

The Aggregating Algorithm and Laissez-Faire Investment¹

Yuri Kalnishkan

Technical Report

CLRC — TR — 09 — 02

September 2009



Department of Computer Science
Egham, Surrey TW20 0EX, England

¹This article has been written for wiki project onlineprediction.net; see <http://onlineprediction.net/?n=Main.TheAAAndLaissez-FaireInvestment> for the up-to-date version. The author was supported by the EPSRC grant EP/F002998 'Practical competitive prediction'.

Abstract

This article describes a financial intuition behind the aggregating algorithm.

Contents

1	Sequential Investment	2
1.1	Sequential Investment: the Setup	2
1.2	Experts	2
1.3	Cover's Universal Portfolios	3
2	Bayesian Estimation	4
3	General Case	5
3.1	Games and Losses	5
3.2	Extension of Laissez-Faire Investment	6
3.3	Optimality	8

The aggregating algorithm was proposed by V. Vovk in [Vov90] (see also [VW98] and [Vov98]). This note aims to explain some intuition behind the aggregating algorithm and to obtain it as an extension of a natural investment strategy.

1 Sequential Investment

1.1 Sequential Investment: the Setup

Suppose that there are M shares (numbered $0, 1, \dots, M - 1$) we can invest into. The prices of shares change in discrete time. Let the vector x_t represents the proportional changes of the values of the shares between the moments $t - 1$ and t ; namely, let $x_{t,j}$ be the ratio of the prices of share i at the moments t and $t - 1$, $t = 1, 2, \dots$. We assume that all components of x_t are non-negative, $x_{t,j} \geq 0$.

We have a sum of money to invest in the shares. An investment decision can be represented by a vector showing how we split the capital among the shares. On step $t - 1$ we spend the fraction $\gamma_{t,j} \in [0, 1]$ of our capital to buy share j , $j = 1, 2, \dots, M$. Assume that we do not keep any money in cash (or that cash is actually one of the shares) and therefore $\sum_{j=0}^{M-1} \gamma_{t,i} = 1$. Possible investment decisions thus fill in the simplex.

Suppose that we have capital W_{t-1} on step $t - 1$. We buy $W_{t-1}\gamma_{t,j}$ worth¹ of share j . As the price changes between times $t - 1$ and t , this sum becomes $W_{t-1}\gamma_{t,j}x_{t,j}$. The total amount of money at our disposal thus changes from W_{t-1} to $W_t = W_{t-1}\langle\gamma_t, x_t\rangle$. We can sell the shares, consolidate the capital, and then make a new investment decision². If the initial capital is $W_0 = 1$, then after T steps we get $W_T = \prod_{t=1}^T \langle\gamma_t, x_t\rangle$.

It is typical to make assumptions about the behaviour of the prices, i.e., price changes x_t . The Black-Scholes theory postulates that a share price follows a geometric Brownian motion. The Capital Asset Price Model is less restrictive, but it still assumes that prices are random variables with means and variances. We will not be making any assumptions of this kind. Instead we are interested in properties of investment strategies that hold in the worst case, i.e., for all possible changes of stock prices.

1.2 Experts

Suppose that there are N experts $\mathcal{E}^{(1)}, \mathcal{E}^{(2)}, \dots, \mathcal{E}^{(N)}$ that suggest investment decisions to us. Before deciding on γ_t , we can observe decisions $\gamma_t^{(i)}$, $i = 1, 2, \dots, N$, output by experts. We make no assumptions as to how the experts arrive at their decisions. They may be algorithms of arbitrary nature and complexity, rely on side information etc. We treat them as black boxes producing decisions. Our goal is to merge their decisions in such a way so that our

¹We assume that we can buy an arbitrary fraction of a share; in real life one can only buy multiples of some fixed amount, but we ignore this.

²In real life buying and selling shares incurs transaction costs, and shares should not be bought or sold too often. We ignore this consideration in our analysis.

capital W_t is not much less than the capital of any expert $W_t^{(i)}$, $i = 1, 2, \dots, N$ (by $W_t^{(i)}$ we mean the capital we would have earned by starting from 1 and following the suggestions of expert i all the time); in other terms we want an inequality of the type $W_t \gtrsim W_t^{(i)}$ to hold uniformly for all i and possibly for all t .

