Universal Portfolio Selection

V. Vovk and C. Watkins
Department of Computer Science
Royal Holloway, University of London
Egham, Surrey TW20 0EX, UK
{vovk,chrisw}@dcs.rhbnc.ac.uk

Abstract

A typical problem in portfolio selection in stock markets is that it is not clear which of the many available strategies should be used. We apply a general algorithm of prediction with expert advice (the Aggregating Algorithm) to two different idealizations of the stock market. One is the well-known game introduced by Cover in connection with his "universal portfolio" algorithm; the other is a more realistic modification of Cover's game introduced in this paper, where market's participants are allowed to take "short positions", so that the algorithm may be applied to currency and futures markets. Besides applying the Aggregating Algorithm to a countable (or finite) family of arbitrary investment strategies, we also apply it, in the case of Cover's game, to the uncountable family of "constant rebalanced portfolios" considered by Cover. We generalize Cover's worst-case bounds for his "universal portfolio" algorithm (which can be regarded as a special case of the Aggregating Algorithm corresponding to learning rate 1) to the case of learning rates not exceeding 1. Finally, we discuss a general approach to designing investment strategies in which, instead of making statistical or other assumptions about the market, natural assumptions of computability are made about possible investment strategies; this approach leads to natural extensions of the notion of Kolmogorov complexity.

1 INTRODUCTION

In recent years, following Littlestone and Warmuth [11], a series of papers devoted to "prediction with expert advice" have appeared. (In this paper, we will usually speak of "actions" instead of "predictions".) Littlestone and Warmuth considered the simplest situation of binary actions and outcomes with the loss function being the number of mistakes made and proposed the Weighted Majority Algorithm for merging several decision strategies (the same algorithm was independently proposed in Vovk [13]). Later the Weighted Majority Algorithm was developed in (at least) two different directions. Vovk [12] (see also [14]) generalized it to a wide class of outcome and action spaces and loss functions. This general algorithm, which we will call the Aggregating Algorithm (AA), was shown to be in a certain sense optimal (see Haussler et al. [5], Vovk [12, 14]). On the other hand, Littlestone et al. [10] and Kivinen and Warmuth [8] developed the Weighted Majority Algorithm in a different direction, introducing what we will call EG-type algorithms.

An interesting special case of the AA is Cover's (see, e.g., Cover and Ordentlich [2]) universal portfolio selection algorithm. The AA, like the Weighted Majority Algorithm and EG-type algorithms, depends on a parameter called "learning rate" and usually denoted by η; $0 < \eta < \infty$. Cover's algorithm coincides with the AA applied to what we call "Cover's game" and the pool of constant rebalanced portfolios with learning rate $\eta = 1$. The usual methods of proving worst-case bounds for the AA give the best results for $\eta = 1$, and it seemed evident that one should take $\eta = 1$. However, in Helmbold et al.'s experiments [6] it was found that the values $\eta = 0.01$ to $\eta = 0.15$ are better than $\eta = 1$ for their EG-type algorithm; moreover, their algorithm with the learning rate in this range outperformed Cover's algorithm. Since the η of Helmbold et al.'s algorithm and the AA's η are of the same provenance, one might suspect that the cause of the inferior performance of Cover's algorithm is a wrong learning rate. This suggests that learning rates different from 1 might also be interesting even for Cover's game, and so we consider general η.

The main contribution of this paper is that, besides Cover's game (Section 3), we consider (in Section 4) a more realistic, in some respects, game, where the market's participants are allowed to take short positions. (This game is applicable, e.g., to investment in currencies and futures, as well as stock, markets.)

COLT 98 Madison WI USA

Cover and Ordentlich [2] are mainly interested in an infinite pool of experts which correspond to constant rebalanced portfolios (for exact definitions, see below). It is not clear why the constant rebalanced portfolios are a good choice (they typically involve a lot of trading, ignore the non-stationarity of markets, etc.), and we are more interested in applying the AA to countable or finite families of investment strategies. Such a family might consist, e.g., of all investment strategies that the investor believes to be reasonable. For a further motivation, see Section 5.

However, we still prove (in Section 3) a generalization of Cover's result for Cover's game and the pool of constant rebalanced portfolios; namely, we obtain a bound of $\frac{N-1}{2\eta}\ln T$ on the extra loss of the algorithm as compared to the best constant rebalanced portfolio (N is the number of stocks and T is the number of trials). For $\eta = 1$ our bound becomes Cover and Ordentlich's [2] bound of $\frac{N-1}{2}\ln T$, whereas Helmbold et al.'s [6] bounds are much worse (for fixed η, their bound on the extra loss of their algorithm is linear in T).

Our use of the word "universal" in the title of this paper is traditional: for example, it is widely used in discussions of Cover's algorithm of merging constant rebalanced portfolios (see, e.g., [2]) and in information theory (see, e.g., [3]). Our preferred interpretation of it is, however, in terms of the theory of Kolmogorov complexity (for an excellent survey of this theory, see Li and Vitányi [9]).

The basic idea of the universal approach to investing may be described as follows. The AA can merge any countable families of investment strategies, which suggests a possibility of applying Solomonoff's idea of universal prediction (see, e.g., [9], Section 5.2). Consider a universal Turing machine $\mathcal{M}$: when supplied with a "program" i, $\mathcal{M}$ simulates Turing machine $\mathcal{M}_i$, and every Turing machine is among the $\mathcal{M}_i$. Every computable investment strategy is computed by some Turing machine $\mathcal{M}_i$, and so we can obtain a universal investment strategy by merging all $\mathcal{M}_i$. The resulting universal strategy will not be computable in the usual sense (though it will be computable "in the limit"), so we need to use some approximations to it; in the usual theory of Kolmogorov complexity such approximations were considered, e.g., by Levin (universal search procedures), Rissanen (MDL principle) and Wallace (MML principle).

In Section 5 we elaborate the idea of universal investment. In this paper we are mainly interested in abstract setting in which the learner's computational resources (such as the allowed computation time) are unlimited and arrive at two modifications of Kolmogorov complexity. The usual Kolmogorov complexity (more accurately, its variant KM, see [9], Subsection 4.5.4) can be interpreted as an intrinsic measure of helpfulness of the learner's adversary in the log-loss game. One of our modifications ("Cover complexity") is a measure of the market's helpfulness in Cover's game; the log-loss game is a restriction of Cover's game and, correspondingly, Kolmogorov complexity is a restriction of Cover complexity (up to a scaling factor). The other modification is a measure of the market's helpfulness in the long-short game.

2 GENERAL AA

The AA is an algorithm which the learner can use to choose an action based on the suggestions from a pool Θ of experts; the set Θ is equipped with a σ-algebra. We will consider the following perfect-information game between three participants, Pool, Learner, and Nature:

FOR $t = 1, 2, \ldots$
  Pool chooses a measurable function $\gamma_t : \Theta \to \Gamma$
  Learner chooses $\gamma_t(\text{Learner}) \in \Gamma$
  Nature chooses $\omega_t \in \Omega$
END FOR.

In this description, Γ is a fixed *action space* (Γ is assumed to be a topological space equipped with the σ-algebra generated by the open sets), $\gamma_t(\theta)$ is the action chosen by expert θ at time t, and Ω is a fixed *outcome space*. We also fix a *loss function* $\lambda : \Omega \times \Gamma \to \mathbb{R}$; the total loss suffered by expert θ over the first T trials is

$$\text{Loss}_T(\theta) := \textstyle\sum_{t=1}^{T} \lambda(\omega_t, \gamma_t(\theta)),$$

and the total loss suffered by Learner over the first T trials is

$$\text{Loss}_T(\text{Learner}) := \textstyle\sum_{t=1}^{T} \lambda(\omega_t, \gamma_t(\text{Learner})).$$

We do not fix the goal of the game; informally, Learner aims at performing almost as well as the best expert. The AA provides a possible strategy for Learner in this game.

First we fix a learning rate $\eta > 0$, put $\beta = e^{-\eta}$, and fix a probability distribution P_0 in the pool Θ; the *prior distribution* P_0 specifies the initial weights assigned to the experts. Besides choosing η and P_0, we also need to specify a "substitution function" in order to be able to apply the AA. A *generalized action* is defined to be any function of the type $\Omega \to \mathbb{R}$ and a *substitution function* is a function Σ that maps every generalized action $g : \Omega \to \mathbb{R}$ into a "permitted action" $\Sigma(g) \in \Gamma$. A permitted action $\gamma \in \Gamma$ is identified with the generalized action g defined by $g(\omega) := \lambda(\omega, \gamma)$. (Note that, abstractly, an action may be identified with its loss function. In the context of portfolio selection, a permitted action is a loss function that can be achieved by a feasible portfolio.) Later we will describe which substitution functions are allowed in the AA; let us assume that we are given some substitution function Σ. Now we have all we need to describe how the AA works.

