $\frac{1}{\eta}\ln\sum_\theta P_0(\theta)Q_N(\theta)$ is known as the exponential potential (see [3, Sections 3.3,3.5]) and plays an important role in the analysis of weighted average algorithms.

In the next section we show that the reason why condition (3) holds is essentially that the Q is a supermartingale.

3 Supermartingales

Let $\mathcal{P}(\Omega)$ be the set of all distributions on Ω. Note that since Ω is finite we can identify $\mathcal{P}(\Omega)$ with a $(|\Omega|-1)$-dimensional simplex in Euclidean space, with the standard distance and topology. Let E be any set (maybe, empty). A function $S\colon (E\times\mathcal{P}(\Omega)\times\Omega)^* \to \mathbb{R}$ is called a (game-theoretic) *supermartingale* if for any N, for any $e_1,\ldots,e_N \in E$, for any $\pi_1,\ldots,\pi_N \in \mathcal{P}(\Omega)$, for any $\omega_1,\ldots,\omega_{N-1} \in \Omega$, it holds that

$$\sum_{\omega\in\Omega} \pi_N(\omega)S(e_1,\pi_1,\omega_1,\ldots,e_{N-1},\pi_{N-1},\omega_{N-1},e_N,\pi_N,\omega) \le S(e_1,\pi_1,\omega_1,\ldots,e_{N-1},\pi_{N-1},\omega_{N-1}) . \quad (4)$$

Remark 1. In the context of algorithmic probability theory (e. g. [10, p. 296]), the word 'supermartingale' is used in the following sense. Let $\mu:\ \Omega^* \to [0,1]$ be a measure on Ω^*. A function $s:\ \Omega^* \to \mathbb{R}_+$ is a supermartingale with respect to μ if for any N and any $\omega_1,\ldots,\omega_{N-1} \in \Omega$ it holds that

$$\sum_{\omega\in\Omega} \mu(\,\omega \mid \omega_1,\ldots,\omega_{N-1})s(\omega_1,\ldots,\omega_{N-1},\omega) \le s(\omega_1,\ldots,\omega_{N-1}) ,$$

where $\mu(\,\omega \mid \omega_1,\ldots,\omega_{N-1}) = \frac{\mu(\omega_1,\ldots,\omega_{N-1},\omega)}{\mu(\omega_1,\ldots,\omega_{N-1})}$. The relation with the game-theoretic supermartingale notion is the following. Let us take any measure μ. Let e_n be any functions of $\omega_1\ldots,\omega_{n-1}$. Let $\pi_n(\omega)$ be $\mu(\,\omega \mid \omega_1,\ldots,\omega_{n-1})$. Having substituted these functions in a game-theoretic supermartingale S, one gets a probabilistic supermartingale with respect to μ.

A supermartingale S is called *forecast-continuous* if for each N, it is continuous as a function of π_N.

3.1 Two Examples of Supermartingales

Let us consider two examples of supermartingales that naturally arise from two widely used games of prediction.

The *logarithmic loss function* is defined by

$$\lambda(\omega,\gamma) := \begin{cases} -\ln\gamma & \text{if } \omega = 1, \\ -\ln(1-\gamma) & \text{if } \omega = 0, \end{cases}$$

where $\omega \in \{0,1\}$ and $\gamma \in [0,1]$ (notice that the loss function is allowed to take value ∞). The losses in the game are $L_N := \sum_{n=1}^N \lambda(\omega_n, \gamma_n)$ for Learner who predicts γ_n and $L_N^\theta := \sum_{n=1}^N \lambda(\omega_n, \gamma_n^\theta)$ for expert θ who predicts γ_n^θ. Consider an exponent of the difference of losses of Learner and the expert:

$$\exp\left(\eta \sum_{n=1}^N \Big(\lambda\left(\omega_n,\gamma_n\right) - \lambda\left(\omega_n,\gamma_n^\theta\right)\Big)\right).$$

Let us assign prediction $\gamma \in [0,1]$ to each probability distribution $(1-\gamma,\gamma)$ on $\{0,1\}$. With this identification, the expression above can be considered as a function on $([0,1] \times \mathcal{P}(\{0,1\}) \times \{0,1\})^*$.

Lemma 1. *For $\eta \in [0,1]$, the function above is a forecast-continuous supermartingale.*

Proof. The continuity is obvious. For the supermartingale property, it suffices to check that

$$p\mathrm{e}^{\eta(-\ln p + \ln g)} + (1-p)\mathrm{e}^{\eta(-\ln(1-p)+\ln(1-g))} \le 1,$$

i.e., that $p^{1-\eta}g^\eta + (1-p)^{1-\eta}(1-g)^\eta \le 1$ for all $p, g \in [0,1]$ (p stands for γ_n and g stands for γ_n^θ). The last inequality immediately follows from the inequality between the geometric and arithmetic means when $\eta \in [0,1]$. (The left-hand side of that inequality is a special case of what is known as the Hellinger integral in probability theory.) □

In the game with *quadratic loss function*, $\omega \in \{0,1\}$ and $\gamma \in [0,1]$ as before, and the losses of Learner and expert θ are $L_N := \sum_{n=1}^N (\gamma_n - \omega_n)^2$ and $L_N^\theta := \sum_{n=1}^N (\gamma_n^\theta - \omega_n)^2$, respectively.