We can think of a merging strategy in the following terms. At time $t - 1$ we split our capital into N parts and invest the i th part as suggested by expert i . Let $p_t^{(i)}$ represent the share of the wealth invested according to the suggestions of expert i .

We need to decide what share of the capital to entrust to expert i on step t . If an expert is doing well, it is natural to increase its share, and if an expert is doing badly then to decrease it. This can be achieved in the following natural way. Let us give each expert a share of the capital at the start and let it operate on what it has earned. This *laissez-faire* approach ensures that successful experts will have a higher influence.

If expert i gets the share $p^{(i)}$ of the initial capital $W_0 = 1$ and $W_t^{(i)}$ is as above, our capital at time t equals $W_t = \sum_{i=1}^N p^{(i)} W_t^{(i)}$. We thus get

$$W_t \geq p^{(i)} W_t^{(i)}$$

for all $i = 1, 2, \dots, N$. If we give each expert an equal share $1/N$, then we get

$$W_t \geq \frac{1}{N} W_t^{(i)} .$$

Note that the bounds hold for all times t and all possible changes in stock prices. We do not even need to know t in advance; the strategy is the same no matter at what point in the future we are going to check its performance.

It is easy to get a formula for $p_t^{(i)}$, the share of our wealth ‘in possession of’ expert i after time t . We have

$$\begin{aligned} p_t^{(i)} &= p^{(i)} W_t^{(i)} / W_t \\ &= p^{(i)} W_t^{(i)} / \sum_{i=1}^N p^{(i)} W_t^{(i)} . \end{aligned}$$

In fact giving a share of the wealth to an expert is a metaphor. We need to produce one investment decision, γ_t , and the formula for it is $\gamma_t = \sum_{i=1}^N p_{t-1}^{(i)} \gamma_t^{(i)}$.

One can update the weights in a simple way. To get the weights for step t we multiply each $p_{t-1}^{(i)}$ by $\langle \gamma_t^{(i)}, x_t \rangle$ and then normalise the vector to ensure the weights sum up to 1.

1.3 Cover’s Universal Portfolios

The bound we have obtained is rather trivial. However it is possible to obtain more involved results arguing in this way.

We can consider some class of investment strategies as experts and obtain a universal strategy for this class. Interesting classes of strategies are infinite.

One can easily expand the above argument to a countable class of experts. The extension to uncountable classes is less straightforward. The sum will be replaced by an integral and we will not be able to obtain a lower bound by simply dropping all terms except for one from the sum. One will need to estimate the integral instead.

Interesting results were obtained in this fashion in [CO96], where the class of *constant rebalanced portfolios* is considered.

A constant rebalanced portfolio is a strategy that sticks to a fixed partitioning of capital among M shares, which is maintained all the time. Consider a partitioning $\gamma = (\gamma_1, \gamma_2, \dots, \gamma_M)$, $\sum_{j=1}^M \gamma_j = 1$. The constant rebalanced portfolio using this γ requires that on each step the share γ_i of our total wealth is invested in share i . Suppose that we have invested the money accordingly at time $t - 1$. By the time t the share prices will change and this requirement will no longer hold, so we will need to buy and sell some shares to ensure that the share γ_i of our total wealth is invested in share i for all i .

We thus need to *rebalance* the portfolio on each step, hence the word ‘rebalanced’. However the rebalancing seeks to achieve a fixed allocation of money to shares, hence the word ‘constant’.

The constant rebalanced portfolios fill the simplex of dimension $M - 1$. In [CO96] two distributions on the simplex are considered. If the distribution is uniform, we get the bound

$$W_t \geq \frac{1}{(t+1)^{M-1}} W_\gamma$$

and for the Dirichlet distribution with parameters $(1/2, \dots, 1/2)$ we get

$$W_t \geq \frac{1}{2^{(t+1)^{(M-1)/2}} W_\gamma$$

for all t and all constant rebalanced portfolios γ .