At every trial $t = 1, 2, \ldots$ Learner updates the experts' weights as follows:

$$P_t(A) := \textstyle\int_A \beta^{\lambda(\omega_t, \gamma_t(\theta))} P_{t-1}(d\theta),$$

where $A \subseteq \Theta$ is a measurable set; P_0 is the prior distribution. (Therefore, the larger the loss $\lambda(\omega_t, \gamma_t(\theta))$ the more sharply the weight of expert θ decreases.) The action chosen by the AA at trial t is obtained from the weighted average of the experts' actions by applying the substitution function: $\gamma_t := \Sigma(g_t)$, where the generalized action g_t is defined by

$$g_t(\omega) := \log_\beta \int_\Theta \beta^{\lambda(\omega,\gamma_t(\theta))} P^*_{t-1}(d\theta) \qquad (1)$$

and P^*_{t-1} are the normalized weights, $P^*_{t-1}(d\theta) := P_{t-1}(d\theta)/P_{t-1}(\Theta)$ (assuming that the denominator is positive; if it is 0, P_0-almost all experts have suffered infinite loss and, therefore, the AA is allowed to choose any action).

When Learner follows $\mathrm{AA}(\eta, P_0)$ (i.e., AA with learning rate η and prior P_0) we will write

$$\mathrm{Loss}_T(\mathrm{AA}(\eta, P_0))$$

in place of Loss_T(Learner). We will also speak of *Aggregating Pseudo-Algorithm* (APA), which makes not permitted actions but generalized actions, in accordance with (1); we will use the self-evident notation $\mathrm{APA}(\eta, P_0)$ and also the notation

$$\mathrm{Loss}_T(\mathrm{APA}(\eta, P_0)) := \textstyle\sum_{t=1}^T g_t(\omega_t),$$

where g_t is the generalized action output by $\mathrm{APA}(\eta, P_0)$ at trial t.

All our proofs are based on the following property of the APA (see [14]).

Lemma 1 *For any learning rate $\eta > 0$, prior P_0, and $T = 1, 2, \ldots$,*

$$\mathrm{Loss}_T(\mathrm{APA}(\eta, P_0)) = \log_\beta \textstyle\int_\Theta \beta^{\mathrm{Loss}_T(\theta)} P_0(d\theta).$$

To finish the description of the AA, we only need to specify the admissible substitution functions Σ. Let $\mathrm{GA}(\eta)$ be the set of all possible generalized actions that can be output by the APA with learning rate η (see (1)); in other words,

$$\mathrm{GA}(\eta) := \left\{g : \Omega \to \mathbb{R} \mid \exists P \forall \omega : g(\omega) = \log_\beta \textstyle\int_\Gamma \beta^{\lambda(\omega,\gamma)} P(d\gamma)\right\},$$

where $\beta := e^{-\eta}$ and P ranges over the set of all probability distributions in Γ. Notice that $\mathrm{GA}(\eta)$ depends only on the loss function λ and does not depend on the pool of experts.

Let us assume that λ takes only nonnegative values. Notice that in this case every $g \in \mathrm{GA}(\eta)$ also takes only nonnegative values. For any such generalized action g we define

$$c(g) = \inf_{\gamma\in\Gamma} \sup_{\omega\in\Omega} \tfrac{\lambda(\omega,\gamma)}{g(\omega)}$$

(with convention $\frac{0}{0} := 0$); we also put

$$c(\eta) := \sup_{g\in \mathrm{GA}(\eta)} c(g).$$

Our assumptions about the game $(\Omega, \Gamma, \lambda)$ will be the same as in [14]:

- Γ is a compact topological space.
- For each ω, the function $\gamma \mapsto \lambda(\omega, \gamma)$ is continuous.
- There exists γ such that, for all ω, $\lambda(\omega, \gamma) < \infty$.
- There exists no γ such that, for all ω, $\lambda(\omega, \gamma) = 0$.

Under these assumptions, there exists a substitution function $\Sigma = \Sigma_\eta$ such that

$$\forall g \in \mathrm{GA}(\eta) \forall \omega \in \Omega : \lambda(\omega, \Sigma(g)) \le c(\eta) g(\omega) \qquad (2)$$

(see Subsection 7.1 below). *The allowed substitution functions in the AA are those satisfying (2).* Notice that, according to Lemma 1 and (2),

$$\mathrm{Loss}_T(\mathrm{AA}(\eta, P_0)) \le c(\eta) \log_\beta \int_\Theta \beta^{\mathrm{Loss}_T(\theta)} P_0(d\theta). \quad (3)$$

We will say that the game $(\Omega, \Gamma, \lambda)$ is *η-mixable* if we can take $c(\eta) = 1$ in (2) (notice that always $c(\eta) \ge 1$). For such games the requirement that the loss function should be nonnegative is superfluous; even when the loss function is allowed to take negative values, we will always have

$$\mathrm{Loss}_T(\mathrm{AA}(\eta, P_0)) \le \log_\beta \int_\Theta \beta^{\mathrm{Loss}_T(\theta)} P_0(d\theta). \qquad (4)$$

Cover's game and the long-short game are η-mixable if and only if $\eta \le 1$. *When discussing η-mixable games, we will never make the assumption that the loss function is nonnegative.*

When Θ is a countable or finite set, (3) and (4) imply the following lemma showing that the AA performs not much worse than any expert. (A countable or finite Θ will always be equipped with the σ-algebra of all subsets of Θ.)

Lemma 2 *If the pool Θ of experts is countable or finite,*

$$\mathrm{Loss}_T(\mathrm{AA}(\eta, P_0)) \le c(\eta)\, \mathrm{Loss}_T(\theta) + a(\eta) \ln \tfrac{1}{P_0\{\theta\}},$$

where $a(\eta) := c(\eta)/\eta$. In the case of an η-mixable game,

$$\mathrm{Loss}_T(\mathrm{AA}(\eta, P_0)) \le \mathrm{Loss}_T(\theta) + \tfrac{1}{\eta} \ln \tfrac{1}{P_0\{\theta\}}.$$

(To prove this lemma it suffices to replace the sum, which is represented by the integral sign, in (3) and (4) with one addend.)

3 COVER'S GAME

Learner is investing in a market of N stocks. The behavior of the market at trial t is described by a nonnegative price relative vector $\omega_t = (\omega_t[0], \ldots, \omega_t[N-1]) \in [0, \infty)^N$. The entry $\omega_t[n]$, $n = 0, 1, \ldots, N-1$, of the tth price relative vector ω_t denotes the ratio of day t closing price to day $t-1$ closing price of the stock n; we will usually assume that at least one of $\omega_t[n]$ is positive. An investment at time t in this market is specified by a portfolio vector $\gamma_t \in [0,1]^N$ with nonnegative entries $\gamma_t[n]$, $n = 0, \ldots, N-1$, summing to 1: $\gamma_t[0] + \cdots + \gamma_t[N-1] = 1$. The entries of γ_t are the proportions of the current wealth invested in each stock at time t. An investment using portfolio γ increases the investor's wealth by a factor of $\gamma \cdot \omega = \sum_{n=0}^{N-1} \gamma[n]\omega[n]$ if the market performs according to the price relative vector $\omega = \omega[0] \ldots \omega[N-1]$. It is natural to define the loss function to be the minus logarithm of this increase:

$$\lambda(\omega, \gamma) := -\ln(\gamma \cdot \omega). \qquad (5)$$

The AA can only deal with a non-negative loss function (unless the latter is η-mixable) and it is usually best to apply it to loss functions satisfying $\inf_\gamma \lambda(\omega,\gamma)=0$ for all ω, so we "normalize" loss function (5) considering "regrets"

$$\lambda^*(\omega,\gamma) := \lambda(\omega,\gamma) - \min_\delta \lambda(\omega,\delta) = \ln \frac{\|\omega\|_\infty}{\gamma\cdot\omega}, \qquad (6)$$

where $\|\omega\|_\infty := \max_n \omega[n]$. It is easy to see that considering loss function (6) is equivalent to considering the initial loss function (5) with the additional restriction that $\|\omega\|_\infty = 1$.

The next lemma gives $c(\eta)$ and the AA's actions for the game just defined (*Cover's game*) and $\eta \le 1$.

Lemma 3 *For every $\eta \le 1$, $c(\eta)=1$. Moreover, for every $\eta \le 1$ and every $g \in \mathrm{GA}(\eta)$, $c(g)=1$. The only action attaining $c(g)=1$ is the average*

$$\gamma^* := \int_\Gamma \gamma P(d\gamma), \qquad (7)$$

where P is a probability distribution in Γ generating g:

$$g(\omega) = \log_\beta \textstyle\int_\Gamma \beta^{\lambda(\omega,\gamma)} P(d\gamma),\ \forall\omega,$$

with $\beta := e^{-\eta}$.

This lemma is proven in Subsection 7.2 below.

Now we can give a relatively explicit description of the AA for $\eta \le 1$:

Algorithm 1 *If $\eta \le 1$, $\mathrm{AA}(\eta, P_0)$'s action γ_T at trial T is*

$$\gamma_T = \frac{\int_\Gamma \delta \prod_{t=1}^{T-1} (\delta\cdot\omega_t)^\eta P_0(d\delta)}{\int_\Gamma \prod_{t=1}^{T-1} (\delta\cdot\omega_t)^\eta P_0(d\delta)}.$$

(Notice that Algorithm 1 becomes Cover's algorithm when $\eta = 1$.)

It seems that the case $\eta > 1$ is less interesting than the case $\eta \le 1$; for the former we only state the following simple result, which shows that $c(\eta) > 1$ when $\eta > 1$.