Lemma 2. *For $\eta \in [0,2]$, the following function on $([0,1] \times \mathcal{P}(\{0,1\}) \times \{0,1\})^*$*

$$\exp\left(\eta \sum_{n=1}^N \left(\left(\gamma_n - \omega_n\right)^2 - \left(\gamma_n^\theta - \omega_n\right)^2\right)\right)$$

is a forecast-continuous supermartingale.

Proof. It is sufficient to check that

$$p\mathrm{e}^{\eta\left((p-1)^2-(g-1)^2\right)} + (1-p)\mathrm{e}^{\eta\left((p-0)^2-(g-0)^2\right)} \le 1$$

for all $p, g \in [0,1]$. If we substitute $g = p + x$, the last inequality will reduce to

$$p\mathrm{e}^{2\eta(1-p)x} + (1-p)\mathrm{e}^{-2\eta p x} \le \mathrm{e}^{\eta x^2}, \quad \forall x \in [-p, 1-p].$$

The last inequality is a simple corollary of Hoeffding's inequality [8, 4.16], which is true for any $h \in \mathbb{R}$ (cf. [3, Lemma A.1]). Indeed, applying Hoeffding's inequality to the random variable X that is equal to 1 with probability p and to 0 with probability $(1-p)$, we obtain $p\exp(h(1-p)) + (1-p)\exp(-hp) \le \exp(h^2/8)$, which the substitution $h := 2\eta x$ reduces to $p\exp(2\eta(1-p)x) + (1-p)\exp(-2\eta p x) \le \exp(\eta^2 x^2/2) \le \exp(\eta x^2)$, the last inequality assuming $\eta \le 2$. □

3.2 A Supermartingale Criterion if the AA is Realizable

In Lemmas 1 and 2, we take a certain function of losses, and consider it as a supermartingale by having identified a prediction γ with a distribution $(1-\gamma, \gamma)$. A similar approach works also in the general case.

Let $\alpha : \mathcal{P}(\Omega) \to \Sigma_\Lambda \subseteq [0,\infty]^\Omega$ map any distribution π to a prediction α_π. Given α, a real $c \ge 1$, and a real $\eta > 0$, let us define the following function on $(\Sigma_\Lambda \times \mathcal{P}(\Omega) \times \Omega)^*$:

$$Q(e_1, \pi_1, \omega_1, \ldots, e_N, \pi_N, \omega_N) = \exp\left(\eta \sum_{n=1}^{N} \left(\frac{\alpha_{\pi_n}(\omega_n)}{c} - e_n(\omega_n)\right)\right) . \qquad (5)$$

Note that this very function is used in Lemmas 1 and 2, and also it is the function $Q_N(\theta)$ from the proof of Theorem 1, with e_n standing for γ_n^θ and α_{π_n} standing for γ_n.

Our next goal (Theorems 2 and 3) is to show that the AA is realizable if and only if there exists α such that the function Q is a supermartingale.

Lemma 3. *Let S be a forecast-continuous supermartingale. For any N, for any $e_1, \ldots, e_N \in E$, for any $\pi_1, \ldots, \pi_{N-1} \in \mathcal{P}(\Omega)$, for any $\omega_1, \ldots, \omega_{N-1} \in \Omega$, it holds that*

$$\exists \pi \in \mathcal{P}(\Omega)\, \forall \omega \in \Omega \quad S(e_1, \pi_1, \omega_1, \ldots, e_N, \pi, \omega) \le S(e_1, \pi_1, \omega_1, \ldots, e_{N-1}, \pi_{N-1}, \omega_{N-1}) .$$

A variant of this lemma was originally proved by Levin [9]. For a full proof see [7, Theorem 6] and [20, Theorem 1].

Note that the property provided by Lemma 3 is essentially the condition (3).