Note that the coefficient depends on t and decreases as t increases. This is the price to pay for having a very large class of experts.

2 Bayesian Estimation

The investment scenario and our laissez-faire strategy turn out to have an important special case.

Let us consider a restricted version of the above scenario. Suppose that there are two “shares” and that the vector of proportional changes can take two values, $(0, 1)$ and $(1, 0)$. Within the investment metaphor this means that one part of our investment always disappears and another part is returned to us intact.

There is a different interpretation of this scenario though. Suppose we need to output the probability of some event. The outcome $x_t = (0, 1)$ corresponds to the event happening and $x_t = (1, 0)$ corresponds to it not happening. The “investment” $\gamma_{t,1}$ can be thought of as the probability of it happening and $\gamma_{t,0}$ as the probability of it not happening (recall that $\gamma_{t,0} + \gamma_{t,1} = 1$). Our “capital”

changes by the factor equal to the probability we assigned to the outcome that actually occurs. We shall use the notation

$$k_t = \begin{cases} 1, & \text{if } x_t = (0, 1); \\ 0, & \text{if } x_t = (1, 0). \end{cases}$$

This variable tells us whether the event has happened ($k_t = 1$) or not ($k_t = 0$).

We can think of experts $\mathcal{E}^{(1)}, \mathcal{E}^{(2)}, \dots, \mathcal{E}^{(N)}$ as of N hypotheses. Each hypothesis suggests a probability distribution: $\Pr(k_t = 1 \mid \mathcal{E}_i) = \gamma_{t,1}^{(i)}$ and $\Pr(k_t = 0 \mid \mathcal{E}_i) = \gamma_{t,0}^{(i)}$.

It is remarkable that the laissez-faire investment in this special case corresponds to the Bayesian approach. Let us think of the weights $p_{t-1}^{(1)}, p_{t-1}^{(2)}, \dots, p_{t-1}^{(N)}$ as of prior probabilities we have for the hypothesis before step t . The Bayes rule gives us posterior probabilities

$$\begin{aligned} p_t^{(i)} &= \Pr(\mathcal{E}^{(i)} \mid x_t) = \frac{\Pr(x_t \mid \mathcal{E}^{(i)}) \Pr(\mathcal{E}^{(i)})}{\Pr(x_t)} \\ &\sim \gamma_{t,k_t} p_{t-1}^{(i)} \\ &= \langle \gamma_t^{(i)}, x_t \rangle p_{t-1}^{(i)} \end{aligned}$$

where k_t is as above and \sim denotes proportionality. To get p_t one should normalise the products $\langle \gamma_t^{(i)}, x_t \rangle p_{t-1}^{(i)}$ to ensure the results sum up to 1. Note that this is exactly the same procedure as described above for investments.

The distribution γ_t is obtained from $\gamma_t^{(i)}$ by means of the product rule:

$$\begin{aligned} \Pr(k_t = k) &= \sum_{i=1}^N \Pr(k_t = k \mid \mathcal{E}^{(i)}) \Pr(\mathcal{E}^{(i)}) \\ &= \sum_{i=1}^N \gamma_{t,k} p_{t-1}^{(i)}, \end{aligned}$$

where $k = 0, 1$, i.e., $\gamma_t = \sum_{i=1}^N p_{t-1}^{(i)} \gamma_t^{(i)}$.

3 General Case

In this section we will generalise the laissez-faire investment to obtain the aggregating algorithm.

3.1 Games and Losses

Let us introduce a general prediction scenario. Suppose that outcomes ω_t occur sequentially in discrete time: $\omega_1, \omega_2, \dots$. We assume that outcomes are drawn from an *outcome space* Ω . The learner outputs a prediction γ_t before seeing ω_t . Predictions may be taken from a *prediction space* Γ . The quality of the predictions is measured by a *loss function* $\lambda : \Omega \times \Gamma \rightarrow [0, +\infty)$. We want

to minimise the cumulative loss $\text{Loss}(T) = \sum_{i=1}^T \lambda(\omega_t, \gamma_t)$. The triple $(\Omega, \Gamma, \lambda)$ describes a prediction environment; we will refer to it as a *game*.