Lemma 4 *When $\eta > 1$, $c(\eta) = \eta$.*

(For a proof, see Subsection 7.3 below.) Lemmas 2, 3 and 4 immediately imply

Theorem 1 *If the pool Θ of experts is countable or finite, the performance of $\mathrm{AA}(\eta,P_0)$ will satisfy, for any T and $\theta\in\Theta$,*

$$\mathrm{Loss}_T(\mathrm{AA}(\eta,P_0)) \le \mathrm{Loss}_T(\theta) + \frac{1}{\eta}\ln\frac{1}{P_0\{\theta\}} \qquad (8)$$

if $\eta \le 1$, and

$$\mathrm{Loss}_T(\mathrm{AA}(\eta,P_0)) \le \eta\,\mathrm{Loss}_T(\theta) + \ln\frac{1}{P_0\{\theta\}} \qquad (9)$$

if $\eta > 1$. In (8) the loss function is either $-\ln(\gamma\cdot\omega)$ or $\ln\frac{\|\omega\|_\infty}{\gamma\cdot\omega}$; in (9) the loss function is $\ln\frac{\|\omega\|_\infty}{\gamma\cdot\omega}$.

We can see that the best bound obtains when $\eta = 1$.

In conclusion of this section, we state a generalization of Cover's result for Cover's game, $\eta \le 1$ and the pool of constant rebalanced portfolios (for a proof, see Subsection 7.4 below). The later pool is defined to be $\Theta = \Gamma$; expert γ's action is always γ. Notice that expert γ's loss is the minus logarithm of the wealth attained by using the same portfolio γ starting with a unit capital. (Expert γ's strategy is called a constant rebalanced portfolio strategy; it actually involves a great deal of trading.)

Theorem 2 *For any learning rate $\eta \le 1$ and number of stocks N there exists a constant $c = c(\eta, N)$ such that always*

$$\mathrm{Loss}_T(\mathrm{AA}(\eta,P_0)) \le \inf\nolimits_{\gamma\in\Gamma} \mathrm{Loss}_T(\gamma) + \tfrac{N-1}{2\eta}\ln T + c, \quad T = 1,2,\ldots,$$

where P_0 is the Dirichlet measure with parameters $\frac12,\ldots,\frac12$. Here the loss function is either $-\ln(\gamma\cdot\omega)$ or $\ln\frac{\|\omega\|_\infty}{\gamma\cdot\omega}$.

Notice that the best bound again corresponds to $\eta = 1$ (Cover's case). This is because this bound is deduced in two steps: we prove that the APA's performance satisfies it and then make use of the fact that $c(g) = 1$ for any generalized action g. For $\mathrm{AA}(1) = \mathrm{APA}(1)$ nothing is lost, but for $\eta < 1$ we ignore the possibility that at the realized outcome ω the permitted action $\Sigma(g)$ can be much better than the generalized action g.

4 LONG-SHORT GAME

Here we consider a modification of Cover's portfolios. Again Learner is investing in a market of N stocks, but the behavior of the market at trial t is described by a vector of returns $\omega_t \in [-1,\infty)^N$ (notice the difference from Cover's game). The entry $\omega_t[n]$, $n = 0,\ldots,N-1$, of the tth vector of returns ω_t denotes the ratio of the day t increase in the price of stock n (i.e., of the difference between its day t closing price and its day $t-1$ closing price) to its day $t-1$ closing price. (The minimal value of this ratio is -1.) An investment at time t in this market is specified by a portfolio vector $\gamma_t \in \mathbb{R}^N$ with entries $\gamma_t[n]$ satisfying

$$\|\gamma_t\|_1 := |\gamma_t[0]| + \cdots + |\gamma_t[N-1]| \le a, \qquad (10)$$

where $a > 0$ is a constant ("prudence coefficient") reflecting how much of her capital Learner is willing to jeopardize. The entries of γ_t are the proportions of the current wealth invested in each stock at time t; $\gamma_t[n]$ being negative means a short position in stock n. An investment using portfolio γ increases the investor's wealth by a factor of $1 + \gamma\cdot\omega = 1 + \sum_{n=0}^{N-1}\gamma[n]\omega[n]$ if the market performs according to the vector of returns $\omega = \omega[0]\ldots\omega[N-1]$. A natural loss function is the minus logarithm of this increase:

$$\lambda(\omega,\gamma) := -\ln(1+\gamma\cdot\omega),$$

but again we will also be interested in Learner's "regret"

$$\lambda^*(\omega,\gamma) := \ln\tfrac{1+a\|\omega\|_\infty}{1+\gamma\cdot\omega}, \qquad (11)$$

where $\|\omega\|_\infty := \max_n |\omega[n]|$ (notice that $-\ln(1 + a\|\omega\|_\infty)$ is the loss Learner would have suffered had he known ω *a priori*). The loss function λ^* is non-negative and satisfies $\inf_\gamma \lambda^*(\omega,\gamma) = 0$ for all ω; it is obtained by "normalizing" λ:

$$\lambda^*(\omega,\gamma) = \lambda(\omega,\gamma) - \inf_\delta \lambda(\omega,\delta).$$

We will make the optimistic assumption that always $\|\omega_t\|_\infty \le A$, where $0 \le A < \frac{1}{a}$; this assumption ensures that Learner will never go bust.

In the next lemma (which is proven in Subsection 7.5 below) we compute $c(\eta)$ and the AA's actions for the long-short game.

Lemma 5 *For every $\eta \le 1$, $c(\eta) = 1$. Moreover, for every $\eta \le 1$ and every $g \in \mathrm{GA}(\eta)$, $c(g) = 1$. The only action attaining $c(g) = 1$ is the average (7), where, as before, P is a probability distribution in Γ generating g. When $\eta > 1$, $c(\eta) > 1$.*

Notice that (7) is indeed an action: the convexity of the absolute value function implies that a linear mixture of portfolios satisfying (10) again satisfies (10).

A more explicit description of the AA for the long-short game with $\eta \le 1$ is as follows:

Algorithm 2 $\mathrm{AA}(\eta, P_0)$*'s action γ_T at trial T is*

$$\gamma_T = \frac{\int_\Gamma \delta \prod_{t=1}^{T-1} (1 + \delta \cdot \omega_t)^\eta \, P_0(d\delta)}{\int_\Gamma \prod_{t=1}^{T-1} (1 + \delta \cdot \omega_t)^\eta \, P_0(d\delta)}.$$

Now we state an analog of Theorem 1, which immediately follows from Lemmas 2 and 5, for the long-short game.

Theorem 3 *For the long-short game, $\eta \le 1$, countable or finite Θ and any P_0 we have*

$$\forall \theta \in \Theta \forall T :$$
$$\mathrm{Loss}_T(\mathrm{AA}(\eta, P_0)) \le \mathrm{Loss}_T(\theta) + \frac{1}{\eta} \ln \frac{1}{P_0\{\theta\}},$$

with the loss function either $-\ln(1+\gamma\cdot\omega)$ *or* $\ln \frac{1+a\|\omega\|_\infty}{1+\gamma\cdot\omega}$.

5 PREDICTIVE COMPLEXITY

So far we have discussed merging virtually arbitrary pools of strategies for Learner; there remains, however, the question of how to choose the pool of strategies to be merged. One (traditional) possibility is to merge a pool which we feel contains a strategy that will perform well. For example, if we believe that in some securities market security prices are generated by an iid process (a dubious assumption), we could choose to merge all constant rebalanced portfolios (see Theorem 2 above). The wider such a pool is, the more justified is the belief that one of the strategies in it will perform well; unfortunately, it is also true that it will be more difficult to compete with the best strategy in the pool. We are more interested, however, in the case where Learner does not have any beliefs and is not willing to make any assumptions about her environment (such as a securities market). At first it might seem that without any limitations on Learner's environment it is impossible to devise any reasonable strategy for Learner. However, since we deal with the system

$$\text{Learner} + \text{Environment},$$

instead of imposing restrictions on the Environment part we can impose restrictions on the Learner part. Indeed, such Learner-side limitations are very natural: we know that she must *compute* her strategy and, moreover, her computational resources are bounded. Instead of pools reflecting our beliefs we can use pools reflecting Learner's limitations such as the pool of all computable strategies or, even better, the pool of all efficiently computable strategies. (Cf. [13].)

In this paper we consider the simplest situation where Learner's strategy is required to be computable but in a somewhat unusual sense: she is allowed to spend infinite amount of time computing her actions. We will prove that there exists the best, up to an additive constant, strategy of this kind. Of course, this strategy cannot be directly used in practice, but it can be a useful (though not achievable) goal: as in the theory of Kolmogorov complexity, we can study different efficiently computable approximations to this best (or *universal*) strategy.

We will mainly concentrate on what we call "predictive complexity": the predictive complexity of a data sequence is the loss suffered by the universal strategy on that sequence. (The word "predictive" is not very suitable for this paper, since Learner's actions in Cover's or the long-loss game can hardly be interpreted as predictions, but it is widely used in related contexts, and we use it as well.) Now we will give formal definitions (partly following [16], where the square-loss game was considered).