Theorem 2. *Let α be a mapping from $\mathcal{P}(\Omega)$ to Σ_Λ and $c \ge 1$ and $\eta > 0$ reals such that Q is a forecast-continuous supermartingale. Then the AA is realizable for c and η.*

Proof. Let $G \subseteq \Lambda$ be an arbitrary finite set. To prove (2) for any distribution ρ on G, we construct a supermartingale Q^ρ on $(\mathcal{P}(\Omega) \times \Omega)^*$ (the set E in the definition of supermartingale may be empty), which is a ρ-average of Q with g substituted for e_1 ($e_2, e_3, \ldots$ may be arbitrary), and apply Lemma 3 for $N = 1$. Namely,

$$Q^\rho(\pi, \omega) = \sum_{g \in G} \rho(g) Q(g, \pi, \omega) .$$

By the lemma, there exists $\pi \in \mathcal{P}(\Omega)$ such that $Q^\rho(\pi, \omega) \le 1$ for all ω, that is,

$$\sum_{g \in G} \rho(g) \exp\left(\eta\left(\frac{\alpha_\pi(\omega)}{c} - g(\omega)\right)\right) \le 1\,.$$

Since $\alpha_\pi \in \Sigma_\Lambda$, there is $\gamma \in \Lambda$ such that $\gamma \le \alpha_\pi$, which completes the proof. □

For the converse statement, we need three assumptions about Λ.

Assumption 1. Λ is a compact set.

Assumption 2. There is $\gamma \in \Lambda$ such that $\gamma(\omega) < \infty$ for all ω.

These assumptions are standard (see [16]). The third assumption is new and very technical. First we introduce some definitions that will be useful also in the proof of the theorem below.

For a given η, the *exp-convex hull* of Σ_Λ is the set $\Xi_\Lambda \supseteq \Sigma_\Lambda$ that consists of all points $g \in [0, \infty]^\Omega$ of the form

$$g(\omega) = \log_{(\mathrm{e}^{-\eta})} \sum_{\gamma \in G} \rho(\gamma) \left(\mathrm{e}^{-\eta}\right)^{\gamma(\omega)} = -\frac{1}{\eta} \ln \sum_{\gamma \in G} \rho(\gamma) \exp(-\eta\gamma(\omega))$$

for all $\omega \in \Omega$, where G is a finite subset of Σ_Λ and ρ is a distribution on G.

Let Ξ' be the set of minimal elements of Ξ_Λ: $f \in \Xi'$ if and only if for any $g \in \Xi$ $\forall\omega\,(g(\omega) \le f(\omega))$ implies $f = g$. Notice that Ξ' is contained in the boundary of Ξ_Λ. It is known that the game is η-mixable if and only if $\Xi_\Lambda \subseteq \Sigma_\Lambda$ (which explains the name) if and only if $\Xi' \subseteq \Lambda$.

For $\pi \in \mathcal{P}(\Omega)$ and $g \in [0, \infty]^\Omega$, denote

$$\mathrm{E}_\pi g = \sum_{\omega \in \Omega} \pi(\omega) g(\omega)\,.$$

Assumption 3. Let $\pi \in \mathcal{P}(\Omega)$ be such that $\pi(\omega_1) = 0$ and $\pi(\omega_2) = 0$ for some $\omega_1 \ne \omega_2$. Let $m = \min_{\gamma \in \Xi'} \mathrm{E}_\pi \gamma$. If $\mathrm{E}_\pi \gamma_1 = m$ and $\mathrm{E}_\pi \gamma_2 = m$ for some $\gamma_1, \gamma_2 \in \Xi'$, then either γ_1 minorizes γ_2 or vice versa.

Assumption 3 is rather awkward. But it holds for a wide class of games. In particular, it holds for all binary games and for all proper scoring rules. (But it does not hold, e. g., for non-binary "absolute-loss" game, where $\lambda(\omega, \gamma) = \sum_{i=1}^n |\omega_i - \gamma_i|$.)

On the other hand, some technical requirement of this kind is unavoidable. It does not appear just from our proof: For almost all $\pi \in \mathcal{P}(\Omega)$ there is a unique α_π such that Q defined by (5) is a supermartingale. And it is easy to construct an example where this correspondence cannot be extended to a continuous mapping from $\mathcal{P}(\Omega)$ to Ξ'. (One possible way is to consider the image of Ξ_Λ under point-wise exponential mapping $g \mapsto \mathrm{e}^{-\eta g}$. Every $\pi \in \mathcal{P}(\Omega)$ can be naturally identified with a family of parallel hyperplanes. The inequality (4) actually means that α_π is a point of tangency, where one of the hyperplanes touches the image of Ξ_Λ. To get an example with a point of discontinuity, one may consider a vertical cylinder in the three-dimensional space and a horisontal hyperplane.) It is not clear how to cope with such cases under the supermartingale approach.

Theorem 3. *Let the game (Ω, Λ) satisfy Assumptions 1–3. If the AA is realizable for certain c and η, then there is a mapping $\alpha : \mathcal{P}(\Omega) \to \Sigma_\Lambda$ such that for these α, c, and η, Q defined by (5) is a forecast-continuous supermartingale.*

Proof. We construct the mapping α in two steps. First, we map $\mathcal{P}(\Omega)$ to Ξ', and then we map Ξ' to the boundary of Σ_Λ. Geometrically, these steps amount to taking a tangent plane to Ξ' and the central projection from Ξ' to the boundary of Σ_Λ.