One may think of prediction γ_t as of an action taken in a situation of uncertainty. As the uncertainty is lifted, the learner confronts the reality ω_t and faces the consequences of its action represented by $\lambda(\omega_t, \gamma_t)$. We will however keep the prediction terminology.

The investment scenario with M shares can be interpreted in this framework. Let the space of outcomes be $\Omega = [0, +\infty)^M$, the space of predictions be the simplex $\Gamma = \{\gamma = (\gamma_0, \gamma_1, \dots, \gamma_M) \in \mathbb{R}^M \mid \sum_{i=0}^{M-1} \gamma_j = 1\}$, and the loss function be $\lambda(\omega, \gamma) = -\ln\langle \omega, \gamma \rangle$. This game is called Cover's game. The cumulative loss is the negative logarithm of the wealth: $\text{Loss}(T) = -\ln W_t$.

Simpler examples are provided by games with finite outcome spaces. In binary games there are two possible outcomes, $\Omega = \{0, 1\}$, and predictions can be drawn from the unit interval $[0, 1]$. The *square-loss game* has the loss function $\lambda(\omega, \gamma) = (\omega - \gamma)^2$, the *absolute-loss game* has the loss function $\lambda(\omega, \gamma) = |\omega - \gamma|^2$, and the *logarithmic-loss game* has

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

The logarithmic-loss game corresponds to sequential estimation discussed above. The outcome ω_t can be interpreted at k_t and the prediction γ_t as $\gamma_{t,1}$.

In the problem of prediction with expert advice we have access to predictions of N experts $\mathcal{E}^{(1)}, \mathcal{E}^{(2)}, \dots, \mathcal{E}^{(N)}$ and want to suffer loss not much greater than the loss of each expert, i.e., to achieve $\text{Loss}(T) \lesssim \text{Loss}_{\mathcal{E}_i}(T)$.

3.2 Extension of Laissez-Faire Investment

For the sequential investment scenario the laissez-faire strategy achieves the wealth $W_t \geq p^{(i)}W_t^{(i)}$ for all $i = 1, 2, \dots, N$ and $t = 1, 2, \dots$. Taking the logarithm yields $\text{Loss}(t) \leq \text{Loss}_{\mathcal{E}_i}(t) + \ln(1/p^{(i)})$.

Let us try and extend the algorithm to an arbitrary game $(\Omega, \Gamma, \lambda)$. We need to consider notional 'wealth' instead of loss. It is natural to define it as the exponent of $-\text{Loss}$. What shall be the basis of this exponent however? Let us postpone this decision till later and consider a parameter $\eta > 0$. Take $W_t^{(i)} = e^{-\eta \text{Loss}_{\mathcal{E}_i}(t)}$ for expert \mathcal{E}_i . After step $t + 1$ it changes by the factor $e^{-\eta \lambda(\omega_{t+1}, \gamma_{t+1}^{(i)})}$.

Let the learner share its initial wealth of $e^{-\eta \cdot 0} = 1$ among the experts so that expert \mathcal{E}_i gets $p^{(i)}$. The 'laissez-faire investment' achieves $W_t = \sum_{i=1}^N p^{(i)}W_t^{(i)}$ (cf. the pseudo-aggregating algorithm from [VW98] and Lemma 1). Let us turn this notional method into a real practical strategy by 'translating' its terms into the 'prediction' language.

In laissez-faire investment we maintain a set of weights for the experts. The weights before step t correspond to the share of our wealth currently possessed by expert \mathcal{E}_i and can be obtained by normalising $p^{(i)}W_{t-1}^{(i)}$ so that they sum up

to 1. This can be easily expressed in prediction terms: $p_{t-1}^{(i)}$ are obtained by normalising $p^{(i)}e^{-\eta \text{Loss}_{\mathcal{E}_i}(t-1)}$ so that $\sum_{i=1}^N p_{t-1}^{(i)} = 1$.