First we formally state a protocol which includes all games we are interested in. It will involve one more set, Σ, which we call the *signal space*; as before, we have the outcome space Ω, the action space Γ, and the loss function $\lambda : \Omega \times \Gamma \to \mathbb{R}$. We will rarely need the signals, and in most of this section the reader can simply ignore mentioning signals (or, alternatively, assume that Σ is a 1-element set and so the signals do not carry any information); sometimes we will even drop mentioning signals. Our protocol is:

FOR $t = 1, 2, \ldots$ Learner
  observes signal $\sigma_t \in \Sigma$
  chooses prediction $\gamma_t \in \Gamma$
  observes the actual outcome $\omega_t \in \Omega$
  suffers loss $\lambda(\omega_t, \gamma_t)$
END FOR.

Such a quadruple $(\Sigma, \Omega, \Gamma, \lambda)$ is called our *game*. We will assume that the sets Σ, Ω and Γ are equipped with some computability structure that allows one to speak of, say, computable functions on $\Omega \times \Gamma$ (in our examples this somewhat vague assumption is obviously satisfied). The function λ is assumed to be computable.

A *data sequence* is defined to be a finite sequence $x \in (\Sigma \times \Omega)^*$ of signal/outcome pairs; instead of

$x = ((\sigma_1, \omega_1), \dots, (\sigma_T, \omega_T))$ we will usually write $x = (\sigma_1 \dots \sigma_T | \omega_1 \dots \omega_T)$. Let S be a *prediction strategy*, i.e., a function that maps every data sequence x into an action $S(x) \in \Gamma$. Our notation for the total loss

$$\sum_{t=1}^{T} \lambda\left(\omega_t, S(\omega_1 \dots \omega_{t-1} | \sigma_1 \dots \sigma_{t-1})\right)$$

incurred over the first T trials by Learner who follows S will be $\mathrm{Loss}_S(x)$, where $x = (\omega_1 \dots \omega_T | \sigma_1 \dots \sigma_T)$ are the realized signals and outcomes. The function $\mathrm{Loss}_S(x)$ of a finite sequence $x \in (\Sigma \times \Omega)^*$ is called the *loss process* of S; a real-valued function on $(\Sigma \times \Omega)^*$ is a *loss process* if it coincides with Loss_S for some prediction strategy S.

Especially important are the loss processes corresponding to computable prediction strategies S; in all our examples these are exactly the computable loss processes. It would be ideal if the class of computable loss processes contained a smallest (say, to within an additive constant) element. Unfortunately, for the loss functions in our games such a smallest element does not exist: given a computable prediction strategy S, it is easy to construct a computable prediction strategy that greatly outperforms S on at least one signal/outcome sequence. Levin suggested (for a particular game, the log-loss game; see below) a very natural solution to the problem of non-existence of a smallest computable loss process.

We will say that a function $k : (\Sigma \times \Omega)^* \to \mathbb{R}$ is a *measure of predictive complexity* if the following two conditions hold:

1. k must be a *superloss process*, which means that $k(\square) = 0$ (where $\square$ is the empty sequence) and
$$\forall x \in (\Sigma \times \Omega)^* \, \exists \gamma \in \Gamma \, \forall \sigma \in \Sigma, \omega \in \Omega : \quad k(x * (\sigma, \omega)) \geq k(x) + \lambda(\omega, \gamma), \tag{12}$$
where $(\omega_1 \dots \omega_T | \sigma_1 \dots \sigma_T) * (\sigma, \omega)$ is defined to be $(\omega_1 \dots \omega_T \omega | \sigma_1 \dots \sigma_T \sigma)$.

2. k must be *semicomputable from above*, which means that there exists a computable sequence of computable functions $k_i : (\Sigma \times \Omega)^* \to \mathbb{R}$ such that, for every $x \in (\Sigma \times \Omega)^*$, $k(x) = \inf_i k_i(x)$.

Requirement 1 means that our measure of predictive complexity must be valid: there must exist a prediction strategy that achieves it. (Notice that if $\geq$ is replaced by $=$ in (12), we will obtain the definition of a loss process.) Requirement 2 means that it must be "computable in the limit"; since we are interested in a universal measure of predictive complexity k, we cannot hope that we will be able to compute it in finite time; all we can do is to output more and more accurate approximations k_i to it so that in the limit we obtain k. (Notice that in item 2 we can assume, without loss of generality, that the sequence k_i is decreasing.)

A smallest, to within an additive constant, measure k^* of predictive complexity will be said to be *universal*. In other words, a measure k^* of predictive complexity is universal if for any other measure k of predictive complexity there exists a constant C such that

$$\forall x \in (\Sigma \times \Omega)^* : k^*(x) \leq k(x) + C. \tag{13}$$

In Subsection 7.6 below we will prove that a universal measure of predictive complexity exists for perfectly mixable games (we say that a game is *perfectly mixable* if it is η-mixable for some $\eta > 0$; as we already mentioned, the games considered in this paper are perfectly mixable).

Lemma 6 *There exist universal measures of predictive complexity for perfectly mixable games.*

(Remember that we always assume that our games satisfy assumptions of computability which we do not specify explicitly.)

For every perfectly mixable game we fix a universal measure of predictive complexity; the notation for the latter will be $\mathcal{K}$ with sub- and/or superscripts to identify the game; these sub-/superscripts will be dropped when clear from the context. For every $x \in (\Sigma \times \Omega)^*$, $\mathcal{K}(x)$ will be called the *predictive complexity* of x.

Remark 1 We have defined predictive complexity only "up to an additive constant" (see (13)). It is not difficult to make this constant more explicit (see Corollary 1 below), but to get rid of it completely we need to specify a concrete universal measure of predictive complexity. Some suggestions on how to choose concrete variants of Kolmogorov complexity can be found in [9], Section 3.2.

An important perfectly mixable game is the N-*outcome log-loss* game $(\Sigma, \Omega, \Gamma, \lambda)$, which is defined as follows:

$$\Omega = \{0, \dots, N-1\}, \ \Gamma = \mathcal{S}_N, \quad \lambda(\omega, \gamma) = -\ln \gamma[\omega],$$

where $\mathcal{S}_N$ is the standard simplex in $\mathbb{R}^N$,

$$\mathcal{S}_N = \left\{ (\gamma[0], \dots, \gamma[N-1]) \in [0,1]^N \mid \gamma[0] + \dots + \gamma[N-1] = 1 \right\}$$

(the same action space as in Cover's game); the signal space can be chosen arbitrarily. Since $\gamma[0]$ is determined by $\gamma[1], \dots, \gamma[N-1]$, it is superfluous and the action space can be taken to be the set of all possible $(\gamma[1], \dots, \gamma[N-1])$; this is often done in the case $N = 2$ usually considered in literature, where the action space reduces to $[0,1]$. Levin [17] proved the existence of a universal measure of predictive complexity for the log-loss game (in which the notion of a universal measure of predictive complexity is tantamount to the minus logarithm of Levin's *a priori* semimeasure). We will use the notation $\mathcal{K}^{\log}$ for the predictive complexity in the log-loss game; in the theory of Kolmogorov complexity binary rather than natural logarithms are usually used, and so there only exists a standard notation for the function $\mathcal{K}^{\log}/\ln 2$, which is KM. It is well-known that the function $KM = \mathcal{K}^{\log}/\ln 2$ is connected with other variants of Kolmogorov complexity by the relations

$$|KM(x) - K(x)| = O(\log |x|)$$

and

$$|KM(x) - C(x)| = O(\log |x|),$$

where K is prefix complexity, C is plain Kolmogorov complexity and $|x|$ is the length of x. (For simplicity we are ignoring signals. Prefix complexity will be defined below.)

It is easy to see that the log-loss game is essentially a special case of Cover's game corresponding to what we call *Kelley* (or *horse-race*) markets, where only outcomes in

$$\hat{\Omega} := \{(\omega[0], \ldots, \omega[N-1]) \in \{0,1\}^N \mid \omega[0] + \cdots + \omega[N-1] = 1\}$$

are observed: we identify every outcome $\omega \in \hat{\Omega}$ with the number n (which is an outcome in the log-loss game) for which $\omega[n] = 1$.

We will use the following notation for predictive complexities in the games we are most interested in: $\mathcal{K}_N^{\mathrm{C}}(x)$ for Cover's game with N stocks and the loss function $-\ln(\gamma \cdot \omega)$ (*Cover complexity*), and $\mathcal{K}_{N,a}^{\text{l-s}}(x)$ for the long-short game with N stocks, "prudence coefficient" a and the loss function $-\ln(1+\gamma\cdot\omega)$ (*long-short complexity*). Notice that we do not need special notation for Cover's game with the loss function $-\ln \frac{\|\omega\|_\infty}{\gamma\cdot\omega}$, since it can be expressed as $\mathcal{K}_N^{\mathrm{C}}(x^*)$, where the "normalization" x^* of a data sequence x is defined in a natural way: we replace every σ_t in $x = (\omega_1 \ldots \omega_T | \sigma_1 \ldots \sigma_T)$ by (σ_t, ω_{t-1}) and multiply every ω_t by the number c_t such that $\|c_t \omega_t\|_\infty = 1$.