The mapping μ from $\mathcal{P}(\Omega)$ to Ξ' is defined by the formula

$$\mu_\pi = \arg\min_{g \in \Xi'} \mathrm{E}_\pi g \,.$$

First, let us prove that μ_π is well-defined. By Assumption 2, there is a finite point in Λ and thus in Ξ', hence the min is finite. The minimum is attained since Ξ_Λ is compact and $\mathrm{E}_\pi g_1 \le \mathrm{E}_\pi g_2$ if $g_1 \le g_2$. Let us prove that the minimum is attained at one point only[1]. Assume the converse: for some π, there are at least two different points f and g where the minimum is attained, and $\mathrm{E}_\pi f = \mathrm{E}_\pi g = m$. By definition of Ξ, the point $-\ln((\mathrm{e}^{-\eta f(\omega)} + \mathrm{e}^{-\eta g(\omega)})/2)/\eta$ also belongs to Ξ. If it does not belong to Ξ', there is a point $h \in \Xi'$ that minorizes it. For any reals x, y, we have $(\mathrm{e}^x + \mathrm{e}^y)/2 \ge \mathrm{e}^{(x+y)/2}$, and the inequality is strict if $x \ne y$. Therefore, $h(\omega) \le (f(\omega) + g(\omega))/2$, and if there exists ω such that $\pi(\omega) \ne 0$ and $f(\omega) \ne g(\omega)$, then $\mathrm{E}_\pi h < \mathrm{E}_\pi (f+g)/2 = m$. This contradiction proves that if $f(\omega) \ne g(\omega)$, then $\pi(\omega) = 0$. If f and g differ at one ω only, then one of them minorizes the other, thus only one can belong to Ξ'. The remaining case is directly forbidden by Assumption 3.

Let us prove that μ is continuous. Consider any sequence of π_i converging to π. First of all, prove that $\mathrm{E}_{\pi_i}\mu_{\pi_i}$ converges to $\mathrm{E}_\pi \mu_\pi$. Indeed, for any $\pi, \pi' \in \mathcal{P}(\Omega)$, we have $\mathrm{E}_{\pi'}\mu_{\pi'} \le \mathrm{E}_{\pi'}\mu_\pi \le \mathrm{E}_\pi \mu_\pi + \max_\omega \mu_\pi(\omega) \sum_\omega |\pi'(\omega) - \pi(\omega)|$. Hence $|\mathrm{E}_{\pi_i}\mu_{\pi_i} - \mathrm{E}_\pi \mu_\pi| \le \sum_\omega |\pi'(\omega) - \pi(\omega)| \times \max\{\, \mu_\pi(\omega), \mu_{\pi_i}(\omega) \mid \omega \in \Omega \,\}$, and the last expression tends to zero as π_i tends to π. We omitted several subtle points: how to bound μ_{π_i} and how to cope with the case of $\mu_\pi(\omega) = \infty$ (e. g., consider an auxiliary sequence of finite points converging to μ_π).

Now assume a sequence π_i converges to π and μ_{π_i} converges to some γ. By Assumption 1, γ belongs to Ξ_Λ. Further, $\mathrm{E}_{\pi_i}\gamma$ converges to $\mathrm{E}_\pi \gamma$ since their difference is bounded by $\max_\omega \gamma(\omega) \sum_\omega |\pi_i(\omega) - \pi(\omega)|$, and $\mathrm{E}_{\pi_i}(\mu_{\pi_i} - \gamma)$ converges to zero since it is bounded by $\max|\mu_{\pi_i} - \gamma|$. Thus, $\mathrm{E}_{\pi_i}\mu_{\pi_i}$ converges to $\mathrm{E}_\pi \gamma$ and to $\mathrm{E}_\pi \mu_\pi$, and $\gamma = \mu_\pi$ due to the uniqueness of μ_π.

Now let us construct a continuous mapping from Ξ' to Σ_Λ. Let $g \in \Xi_\Lambda$. By definition, it is a positive combination of some $\gamma \in \Sigma_\Lambda$. For $g(\omega) = 0$ for some ω, then $\gamma(\omega) = 0$ too for any γ from the combination. For ω such that $g(\omega) \ne 0$, there is a constant $c > 0$ such that $cg(\omega) \ge \gamma(\omega)$ for all γ from the combination. Taking the maximal such c over all ω, we get that for any $g \in \Xi_\Lambda$ these is a constant $c > 0$ such that $cg \in \Sigma_\Lambda$. Now take the minimal $c > 0$ such that $cg \in \Sigma_\Lambda$, this c will be called $C(g)$ (the minimum is attained due to

[1] In Bayesian framework this uniqueness means that the loss function is a strictly proper scoring rule, cf. [4].