In the sequential investment scenario we simply took our investment decision to be the weighted sum of experts' decisions: $\gamma_t = \sum_{i=1}^N p_{t-1}^{(i)} \gamma_t^{(i)}$. We do not have this linearity here. Let us write down a suitable replacement.

What we really need is to achieve the 'wealth' greater than or equal to W_t , i.e., $e^{-\eta \text{Loss}(t)} \geq W_t$. One way to ensure this is to maintain the inequality $e^{-\eta \lambda(\omega_t, \gamma_t)} \geq W_t/W_{t-1}$ for all t . We have $W_{t-1} = \sum_{i=1}^N p_{t-1}^{(i)} W_{t-1}$, where $p_{t-1}^{(i)} W_{t-1}$ is the 'money in possession of' expert \mathcal{E}_i ; on the next step this share of the money grows to $p_{t-1}^{(i)} W_{t-1} e^{-\eta \lambda(\omega_t, \gamma_t^{(i)})}$ and therefore

$$\begin{aligned} \frac{W_t}{W_{t-1}} &= \frac{\sum_{i=1}^N p_{t-1}^{(i)} W_{t-1} e^{-\eta \lambda(\omega_t, \gamma_t^{(i)})}}{W_{t-1}} \\ &= \sum_{i=1}^N p_{t-1}^{(i)} e^{-\eta \lambda(\omega_t, \gamma_t^{(i)})} . \end{aligned}$$

The inequality $e^{-\eta \lambda(\omega_t, \gamma_t)} \geq W_t/W_{t-1}$ thus transforms into

$$e^{-\eta \lambda(\omega_t, \gamma_t)} \geq \sum_{i=1}^N p_{t-1}^{(i)} e^{-\eta \lambda(\omega_t, \gamma_t^{(i)})}$$

or, equivalently, into

$$\lambda(\omega_t, \gamma_t) \leq -\frac{1}{\eta} \ln \sum_{i=1}^N p_{t-1}^{(i)} e^{-\eta \lambda(\omega_t, \gamma_t^{(i)})} .$$

Note that we do not know ω_t when we choose γ_t so we must ensure that the inequality holds for *all* $\omega_t \in \Omega$.

Depending on the game and particular values of the variables, this may be possible or not. One can easily formulate a sufficient condition for the existence of γ_t .

Consider $[0, +\infty]^\Omega$, the linear space of all functions from Ω to $[0, +\infty]$ (if Ω is finite, it is finite-dimensional). The game $(\Omega, \Gamma, \lambda)$ defines subsets $P = \{f : \Omega \rightarrow [0, +\infty] \mid \exists \gamma \in \Gamma : f(\cdot) = \lambda(\omega, \cdot)\}$ and $S = \{f : \Omega \rightarrow [0, +\infty] \mid \exists g \in P \forall \omega \in \Omega : f(\omega) \geq g(\omega)\}$. The elements of P may be identified with predictions γ : each element of P may be thought as possible losses of a prediction. The elements of S are called *superpredictions*: they majorise elements of P . Now consider the set $e^{-\eta S}$; it consists of functions $e^{-\eta f}$, where $f \in S$. If this set is convex, the game is called η -mixable.

Suppose that our game is η -mixable. One can easily check that for every array of non-negative weights $p^{(1)}, p^{(2)}, \dots, p^{(N)}$ summing to 1 and every array of experts' predictions $\gamma^{(1)}, \gamma^{(2)}, \dots, \gamma^{(N)}$ there is $\gamma \in \Gamma$ such that

$$\lambda(\omega, \gamma) \leq -\frac{1}{\eta} \ln \sum_{i=1}^N p^{(i)} e^{-\eta \lambda(\omega, \gamma^{(i)})} .$$

for all $\omega \in \Omega$.