Cover complexity contains as a special case not only *KM* but also prefix complexity K (see [9], Chapter 3) if we allow the number N of stocks to be infinite. Therefore, we define Cover's game $(\Omega, \Gamma, \lambda)$ with infinitely many stocks as follows:

$$\Omega = \{(\omega[0], \omega[1], \ldots) \in [0,\infty)^\infty \mid \sup_n \omega[n] < \infty\},$$
$$\Gamma = \{(\gamma[0], \gamma[1], \ldots) \in [0,\infty)^\infty \mid \textstyle\sum_{n=0}^\infty \gamma[n] = 1\},$$
$$\lambda(\omega,\gamma) = \textstyle\sum_{n=0}^\infty \gamma[n]\omega[n].$$

We define the "unstructured" complexity $\mathcal{K}^{\mathrm{int}}$ of an integer $n \geq 0$ to be $\mathcal{K}_\infty^{\mathrm{C}}(\nu)$, where ν is a 1-element sequence whose only element is n (similarly to what we had before, n is identified with the vector $(0, \ldots, 0, 1, 0, \ldots) \in \mathbb{R}^\infty$ where the index of the only 1 is n). Theorem 4.3.3 of [9] shows that $K \stackrel{+}{=} \mathcal{K}^{\mathrm{int}}/\ln 2$ (we let $\stackrel{+}{=}$, $\stackrel{+}{\geq}$ and $\stackrel{+}{\leq}$ stand for the equality and the corresponding inequalities to within an additive constant). As usual, we define the complexity of a computable object (such as a prediction strategy) as the complexity of a simplest program (encoded as an integer) for computing that object; the complexity of a non-computable object is always ∞.

The following simple corollary of Theorems 1 and 3 shows how the constant C in the definition of Cover and long-short complexity (see (13)) depends on the competing investment strategy.

Corollary 1 *Fix the number of stocks N and the prudence coefficient $a > 0$. There exists a constant c such that, for any finite sequence $x \in \Omega^*$ and any computable prediction strategy S in Cover's game or the long-short game,*

$$\mathcal{K}(x) \leq \mathrm{Loss}_S(x) + \mathcal{K}^{\mathrm{int}}(S) + c. \quad (14)$$

This corollary follows from the inequalities of Theorems 1 and 3 with $\eta = 1$. It is interesting that, though any learning rate $\eta \leq 1$ can be used for defining the predictive complexity (see Subsection 7.6 below), inequality (14) holds no matter which value of η is used.

Corollary 2 shows how the performance of the algorithms $\mathrm{AA}(P_0, \eta)$ (see Theorems 1 and 3) for different η bounds predictive complexity.

Corollary 2 *There exists a constant c such that, for any finite sequence $x \in \Omega^*$ and any learning rate η in Cover's or the long-short game,*

$$\mathcal{K}(x) \leq \mathrm{Loss}_{\mathrm{AA}(P_0,\eta)}(x) + \mathcal{K}^{\mathrm{int}}(\eta) + c.$$

The next theorem gives an upper estimate of Cover complexity in terms of Kolmogorov complexity. We fix a computable signal $\perp \in \Sigma$ (it will be used to represent "no signal"); recall that the symbols $\stackrel{+}{\leq}$ and $\stackrel{+}{\geq}$ mean that the corresponding inequalities hold to within an additive constant.

Theorem 4 *For any fixed N,*

$$\mathcal{K}_N^{\mathrm{C}}(\omega_1 \ldots \omega_T) \stackrel{+}{\leq} -\ln \textstyle\sum_{(n_1,\ldots,n_T)\in\{0,\ldots,N-1\}^T} \omega_1[n_1]\ldots\omega_T[n_T] 2^{-KM(n_1\ldots n_T|\perp\omega_1\ldots\omega_{T-1})}. \quad (15)$$

For any computable prediction strategy S in Cover's game,

$$\mathrm{Loss}_S(\omega_1 \ldots \omega_T) \stackrel{+}{\geq} -\ln \textstyle\sum_{(n_1,\ldots,n_T)\in\{0,\ldots,N-1\}^T} \omega_1[n_1]\ldots\omega_T[n_T] 2^{-KM(n_1\ldots n_T|\perp\omega_1\ldots\omega_{T-1})}. \quad (16)$$

It is an open problem to find out if the inequality opposite to (15) is true; (16) shows that it is "close to being true". In the case of horse races the inequality opposite to (15) is true, since for horse races Cover complexity coincides with $KM \ln 2$.

The theory of Kolmogorov complexity contains three major parts:

- theory of Kolmogorov complexity proper;
- theory of randomness;
- theory of information

(see, e.g., the early survey by Zvonkin and Levin [17]). In the rest of this section we will briefly discuss analogs of randomness and information in Cover's and the long-short game. (In what follows we will not mention the game explicitly; what we say will be applicable to both games.)

We will say that a computable strategy S is *efficient* for a data sequence $\omega_1\omega_2\ldots \in \Omega^\infty$ if

$$\exists C \forall T : \mathrm{Loss}_S(\omega_1 \ldots \omega_T) \leq \mathcal{K}(\omega_1 \ldots \omega_T) + C$$

(in other words, if the loss suffered by S is within an additive constant of the predictive complexity on $\omega_1\omega_2\ldots$). For example, one possible interpretation of the market efficiency can be that the market portfolio is efficient for the actual sequence of security prices.

This market efficiency assumption has many interesting consequences: for example, it is possible to state variants (which are stronger than the corresponding limit theorems for sequences of events) of the strong law of large numbers and the upper half of the law of the iterated logarithm as assertions about efficient securities markets.

In the case of the log-loss game our notion of predictive efficiency becomes equivalent to Martin-Löf randomness: a prediction strategy S in the log-loss game is efficient for a sequence in $\{0, \dots, N-1\}^\infty$ if and only if this sequence is Martin-Löf random with respect to S.

Next we discuss the analog of the notion of information in the case of our two games. Suppose, for example, that, along with a dataset describing the security prices at some stock exchange, we also have some inside information about one of the companies whose stock is traded at this stock exchange; we represent this inside information as $\sigma_1 \dots \sigma_T$ (σ_t can be, say, the decisions taken by the company's CEO on day $t-1$). How can we define an objective measure of the usefulness of this information? A natural answer is

$$\begin{aligned} I(\sigma_1 \dots \sigma_T : \omega_1 \dots \omega_T) := \mathcal{K}(\omega_1 \dots \omega_T) \\ -\mathcal{K}(\omega_1 \dots \omega_T | \sigma_1 \dots \sigma_T), \end{aligned} \tag{17}$$

the best achievable loss without the side information $\sigma_1 \dots \sigma_T$ minus the best achievable loss with the side information. This is a general definition; in the case of Cover's and the long-short games the predictive information $I(\sigma_1 \dots \sigma_T : \omega_1 \dots \omega_T)$ is the logarithm of the increase in the final value of the universal portfolio (worth, say, £1 initially) due to the use of the side information.

In the case of the log-loss game, the notion of predictive information (17) is similar to (but different from) the standard notion of the information about $\omega_1 \dots \omega_T$ contained in $\sigma_1 \dots \sigma_T$ (see [9], Section 2.8): in defining the latter, the whole sequence $\sigma_1 \dots \sigma_T$ is assumed to be given in advance (rather than to be supplied elementwise).

6 DISCUSSION

In the previous section we generalized Kolmogorov complexity to the case of Cover's game and extended it to the case of the long-short game. It is easy to see that we can generalize even further taking into account transaction costs and bid-ask spreads. In the case of transaction costs, we only need to make assumption 3 in Subsection 2.1 of Blum and Kalai [1]: an investment strategy I which invests an initial fraction α of its money according to investment strategy I_1 and an initial fraction $1-\alpha$ according to I_2, should achieve at least α times the wealth of I_1 plus $1-\alpha$ times the wealth of I_2. For example, this property is satisfied if we have a fixed percentage commission $c \in (0,1)$ charged for buying each stock, and we can define the modifications $\mathcal{K}^{\mathrm{C}}_{N,c}$ and $\mathcal{K}^{\text{l-s}}_{N,a,c}$ of the complexities $\mathcal{K}^{\mathrm{C}}_N$ and $\mathcal{K}^{\text{l-s}}_{N,a}$ introduced above.

In the case where bid-ask spreads are allowed, we need to consider a more complicated outcome space,

$$\Omega = \{(\omega^{\mathrm{B}}[0], \omega^{\mathrm{A}}[0], \dots, \omega^{\mathrm{B}}[N-1], \omega^{\mathrm{A}}[N-1]) \in [0,\infty)^N \mid \omega^{\mathrm{B}}[0] \le \omega^{\mathrm{A}}[0], \dots, \omega^{\mathrm{B}}[N-1] \le \omega^{\mathrm{A}}[N-1]\},$$

where $\omega^{\mathrm{A}}[n]$ are the ask prices and $\omega^{\mathrm{B}}[n]$ are the bid prices, and to change the loss function λ correspondingly. It is easy to see that predictive complexity can be defined for the modifications of Cover's and the long-short games with bid-ask spreads.

It is interesting that when transaction costs and bid-ask spreads are taken into account, our two main games (Cover's and long-short) do not fit into the standard protocol (see, e.g., [14]) of the theory of prediction with expert advice any longer: the current action influences not only the immediate losses, but there also appears the possibility of delayed losses or rewards, since the choice of the portfolio at trial t will make it easier or more difficult to switch to desirable portfolios at later trials $t+1, t+2, \dots$.