Assumption 1). Clearly, the mapping $g \mapsto C(g)g$ is a continuous mapping from Ξ_Λ (and thus from Ξ') to the boundary of Σ_Λ. If the AA is realizable for some c (and η—recall that Ξ' depends on η), then $C(g) \le c$ for all $g \in \Xi'$.

Let $\alpha_\pi = C(\mu_\pi)\mu_\pi$. Clearly, α_π is continuous and so is Q.

It remains to check that Q is a supermartingale, that is, satisfy (4). Dividing both sides by the right-hand side, we get that it suffices to check the following:

$$\sum_{\omega\in\Omega} \pi(\omega) \exp\left(\eta\left(\frac{\alpha_\pi(\omega)}{c} - \gamma(\omega)\right)\right) \le 1,$$

for any $\gamma \in \Sigma_\Lambda$ and any $\pi \in \mathcal{P}(\Omega)$. This inequality will follow from

$$\sum_{\omega\in\Omega} \pi(\omega) \mathrm{e}^{\eta(\mu_\pi(\omega)-\gamma(\omega))} \le 1 \tag{6}$$

since $\alpha_\pi \le c\mu_\pi$ (here we use that the AA is realizable for c and η). Take $\epsilon > 0$ Consider $-\frac{1}{\eta}\ln\left((1-\epsilon)\mathrm{e}^{-\eta\mu_\pi} + \epsilon \mathrm{e}^{-\eta\gamma}\right)$. By definition, this mixture belongs to Ξ.

When $\epsilon \to 0$, we have

$$\begin{aligned} -\frac{1}{\eta}\ln\left((1-\epsilon)\mathrm{e}^{-\eta\mu_\pi} + \epsilon\mathrm{e}^{-\eta\gamma}\right) &= \mu_\pi - \frac{1}{\eta}\ln\left(1 + \epsilon\left(e^{\eta(\mu_\pi-\gamma)} - 1\right)\right) \\ &= \mu_\pi - \frac{\epsilon}{\eta}\left(e^{\eta(\mu_\pi-\gamma)} - 1\right) + o(\epsilon^2). \end{aligned}$$

Take the expectation E_π of the last expression. If (6) does not hold, this expectation is less than $\mathrm{E}_\pi\mu_\pi$ for sufficiently small ϵ, which contradicts the definition of μ_π. □

Remark 2. For any $g \in \Xi'$, we have $C(g) \ge 1$ since $\Xi_\Lambda \supseteq \Sigma_\Lambda$ and g is minimal in Ξ_Λ. Thus for η-mixable games ($c = 1$) we have $\alpha_\pi = \mu_\pi$ for all π, and the image of the mapping α is included in Λ. For arbitrary games, the image is included in the boundary of Σ_Λ. Note also that α is continuous.

3.3 Defensive Forecasting

Now we describe the *Defensive Forecasting* algorithm (DF) for the game of prediction with expert advice. It has five parameters: reals $c \ge 1$, $\eta > 0$, a function $\alpha : \mathcal{P}(\Omega) \to \Sigma_\Lambda$, a distribution P_0 on Θ, and a substitution function $\sigma : \Sigma_\Lambda \to \Lambda$ such that $\sigma(g) \le g$ for all $g \in \Sigma_\Lambda$.

The parameters c, η, and α are such that the function Q is a forecast-continuous supermartingale. Let

$$Q^{P_0}(\{\gamma_1^\theta\}_{\theta\in\Theta}, \pi_1, \omega_1, \ldots) = \sum_{\theta\in\Theta} P_0(\theta) Q(\gamma_1^\theta, \pi_1, \omega_1, \ldots)\,.$$

Clearly, Q^{P_0} is also a forecast-continuous supermartingale, therefore Lemma 3 applies.

At step N, the DF takes any π_N that satisfy the conclusion of Lemma 3, and then announces $\gamma_N = \sigma(\alpha_{\pi_N})$ as the next prediction of Learner.

Theorem 4. *If Q is a forecast-continuous supermartingale for c, η, and α, then the DF with parameters c, η, α, P_0, and σ guarantees that at each step n for all experts θ*

$$L_n \le cL_n^\theta + \frac{c}{\eta} \ln \frac{1}{P_0(\theta)} .$$

Proof. Lemma 3 guarantees that at each step Q^{P_0} is not greater than its initial value, 1. Thus,

$$\exp\left(\eta \sum_{n=1}^{N} \left(\frac{\alpha_{\pi_n}(\omega_n)}{c} - \gamma_n^\theta(\omega_n) \right)\right) \le \frac{1}{P_0(\theta)} ,$$

and therefore

$$\sum_{n=1}^{N} \alpha_{\pi_n}(\omega_n) \le cL_n^\theta + \frac{c}{\eta} \ln \frac{1}{P_0(\theta)} .$$

It remains to note that $\gamma_n = \sigma(\alpha_{\pi_n}) \le \alpha_{\pi_n}$, and hence $L_n \le \sum_{n=1}^{N} \alpha_{\pi_n}(\omega_n)$. □

As we have seen, the AA and the DF are very close in the loss bound and also in their procedure. We can say even more: with the same parameters and inputs, they give the same predictions. More precisely, two sets coincide: the set of γ_N satisfying (1) and the set of γ_N such that they minorize α_{π_N} for π_N satisfying the conclusion of Lemma 3. Thus, within the standard setting of the prediction with expert advice, the DF is just another way of looking at the AA. In the next section, we consider a setting where these two algorithms differ.