We can now formulate the aggregating algorithm for η -mixable games. Assign to the experts initial weights $p^{(i)}$. On step t calculate the weights $p_{t-1}^{(i)}$ for the experts and find any γ_t satisfying

$$\lambda(\omega, \gamma_t) \leq -\frac{1}{\eta} \ln \sum_{i=1}^N p_{t-1}^{(i)} e^{-\eta \lambda(\omega, \gamma_t^{(i)})}$$

for all $\omega \in \Omega$. By reverting the above calculations one can easily check that the loss of this algorithm satisfies

$$\begin{aligned} e^{-\eta \text{Loss}(t)} &\geq \sum_{i=1}^N p^{(i)} e^{-\eta \text{Loss}_{\mathcal{E}_i}(t)} \\ &\geq p^{(i)} e^{-\eta \text{Loss}_{\mathcal{E}_i}(t)} \end{aligned}$$

and

$$\text{Loss}(t) \leq \text{Loss}_{\mathcal{E}_i}(t) + \frac{1}{\eta} \ln \frac{1}{p^{(i)}}$$

for all $i = 1, 2, \dots, N$.

Suppose that the game is not η -mixable, i.e., the convex hull $\mathcal{H}(e^{-\eta S})$ is not a subset of $e^{-\eta S}$, or, equivalently, that $-\frac{1}{\eta} \ln \mathcal{H}(e^{-\eta S})$ is not a subset of S .

The set S is defined in such a way that for every $c > 1$ we have $cS \subseteq S$. The set $\frac{1}{c}S$ includes S and may be larger than S . It may even include $-\frac{1}{\eta} \ln \mathcal{H}(e^{-\eta S})$.

Let us assume that S is closed and we can take the minimum of all $c \geq 1$ such that $-c\frac{1}{\eta} \mathcal{H}(e^{-\eta S}) \subseteq S$. This minimum is denoted by $c(\eta)$. Now for every array of non-negative weights $p^{(1)}, p^{(2)}, \dots, p^{(N)}$ summing to 1 and every array of experts' predictions $\gamma^{(1)}, \gamma^{(2)}, \dots, \gamma^{(N)}$ there is $\gamma \in \Gamma$ such that

$$\lambda(\omega, \gamma) \leq -c(\eta) \frac{1}{\eta} \ln \sum_{i=1}^N p^{(i)} e^{-\eta \lambda(\omega, \gamma^{(i)})} .$$

for all $\omega \in \Omega$ and the aggregating algorithm with a natural modification achieves the loss

$$\text{Loss}(t) \leq c(\eta) \text{Loss}_{\mathcal{E}_i}(t) + c(\eta) \frac{1}{\eta} \ln \frac{1}{p^{(i)}}$$

for all $i = 1, 2, \dots, N$.

One can show (cf. [Vov98]) that for the square-loss game $c(\eta) = 1$, i.e., the game is η -mixable, for $0 < \eta \leq 2$. The absolute-loss game is not mixable for any $\eta > 0$; we have

$$c(\eta) = \eta / \left(2 \ln \frac{2}{1 + e^{-\eta}} \right) > 1 .$$

The logarithmic-loss and Cover's game are mixable for $0 < \eta \leq 1$ and the laissez-faire bound cannot be improved by the aggregating algorithm.

3.3 Optimality

We have constructed some algorithms so far and obtained performance bounds for them. We have not explained why those bounds and the algorithms are good in any sense. In [Vov98] an optimality result is obtained.

References

- [CO96] T. M. Cover and E. Ordentlich. Universal portfolios with side information. *IEEE Transactions on Information Theory*, 42(2), 1996.
- [Vov90] V. Vovk. Aggregating strategies. In *Proceedings of the 3rd Annual Workshop on Computational Learning Theory*, pages 371–383, San Mateo, CA, 1990. Morgan Kaufmann.
- [Vov98] V. Vovk. A game of prediction with expert advice. *Journal of Computer and System Sciences*, 56:153–173, 1998.
- [VW98] V. Vovk and C. J. H. C. Watkins. Universal portfolio selection. In *Proceedings of the 11th Annual Conference on Computational Learning Theory*, pages 12–23. ACM Press, 1998.