There also remain several interesting questions related to the pool of constant rebalanced portfolios. First, it would be nice to prove an analog of Theorem 2 above for the long-short game. A potentially promising direction of further research would be to study the possibility of tracking the best constant rebalanced portfolio (see Herbster and Warmuth [7] and also Vovk [15]).

7 SOME PROOFS

7.1 EXISTENCE OF SUBSTITUTION FUNCTIONS SATISFYING (2)

In this subsection we will prove that, under our assumptions on $(\Omega, \Gamma, \lambda)$, the inf in the definition of $c(g)$ exists.

The case $c(g) = \infty$ is trivial, so we assume that $c(g)$ is finite. Let $c_1 c_2 \dots$ be a decreasing sequence such that $c_k \to c(g)$ as $k \to \infty$. By the definition of $c(g)$, for each k there exists $\delta_k \in \Gamma$ such that

$$\forall \omega\colon \frac{\lambda(\omega, \delta_k)}{g(\omega)} \le c_k.$$

Let δ be a limit point (whose existence follows from the compactness of Γ) of the sequence $\delta_1 \delta_2 \dots$. Then, for each ω, $\lambda(\omega, \delta)$ is a limit point of the sequence $\lambda(\omega, \delta_k)$ (by the continuity of λ) and, therefore,

$$\frac{\lambda(\omega, \delta)}{g(\omega)} \le c(g).$$

(This means that we can put $\Sigma(g) := \delta$.)

7.2 PROOF OF LEMMA 3

We assume that the loss function is $-\ln(\gamma \cdot \omega)$ and do not assume $\|\omega\|_\infty = 1$.

Let $g \in \mathrm{GA}(\eta)$ ($\eta \le 1$) be generated by a probability distribution P:

$$g(\omega) = \log_\beta \int_\Gamma \beta^{\lambda(\omega,\gamma)} P(d\gamma);$$

we are required to prove that $c(g) = 1$ and that (7) is the only action for which

$$\forall \omega : \lambda(\omega, \gamma^*) \le g(\omega).$$

First we prove that $c(g) \geq 1$. Let an action $\gamma = (\gamma[0], \ldots, \gamma[N-1])$ satisfy

$$\sup_\omega \lambda(\omega, \gamma) \leq c g(\omega)$$

for all ω. Taking

$$\omega[n] := \begin{cases} 1-t, & \text{if } n = j, \\ 1, & \text{otherwise,} \end{cases}$$

where $j \in \{0, \ldots, N-1\}$ and $t \in (0,1)$, we obtain

$$\begin{gathered} -\ln(1-\gamma[j] + (1-t)\gamma[j] \\ \leq c \log_\beta \int_\Gamma \beta^{-\ln(1-\delta[j]+(1-t)\delta[j])} P(d\delta); \\ \ln(1-\gamma[j] + (1-t)\gamma[j]) \\ \geq c\tfrac{1}{\eta} \ln \int_\Gamma (1-\delta[j] + (1-t)\delta[j])^\eta P(d\delta). \end{gathered}$$

For $t = 0$, the last inequality turns into the equality $0 = 0$; therefore, we can replace the left-hand and right-hand sides of this inequality by their derivatives in t at the point $t = 0$. We find:

$$\begin{gathered} \frac{-\gamma[j]}{1-\gamma[j] + (1-t)\gamma[j]} \\ \geq c \frac{1}{\eta} \frac{\int_\Gamma \eta (1-\delta[j] + (1-t)\delta[j])^{\eta-1} (-\delta[j]) P(d\delta)}{\int_\Gamma (1-\delta[j] + (1-t)\delta[j])^\eta P(d\delta)}, \end{gathered}$$

i.e.,

$$\gamma[j] \leq c \int_\Gamma \delta[j] P(d\delta). \tag{18}$$

Summing over j, we can see that $c \geq 1$ and that $c = 1$ is possible only for the average portfolio (7).

It remains to prove that (7) satisfies

$$\lambda(\omega, \gamma^*) \leq \log_\beta \int_\Gamma \beta^{\lambda(\omega,\gamma)} P(d\gamma), \tag{19}$$

for any $\omega \in \Omega$. Let us fix ω and put $f(\gamma) := \beta^{\lambda(\omega,\gamma)}$. Inequality (19) can be rewritten as

$$f\left(\int_\Gamma \gamma P(d\gamma)\right) \geq \int_\Gamma f(\gamma) P(d\gamma);$$

the last inequality follows from the concavity of $f(\gamma) = (\gamma \cdot \omega)^\eta$ (recall that $\eta \leq 1$).

7.3 PROOF OF LEMMA 4

In this subsection we will assume that the loss function is $-\ln(\gamma \cdot \omega)$ under the constraint $\|\omega\|_\infty = 1$ (as we already mentioned, this is equivalent to considering loss function (6)). So we consider the game $(\Omega, \Gamma, \lambda)$, where

$$\Omega = \left\{\omega \in [0,\infty)^N \mid \|\omega\|_\infty = 1\right\}, \quad \Gamma = \mathcal{S}_N,$$

$$\lambda(\omega, \gamma) := -\ln(\gamma \cdot \omega)$$

(recall that $\mathcal{S}_N$ is the standard simplex in $\mathbb{R}^N$).

Fix $\eta > 1$. We start by proving that $c(\eta) \geq \eta$. *Degenerate* outcomes $\omega \in \Omega$ and portfolios $\gamma \in \Gamma$ are defined as those of the form $(0, \ldots, 0, 1, 0, \ldots, 0)$; the *nth degenerate* outcome or portfolio contains n zeroes before the only 1. Consider the probability distribution in Γ assigning equal weights $\frac{1}{N}$ to the N degenerate portfolios. The loss of the corresponding generalized action will be

$$\begin{aligned} &\log_\beta \frac{1}{N} \left(\beta^{-\ln 1} + \beta^{-\ln 0} + \cdots + \beta^{-\ln 0}\right) \\ &= \log_\beta \frac{1}{N} = \frac{\ln N}{\eta} \end{aligned}$$

at every degenerate outcome. On the other hand, the loss of any real portfolio $\gamma = (\gamma[0], \ldots, \gamma[N-1])$ at the nth degenerate outcome is $-\ln \gamma[n]$; therefore, it is at least $-\ln \frac{1}{N}$ for at least one of the degenerate outcomes. We can see that

$$c(\eta) \geq \frac{-\ln \frac{1}{N}}{\frac{\ln N}{\eta}} = \eta.$$

It remains to prove that for any probability distribution P in Γ there exists $\gamma \in \Gamma$ such that, for any $\omega \in \Omega$,

$$-\ln(\gamma \cdot \omega) \leq \eta \log_\beta \int_\Gamma \beta^{-\ln(\delta\cdot\omega)} P(d\delta).$$

This inequality is equivalent to

$$\ln(\gamma \cdot \omega) \geq \ln \int_\Gamma e^{\eta \ln(\delta\cdot\omega)} P(d\delta),$$

i.e.,

$$\gamma \cdot \omega \geq \int_\Gamma (\delta \cdot \omega)^\eta P(d\delta).$$

The last inequality is obviously true for $\gamma = \int_\Gamma \delta P(d\delta)$.

7.4 PROOF OF THEOREM 2

By Lemmas 1 and 3, we are required to prove, for all $\delta \in \Gamma$,

$$\log_\beta \int_\Gamma \beta^{\mathrm{Loss}_T(\gamma)} P_0(d\gamma) \leq \mathrm{Loss}_T(\delta) + \frac{N-1}{2\eta} \ln T + c,$$

i.e., to prove

$$\int_\Gamma \prod_{t=1}^T (\gamma \cdot \omega_t)^\eta P_0(d\gamma) \geq \prod_{t=1}^T (\delta \cdot \omega_t)^\eta \times T^{-\frac{N-1}{2}}.$$

As in Cover and Ordentlich [2], we will reduce our general problem to Kelley markets. To do so, it is sufficient to notice that the function

$$\left(\int_\Gamma (\gamma \cdot \omega)^\eta P(d\gamma)\right)^{1/\eta}$$

is concave in ω (when $\eta \leq 1$): this is a special case of Minkowski's inequality $\|f_1 + f_2\|_p \geq \|f_1\|_p + \|f_2\|_p$, where f_1 and f_2 are nonnegative functions and $p \in (0,1]$ (cf. [4], Theorem 24).

So we shall consider only Kelley markets. We need to estimate

$$\frac{\left(\int_\Gamma \prod_{t=1}^T (\gamma \cdot \omega_t)^\eta P_0(d\gamma)\right)^{1/\eta}}{\prod_{t=1}^T (\delta \cdot \omega_t)}$$

from below by $\epsilon T^{-\frac{N-1}{2\eta}}$ for some $\epsilon = \epsilon(N, \eta) > 0$.