4 Second-Guessing Experts

Let us consider an extension of the protocol of prediction with expert advice. The game is specified by the same elements (Ω, Λ) as above.

PREDICTION WITH SECOND-GUESSING EXPERT ADVICE

$L_0 := 0$.
$L_0^\theta := 0$, $\theta \in \Theta$.
FOR $n = 1, 2, \dots$:
  Experts $\theta \in \Theta$ announce $\gamma_n^\theta : \Lambda \to \Lambda$.
  Learner announces $\gamma_n \in \Lambda$.
  Reality announces $\omega_n \in \Omega$.
  $L_n := L_{n-1} + \gamma_n(\omega_n)$.
  $L_n^\theta := L_{n-1}^\theta + \gamma_n^\theta(\gamma_n)(\omega_n)$.
END FOR.

The new protocol contains only one substantial change. Informally speaking, now the experts announce not their actual predictions, but conditional statements that specify their predictions depending on Learner's next step. Therefore, the loss of each expert is determined by the prediction of Learner as well as by the outcome chosen by Reality. We will call the experts in this protocol *second-guessing experts*. Second-guessing experts are a generalization of experts in the standard protocol: a standard expert can be interpreted in the new protocol as a constant function.

The phenomenon of "second-guessing experts" occurs in real-world finance. For example, commercial banks serve as "second-guessing experts" for a central bank when they use variable interest rates (that is, the interest rate for the next period is announced as an explicit function of the central bank base rate).

In game theory, there is a notion of internal regret [6], which is related to the idea of second-guessing experts. The internal regret appears in the framework where for each prediction, which is called action in this context, there is an expert that consistently recommends this action, and Learner follows one of the experts at each step. The internal regret for a pair of experts (i, j) shows by how much Learner could decrease its loss by having followed expert j each time it followed expert i. This can be modelled by a second-guessing expert that agrees with Learner if Learner does not follow i, and recommends following j when the Learner follows i.

The internal regret is studied in randomized prediction protocols. In our case of deterministic Learner's predictions, one cannot hope to get any interesting loss bound without additional assumptions. Indeed, Experts can always suggest exactly the "opposite" to the Learner's prediction (for example, in the log loss game, predict 0 if Learner predicts $\gamma_n \leq 0.5$ and 1 otherwise), and Reality can "agree" with them; then the Experts' losses remain zero, but the Learner's loss grows linearly in the number of steps. In this paper we consider second-guessing experts that depend on the prediction of Learner *continuously*.

4.1 The DF for Second-Guessing Experts

First consider the case when γ_n^θ are continuous mappings from Λ to Λ. Surprisingly, the DF requires virtually no modifications.

Theorem 5. *Suppose Q defined by (5) is a forecast-continuous supermartingale for some c, η, and a continuous $\alpha : \mathcal{P}(\Omega) \to \Lambda$. Then a DF algorithm can be applied in the protocol of prediction with second-guessing expert advice for continuous experts; it guarantees the same loss bound as in the prediction with expert advice protocol.*

Proof. Let Q be a forecast-continuous supermartingale as a function on $(\Sigma_\Lambda \times \mathcal{P}(\Omega) \times \Omega)^*$. Let γ_n^θ be continuous functions $\Lambda \to \Lambda$. Since α is continuous, $\gamma_n^\theta(\alpha_\pi)$ is a continuous function $\mathcal{P}(\Omega) \to \Lambda$. Then Q with $\gamma_n^\theta(\alpha_\pi)$ substituted for e_n is a forecast-continuous supermartingale as a function on $(C(\Lambda \to \Lambda) \times \mathcal{P}(\Omega) \times \Omega)^*$, where $C(\Lambda \to \Lambda)$ is the set of continuous functions on Λ. We take the identity for a substitution function since α takes values already in Λ, and for the DF with the supermartingale above the proof of Theorem 4 applies. □

Recall that for mixable games the mapping α constructed in the proof of Theorem 3 satisfies the conditions of the last theorem.

For games that are not mixable, it may happen that such α does not exist. Even more, for games where Λ is not connected (e. g., the simple game of prediction), the continuity of experts does not rule out the example with “opposite” predictions.

To cope with such cases, we modify the protocol of the game. Namely, we allow Experts and Learner to announce predictions from the boundary Σ'_Λ of Σ_Λ. That is, Experts are $\gamma_n^\theta : \Sigma'_\Lambda \to \Sigma'_\Lambda$, and Learner is $\gamma_n \in \Sigma'_\Lambda$. Theorem 5 requires minimal changes: now α is a function from $\mathcal{P}(\Omega)$ to Σ'_Λ, and the proof is modified accordingly. Theorem 3 supplies us with the α required.

4.2 The AA for Second-Guessing Experts

The AA cannot be applied to the second-guessing protocol directly. Recall the AA is based on the inequality (1), which is already resolved for γ_N. In the second-guessing protocol, the inequality contains γ_N on both sides:

$$\gamma_N(\omega_N) \le -\frac{c}{\eta} \ln \sum_{\theta\in\Theta} \frac{P_{N-1}(\theta)}{\sum_{\theta\in\Theta} P_{N-1}(\theta)} \exp(-\eta\gamma_N^\theta(\gamma_N)(\omega_N)) . \tag{7}$$

The DF implicitly solves this inequality within (the proof of) Lemma 3, which is a kind of fixed point theorem. We present a modification of the AA which uses a fixed point theorem explicitly. We use the following theorem (see e. g. [1]).

Theorem 6 (Brouwer). *If X is homeomorphic to a closed simplex, and $F : X \to X$ is a continuous function, then F has a fixed point.*

Our goal is to find a subset X of possible Learner’s predictions such that X is homeomorphic to a closed simplex, and a continuous function $G : X \to X$ such that for any $\gamma \in X$, the point $G(\gamma)$ is not greater than the right-hand side of (7) with γ substituted for γ_N. Then the modified AA works as follows: at each step, the AA constructs G and X (they may depend on the step and history), finds a fixed point $\gamma = G(\gamma)$ via the Brouwer theorem, and announces γ as the next prediction of Learner. Since any fixed point of G is a solution of (7), the standard analysis of the AA applies providing the same loss bound.

We construct the domain X and the function G in a different manner for mixable games with the original second-guessing protocol and for non-mixable games and the modified protocol.

Obviously, Ξ' is homeomorphic to a closed simplex (note that the “exponential image” $\mathrm{e}^{-\eta\Xi_\Lambda}$ of Ξ_Λ is a bounded convex subset of $\mathbb{R}^{|\Omega|}$). For mixable games, $\Xi' \subseteq \Lambda$, and we let $X = \Xi'$ in this case. For non-mixable games, we consider the mapping $g \mapsto C(g)g$ used in the proof of Theorem 3, which is a homeomorphism of Ξ' to a part of Σ'_Λ. This part we take as X.

Now let us construct the function G. The beginning is common for both versions. A point $\gamma \in X$ is mapped to the point g such that

$$g(\omega) = -\frac{1}{\eta} \ln \sum_{\theta\in\Theta} \frac{P_{N-1}(\theta)}{\sum_{\theta\in\Theta} P_{N-1}(\theta)} \exp(-\eta\gamma_N^\theta(\gamma)(\omega))$$

for all ω. The point g belongs to Ξ_Λ by definition.

Lemma 4. *There is a continuous mapping* $F : \Xi_\Lambda \to \Xi'$ *such that for any* $g \in \Xi_\Lambda$*, it holds that* $F(g) \le g$*.*

We postpone the proof and continue the construction of G. This F maps g to $F(g) \in \Xi'$. For mixable games, we are done. For non-mixable games, we need one more step: apply to $F(g)$ the homeomorphism from Ξ' to Σ'_Λ. The function G has the correct range by construction and is continuous as a composition of continuous mappings. The point $G(\gamma)$ is not greater than the right-hand side of (7) since $C(F(g))F(g) \le cF(g) \le cg$. Thus, we obtained the following

Theorem 7. *If the AA is realizable for the game* (Ω, Λ) *in the prediction with expert advice for some* c *and* η*, then the AA with fixed point is realizable for the same game for the prediction with second-guessing expert advice protocol (modified—for non-mixable games) for the same* c *and* η*, and guarantees the same loss bounds.*

Proof of Lemma 4. We construct a continuous mapping $F : \Xi_\Lambda \to \Xi'$ as a composition of mappings F_ω for all $\omega \in \Omega$. Each F_ω when applied to $g \in \Xi_\Lambda$ preserves the values of $g(o)$ for $o \ne \omega$ and decreases as far as possible the value $g(\omega)$ so that the result is still in Ξ_Λ. Formally, $F_\omega(g) = g'$ such that $g'(o) = g(o)$ for $o \ne \omega$ and $g'(\omega) = \min\{ f(\omega) \mid f \in \Xi_\Lambda, \forall o \ne \omega\ f(o) = g(o) \}$.