Recall that our prior is the Dirichlet distribution with parameters $\frac{1}{2}$. We let $\mathcal{E}_{N-1}$ stand for the standard solid simplex in $\mathbb{R}^{N-1}$,

$$\mathcal{E}_{N-1} = \{(\gamma[1], \dots, \gamma[N-1]) \in [0,1]^{N-1} \mid \gamma[1] + \dots + \gamma[N-1] \le 1\},$$

and let $\gamma[0]$ stand for $1 - \gamma[1] - \dots - \gamma[N-1]$; $\stackrel{\times}{=}$ means equality to within a constant factor (which may depend on η and N). If T_n is the number of times the nth degenerate outcome occurs (and so $T_0 + \dots + T_{N-1} = T$), we have:

$$\begin{aligned}
&\frac{\left(\int_\Gamma \prod_{t=1}^T (\gamma\cdot\omega_t)^\eta P_0(d\gamma)\right)^{1/\eta}}{\prod_{t=1}^T (\delta\cdot\omega_t)} \\
&\ge \Big(\int_{\mathcal{E}_{N-1}} (\gamma[0])^{T_0\eta} \dots (\gamma[N-1])^{T_{N-1}\eta} \\
&\times (\gamma[0])^{-1/2} \dots (\gamma[N-1])^{-1/2} d\gamma[1] \dots d\gamma[N-1]\Big)^{1/\eta} \\
&\times B^{-1/\eta}(1/2, \dots, 1/2) \left(T_0/T\right)^{-T_0} \left(T_{N-1}/T\right)^{-T_{N-1}} \\
&\stackrel{\times}{=} B^{1/\eta}\left(T_0\eta + \tfrac{1}{2}, \dots, T_{N-1}\eta + \tfrac{1}{2}\right) \frac{T^T}{T_0^{T_0} \dots T_{N-1}^{T_{N-1}}} \\
&\stackrel{\times}{=} \left(\frac{\Gamma(T_0\eta + 1/2) \dots \Gamma(T_{N-1}\eta + 1/2)}{\Gamma(T\eta + N/2)}\right)^{1/\eta} \\
&\times \frac{T^T}{T_0^{T_0} \dots T_{N-1}^{T_{N-1}}};
\end{aligned}$$

using Stirling's formula $\Gamma(z) \stackrel{\times}{=} z^{z-1/2} e^{-z}$, we continue:

$$\begin{aligned}
&\stackrel{\times}{=} \left(\frac{(T_0\eta + 1/2)^{T_0\eta} \dots (T_{N-1}\eta + 1/2)^{T_{N-1}\eta}}{(T\eta + N/2)^{T\eta + N/2 - 1/2}} e^0\right)^{1/\eta} \\
&\times T^T T_0^{-T_0} \dots T_{N-1}^{-T_{N-1}} \\
&= \left(\frac{\left(1 + \frac{1}{2T_0\eta}\right)^{T_0\eta} \dots \left(1 + \frac{1}{2T_{N-1}\eta}\right)^{T_{N-1}\eta}}{(T\eta + N/2)^{N/2 - 1/2} \left(1 + \frac{N}{2T\eta}\right)^{T\eta}}\right)^{1/\eta} \\
&\stackrel{\times}{=} T^{-\frac{N-1}{2\eta}}.
\end{aligned}$$

7.5 PROOF OF LEMMA 5

In this subsection we only consider the loss function $\ln \frac{1+a\|\omega\|_\infty}{1+\gamma\cdot\omega}$ (the proof for the loss function $-\ln(1+\gamma\cdot\omega)$ is obtained by formally setting $a := 0$).

Let $g \in \mathrm{GA}(\eta)$ ($\eta \le 1$) be generated by a probability distribution P:

$$\begin{aligned}
g(\omega) &= \log_\beta \int_\Gamma \beta^{\lambda(\omega,\gamma)} P(d\gamma) \\
&= \log_\beta \int_\Gamma e^{-\eta \ln \frac{1+a\|\omega\|_\infty}{1+\gamma\cdot\omega}} P(d\gamma) \\
&= \log_\beta \int_\Gamma \left(\frac{1+\gamma\cdot\omega}{1+a\|\omega\|_\infty}\right)^\eta P(d\gamma) \\
&= \log_\beta \frac{\int_\Gamma (1+\gamma\cdot\omega)^\eta P(d\gamma)}{(1+a\|\omega\|_\infty)^\eta}; \qquad (20)
\end{aligned}$$

we are required to prove that $c(g) = 1$ and that (7) is the only action for which

$$\forall \omega : \lambda(\omega, \gamma^*) \le g(\omega).$$

(Notice that this is obvious for $\eta = 1$.)

First we prove that $c(g) \ge 1$. Setting

$$\omega[n] := \begin{cases} t, & \text{if } n = j, \\ 0, & \text{otherwise,} \end{cases}$$

where $j \in \{0, \dots, N-1\}$, letting $t \to \pm 0$ and replacing the numerator and denominator of (20) by their first-order Taylor approximations at $t = 0$, we obtain

$$g(\omega) \sim \log_\beta \frac{1 + \eta \left(\int_\Gamma \gamma[j] P(d\gamma)\right) t}{1 \pm \eta a t},$$

where $\sim$ stands for asymptotic equality as $t \to 0$. Since $\gamma^* = \int \gamma P(d\gamma)$, we further obtain

$$\begin{aligned}
g(\omega) &\sim \log_\beta \tfrac{1+\eta\gamma^*[j]t}{1+\eta a|t|} \\
&\sim \tfrac{\eta\gamma^*[j]t - \eta a|t|}{\ln\beta} = -\gamma^*[j]t + a|t|.
\end{aligned}$$

Comparing this expression with the loss of any other action $\gamma \ne \gamma^*$,

$$\ln \frac{1+a|t|}{1+\gamma[j]t} \sim -\gamma[j]t + a|t|,$$

we can see that always $c(g) \ge 1$ and $c(g) = 1$ can be attained by no $\gamma \ne \gamma^*$.

It remains to prove that γ^* attains $c(g) = 1$ if and only if $\eta \le 1$, i.e., to prove (cf. (20)) that

$$\ln \frac{1+a\|\omega\|_\infty}{1+\gamma^*\cdot\omega} \le \log_\beta \frac{\int_\Gamma (1+\gamma\cdot\omega)^\eta P(d\gamma)}{(1+a\|\omega\|_\infty)^\eta} \qquad (21)$$

is equivalent to $\eta \le 1$. Inequality (21) is equivalent to:

$$\begin{aligned}
\eta \ln \frac{1+\gamma^*\cdot\omega}{1+a\|\omega\|_\infty} &\ge \ln \frac{\int_\Gamma (1+\gamma\cdot\omega)^\eta P(d\gamma)}{(1+a\|\omega\|_\infty)^\eta}; \\
\eta \ln(1+\gamma^*\cdot\omega) &\ge \ln \int_\Gamma (1+\gamma\cdot\omega)^\eta P(d\gamma); \\
(1+\gamma^*\cdot\omega)^\eta &\ge \int_\Gamma (1+\gamma\cdot\omega)^\eta P(d\gamma). \qquad (22)
\end{aligned}$$

If $\eta \le 1$, (22) immediately follows from the fact that the function $(1+\gamma\cdot\omega)^\eta$ is concave in γ: it is the composition of the linear function $\gamma \mapsto (1+\gamma\cdot\omega)$ and the concave function $t \mapsto t^\eta$. If $\eta > 1$, (22) is violated for non-trivial ω and P because the function $t \mapsto t^\eta$ is strictly convex for $\eta > 1$.

7.6 PROOF OF LEMMA 6

For simplicity we will assume that the sets Σ and Ω are finite. To satisfy this assumption we might replace in $\Omega = [0,\infty)^N$ the ray $[0,\infty)$ by a finite (but dense) subset; it will also be clear from the proof that it works for $\Omega = [0,\infty)^N$ as well if some of the definitions are (routinely but tediously) modified.

The idea behind our proof is that we merge all "semi-computable prediction strategies" using the AA. To simplify presentation, however, the explicit mentioning of the AA has been "compiled out" of our construction.

Our proof consists of two steps:

1. We observe that there exists a computable sequence $k_1, k_2, \dots$ of measures of predictive complexity.

2. We notice that

$$k(x) := \log_\beta \sum_{i=1}^\infty \beta^{k_i(x)} 2^{-i}$$

is a universal measure of predictive complexity.

Proof of 1

Choose a universal partial recursive function $M(i,j)$, where i and j range over the positive integers, taking positive integer values; the universality of M means that the set of functions M_i defined by $M_i(j) := M(i,j)$ contains all partial recursive functions of the type $\{1,2,\ldots\} \to \{1,2,\ldots\}$.

We will say that a measure N of predictive complexity is *finitary* if:

- it takes values in the set of rational numbers extended by adding ∞ (which is defined to be the largest element of the extended set of rational numbers);
- the set
$$\mathrm{carrier}(N) := \{x \mid N(x) < \infty\}$$
is finite.

Notice that there are only countably many finitary measures of predictive complexity; let N_j ($j = 1,2,\ldots$) be a computable enumeration of the finitary measures of predictive complexity.

Define $M^*(i,j)$ to be $M(i,j)$ if

- $M(i,j')$ is defined for all $j' \le j$, and
- $N_{M(i,j'+1)} \le N_{M(i,j')}$ for all $j' < j$;

$M^*(i,j)$ is undefined if either of these conditions fails. Notice that M^* is a partial recursive function. It remains to set
$$k_i := \inf_j N_{M^*(i,j)}$$
(the inf is over the j for which $M^*(i,j)$ is defined).

We omit the proof that the sequence $k_1, k_2, \ldots$ contains all measures of predictive complexity.