Let us show that F_ω is continuous. It suffices to show that $F_\omega(g)(\omega)$ depends continuously on g, since the other coordinates do not change. We will show that $F_\omega(g)(\omega)$ is concave in g, continuity follows easily (see, e. g. [11]). Indeed, take any $t \in [0,1]$, and $f, g \in \Xi_\Lambda$. Since Ξ_Λ is convex (and even exp-convex), then $tf + (1-t)g \in \Xi_\Lambda$ and $tF_\omega(f) + (1-t)F_\omega(g) \in \Xi_\Lambda$. The latter point has all the coordinates $o \ne \omega$ the same as the former. Thus, by definition of F_ω, we get $F_\omega(tf+(1-t)g)(\omega) \le (tF_\omega(f)+(1-t)F_\omega(g))(\omega) = tF_\omega(f)(\omega)+(1-t)F_\omega(g)(\omega)$, which was to be shown.

All F_ω do not increase the coordinates. Since the set Ξ_Λ contains any point g with all its majorants, $F_\omega(f) = f$ implies that $F_\omega(g) = g$ for any g obtained from f by applying any $F_{\omega'}$. Therefore, the image of a composition of F_ω over all $\omega \in \Omega$ is included in Ξ'. □

Remark 3. Actually, Lemma 4 constructs a continuous substitution function. In many natural games, the standard substitution functions appear to be continuous. In particular, for the log loss function,

$$(g_0, g_1) \mapsto ((-g_0 + g_1)\ln(\mathrm{e}^{-g_0+g_1} + 1), (-g_0 + g_1)\ln(\mathrm{e}^{g_0-g_1} + 1))\,.$$

For the quadratic loss function,

$$(g_0, g_1) \mapsto \left(\left(\frac{1+g_0-g_1}{2}\right)^2, \left(\frac{1-g_0+g_1}{2}\right)^2\right)\,.$$

Acknowledgements

This work was partly supported by EPSRC grant EP/F002998/1. Comments by Alex Gammerman, Glenn Shafer, Alexander Shen, and the anonymous referees have helped us improve the presentation.

References

1. R. Agarwal, M. Meehan, D. O'Regan. *Fixed Point Theory and Applications*, volume 141 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, England, 2001.
2. A. Blum, Y. Mansour. From External to Internal Regret. *J. Mach. Learn. Res.*, 8:1307–1324, 2007.
3. N. Cesa-Bianchi, G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, Cambridge, England, 2006.
4. A. P. Dawid. The geometry of proper scoring rules. *Annals of the Institute of Statistical Mathematics*, 59:77–93, 2007.
5. D. Foster, R. Vohra. Asymptotic calibration. *Biometrika*, 85:379–390, 1998.
6. D. Foster, R. Vohra. Regret in the online decision problem. *Games Econ. Behav.*, 29:104–130, 1999.
7. P. Gács. Uniform test of algorithmic randomness over a general space. *Theoretical Computer Science*, 341:91–137, 2005.
8. W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58:13–30, 1963.
9. L. Levin. Uniform tests of randomness. *Soviet Mathematics Doklady*, 17:337–340, 1976.
10. M. Li, P. Vitányi, *An Introduction to Kolmogorov Complexity and Its Applications*, 2nd edition, New York: Springer, 1997.
11. R. Rockafellar. *Convex Analysis.* Princeton University Press, 1996.
12. G. Shafer, V. Vovk. *Probability and Finance: It's Only a Game!* Wiley, New York, 2001.
13. G. Stoltz, G. Lugosi. Internal Regret in On-Line Portfolio Selection. *Machine Learning*, 59:125–159, 2005.
14. G. Stoltz, G. Lugosi. Learning correlated equilibria in games with compact sets of strategies. *Games and Economic Behavior*, 59:187–209, 2007.
15. V. Vovk. Aggregating Strategies. In M. Fulk, J. Case, editors, *Proceedings of the Third Annual Workshop on Computational Learning Theory*, pp. 371–383, San Mateo, CA, 1990.
16. V. Vovk. A game of prediction of expert advice. *Journal of Computer and System Sciences*, 56:153-173, 1998
17. V. Vovk. Competitive on-line learning with a convex loss function. Technical Report `arXiv:cs/0506041v3` [cs.LG], `arXiv.org` e-Print archive, September 2005.
18. V. Vovk. On-line regression competitive with reproducing kernel Hilbert spaces. Technical Report `arXiv:cs/0511058v2` [cs.LG], `arXiv.org` e-Print archive, January 2006.
19. V. Vovk. Metric entropy in competitive on-line prediction. Technical Report `arXiv:cs/0609045v1` [cs.LG], `arXiv.org` e-Print archive, September 2006.
20. V. Vovk. Continuous and randomized defensive forecasting: unified view. Technical Report `arXiv:0708.2353v2` [cs.LG], `arXiv.org` e-Print archive, August 2007.