Proof of 2

Recall that we assume that our game is η-mixable for some η; fix such η and set $\beta := e^{-\eta}$. We say that a generalized action g is a *superaction* if there exists an action $\gamma \in \Gamma$ such that,
$$\forall \omega \in \Omega : g(\omega) \ge \lambda(\omega,\gamma)$$
(i.e., g is a permitted action, perhaps plus some extra losses). Since our game is η-mixable, every η-mixture of actions is a superaction.

Let us prove now that k is a measure of predictive complexity. Its upper semicomputability immediately follows from the upper semicomputability of all k_i and the computability of their sequence $k_1, k_2, \ldots$, so we only need to prove that k is a superloss process, i.e., that for all $x \in (\Sigma \times \Omega)^*$ and $\sigma \in \Sigma$ the difference
$$g(\omega) := k(x * (\sigma,\omega)) - k(x)$$
is a superaction. To see why this is true, we transform this difference as follows:
$$k(x*(\sigma,\omega)) - k(x)$$
$$= \log_\beta \sum_{i=1}^\infty \beta^{k_i(x*(\sigma,\omega))} 2^{-i} - \log_\beta \sum_{i=1}^\infty \beta^{k_i(x)} 2^{-i}$$
$$= \log_\beta \sum_{i=1}^\infty \frac{\beta^{k_i(x)}2^{-i}}{\sum_{j=1}^\infty \beta^{k_j(x)}2^{-j}} \beta^{k_i(x*(\sigma,\omega))-k_i(x)}.$$
We can see that this difference is an η-mixture of superactions and, therefore, is a superaction itself.

7.7 PROOF OF THEOREM 4

We can define, analogously to the *a priori* semimeasure, the *Cover exponential complexity* as follows:
$$\mathcal{M}(x) := e^{-\mathcal{K}_N^{\mathrm{C}}(x)};$$
notice that $\mathcal{M}$ can be defined to be the largest, up to a constant factor, process satisfying
$$\forall x \exists \gamma \forall \omega : \mathcal{M}(x*\omega) \le (\gamma \cdot \omega)\mathcal{M}(x)$$
(such processes will be called Cover exponential superloss processes in this proof). Notice that the Cover exponential complexity is an extension of the *a priori* semimeasure: if we consider only the degenerate outcomes (horse races), we obtain the definition of the *a priori* semimeasure.

We rewrite inequality (15) as follows:
$$\begin{aligned}\mathcal{M}(\omega_1\ldots\omega_T) &\overset{\times}{\ge} \textstyle\sum_{(n_1,\ldots,n_T)\in\{0,\ldots,N-1\}^T} \\ &\omega_1[n_1]\ldots\omega_T[n_T] M(n_1\ldots n_T \mid \perp\omega_1\ldots\omega_{T-1}),\end{aligned} \tag{23}$$
where M is the obvious "conditional" modification of the *a priori* semimeasure and $\overset{\times}{\ge}$ is "more than or equal to to within a constant factor".

To prove (23) it is sufficient to show that the function
$$\textstyle\sum_{(n_1,\ldots,n_T)\in\{0,\ldots,N-1\}^T} \omega_1[n_1]\ldots\omega_T[n_T] \times M(n_1\ldots n_T\mid\perp\omega_1\ldots\omega_{T-1})$$
is a Cover exponential superloss process, i.e., to show that
$$\begin{gathered}\forall(\omega_1\ldots\omega_T)\exists\gamma\forall\omega : \\ \textstyle\sum_{(n_1,\ldots,n_T)\in\{0,\ldots,N-1\}^T, n\in\{0,\ldots,N-1\}} \\ \omega_1[n_1]\ldots\omega_T[n_T]\omega[n] M(n_1\ldots n_T n\mid\perp\omega_1\ldots\omega_T) \\ \le (\gamma\cdot\omega) \textstyle\sum_{(n_1,\ldots,n_T)\in\{0,\ldots,N-1\}^T} \omega_1[n_1]\ldots\omega_T[n_T] \\ \times M(n_1\ldots n_T\mid\perp\omega_1\ldots\omega_{T-1}).\end{gathered}$$
Rewriting the last inequality as
$$\begin{gathered}\textstyle\sum_{(n_1,\ldots,n_T)\in\{0,\ldots,N-1\}^T} \omega_1[n_1]\ldots\omega_T[n_T] \\ \times\big(\textstyle\sum_{n\in\{0,\ldots,N-1\}} \omega[n] M(n_1\ldots n_T n\mid\perp\omega_1\ldots\omega_T) \\ -(\gamma\cdot\omega) M(n_1\ldots n_T\mid\perp\omega_1\ldots\omega_{T-1})\big) \le 0,\end{gathered}$$
we can see that it is sufficient to establish the existence of γ such that, for every ω,
$$\begin{gathered}\textstyle\sum_{n\in\{0,\ldots,N-1\}} \omega[n] M(n_1\ldots n_T n\mid\perp\omega_1\ldots\omega_T) \\ -(\gamma\cdot\omega) M(n_1\ldots n_T\mid\perp\omega_1\ldots\omega_{T-1}) \le 0.\end{gathered}$$
It is clear that we can take any
$$\gamma[n] \ge \frac{M(n_1\ldots n_T n\mid\perp\omega_1\ldots\omega_T)}{M(n_1\ldots n_T\mid\perp\omega_1\ldots\omega_{T-1})};$$
such γ exists by the definition of the *a priori* semimeasure. This completes the proof of inequality (15).

Inequality (16) is simple, and we omit its proof.

ACKNOWLEDGMENTS

We are grateful to Philip Dawid for enlightening discussions. The comments by the members of the program committee were very useful for improving the presentation. V. Vovk thanks EPSRC for providing financial support through grant GR/L35812.

References

[1] A. Blum and A. Kalai. Universal portfolios with and without transaction costs. In *Proceedings of the 10th Annual Conference on Computational Learning Theory*, pages 309–313, New York, 1997. Assoc. Comput. Mach.

[2] T. Cover and E. Ordentlich. Universal portfolios with side information. *IEEE Trans. Inform. Theory*, 42:348–363, 1996.

[3] M. Feder, N. Merhav, and M. Gutman. Universal prediction of individual sequences. *IEEE Trans. Inform. Theory*, 38:1258–1270, 1992.

[4] G. H. Hardy, J. E. Littlewood, and G. Pólya. *Inequalities*. Cambridge University Press, Cambridge, 1967.

[5] D. Haussler, J. Kivinen, and M. K. Warmuth. Tight worst-case loss bounds for predicting with expert advice. Technical Report UCSC-CRL-94-36, University of California at Santa Cruz, revised December 1994. Short version in P. Vitányi, editor, *Computational Learning Theory*, Lecture Notes in Computer Science, volume 904, pages 69–83, Springer, Berlin, 1995.

[6] D. P. Helmbold, R. E. Schapire, Y. Singer, and M. K. Warmuth. On-line portfolio selection using multiplicative updates. In *Proceedings of the 13th International Conference on Machine Learning*, 1996.

[7] M. Herbster and M. Warmuth. Tracking the best expert. In *Proceedings of the 12th International Conference on Machine Learning*, pages 286–294. Morgan Kaufmann, 1995.

[8] J. Kivinen and M. K. Warmuth. Exponential Gradient versus Gradient Descent for linear predictors. *Inform. Computation*, 132:1–63, 1997.

[9] M. Li and P. Vitányi. *An Introduction to Kolmogorov Complexity and Its Applications*. Springer, New York, 2nd edition, 1997.

[10] N. Littlestone, P. M. Long, and M. K. Warmuth. On-line learning of linear functions. *J. Comput. Complexity*, 5:1–23, 1995.

[11] N. Littlestone and M. K. Warmuth. The Weighted Majority Algorithm. *Inform. Computation*, 108:212–261, 1994.

[12] V. Vovk. Aggregating strategies. In M. Fulk and J. Case, editors, *Proceedings of the 3rd Annual Workshop on Computational Learning Theory*, pages 371–383, San Mateo, CA, 1990. Morgan Kaufmann.

[13] V. Vovk. Universal forecasting algorithms. *Inform. Computation*, 96:245–277, 1992.

[14] V. Vovk. A game of prediction with expert advice. In *Proceedings, 8th Annual Conference on Computational Learning Theory*, pages 51–60, New York, 1995. Assoc. Comput. Mach. To be published in *J. Comput. Syst. Sci.*

[15] V. Vovk. Derandomizing stochastic prediction strategies. In *Proceedings of the 10th Annual Conference on Computational Learning Theory*, pages 32–44, New York, 1997. Assoc. Comput. Mach.

[16] V. Vovk. Probability theory for the Brier game. Accepted for publication in *Theoretical Computer Science*, 1998. Preliminary version in M. Li and A. Maruoka, editors, *Algorithmic Learning Theory*, Lecture Notes in Computer Science, volume 1316, pages 323–338, 1997. Full version: Technical Report CSD-TR-97-09, Department of Computer Science, Royal Holloway, University of London, revised February 1998.

[17] A. K. Zvonkin and L. A. Levin. The complexity of finite objects and the development of the concepts of information and randomness by means of the theory of algorithms. *Russian Math. Surveys*, 25:83–124, 1970.