Practical Investment with the Long-Short Game

Najim Al-baghdadi[1], Yuri Kalnishkan[1], David Lindsay[2], and Siân Lindsay[2]

[1]Department of Computer Science, Royal Holloway, University of London, Egham, Surrey, TW20 0EX, United Kingdom, email: najim.al-baghdadi.2019@live.rhul.ac.uk, yuri.kalnishkan@rhul.ac.uk

[2]AlgoLabs, Ocean House, Bracknell, Berkshire, RG12 1AX, United Kingdom, email: {david, sian}@algolabs.com

April 25, 2023

Abstract

In this paper we apply methods of prediction with expert advice to real-world foreign exchange trading data with the aim of finding effective investment strategies. We start with the framework of the long-short game, introduced by Vovk and Watkins (1998), and then propose modifications aimed at improving the performance with respect to standard portfolio performance indicators.

1 Introduction

Since modern portfolio theory was first introduced by Markowitz [1952], the problem of portfolio selection has become increasingly prominent. We approach this problem using the framework of on-line learning and apply methods of prediction with expert advice, where an investor makes investment decisions based on the observations of a pool of investors' strategies.

A well-known formalisation of the investment process is Cover's game (section 2.2), where an investor partitions the available money between the assets. Cover and Ordentlich [1996] construct a universal investment strategy for Cover's game: it performs nearly as well as any constant rebalanced portfolio. This approach is a special case of the more general aggregating algorithm, which is capable of combining a finite or infinite pool of portfolios.

The aggregating algorithm can be applied to a general problem of prediction with expert advice and is an evolution of the weighted majority algorithm introduced by Littlestone and Warmuth [1994]. The aggregating algorithm was developed by Vovk [1990, 1998] to include a more general concept of a loss function on prediction and outcome spaces. Given a series of predictions from a pool of experts, a learner following the aggregating algorithm assigns to each expert weights quantifying its trust in them and then combines experts' predictions according to the weights.

The framework of Cover's game is popular in mathematical finance, but it is very restrictive and does not capture important aspects of trading such as short positions and investments on a margin. Vovk and Watkins [1998] took steps to consider more realistic trading scenarios. A trader does not partition their money between the assets as in Cover's game. Instead they open positions, long and short, within some limits set by the exchange or the intermediary providing market access. Vovk and Watkins [1998] introduce a modification of Cover's game, namely, the long-short game (see Section 2.3 of this paper). It admits long positions exceeding the trader's capital (within specified limits) and short positions. A major feature of this framework is the possibility of bankruptcy. While in Cover's game the investor may lose all their capital only in a totally unlikely event of all stocks simultaneously plunging to zero, with the long-short game losing all the money is a much more realistic prospect.

In this paper, we apply the aggregating algorithm to the long-short game in the case of the currency exchange market. The experts are based on the trading activity of 100 clients who used demo trading accounts to trade so-called basket orders of 55 of the most liquid currency pairs during the period from September 2019 to January 2020. We describe a method of deriving predictions from raw trade data using the data staging algorithm DAPRA [Al-baghdadi et al., 2019]. We evaluate the performance of the aggregating algorithm at the long-short game and in Section 4 we propose modifications aimed at improving the practical performance of the resulting portfolio.

Substantial literature exists on applications of prediction with expert advice to investment, but it usually concentrates on Cover's game with no short positions or uses different techniques and approaches. Helmbold et al. [1998] and Györfi et al. [2008] carry out extensive computational experiments with universal strategies competitive with constant rebalanced portfolios (no short positions are allowed). Zhang and Yang [2017], Yang et al. [2020] consider portfolio selection methods based on weak aggregating algorithm merging finite and infinite pools of experts. In their computational experiments real stock market data is used but without short positions.

V'yugin and Trunov [2012] consider universal investment strategies involving short positions and carry out computational experiments. The methods used by V'yugin and Trunov [2012] to construct universal strategies are based on calibration and defensive forecasting. V'yugin et al. [2017] apply a different class of prediction with expert advice methods, namely, AdaHedge-type algorithms, to stock trading in a different context. The algorithms are used to predict stock values and then predictions are fed to other automated trading algorithms.

The main contributions of the paper, is the introduction of return scaling and downside loss. These being modifications to the loss function of the Long-Short game. This is also the first time that we consider the use of combined loss, within the frame work of the aggregating algorithm.

The organisation of the paper is as follows. Section 2 reviews the key definitions and aggregating algorithm. Section 3 introduces a novel data set based on the currency market trades of 100 clients over a 4-month period. We then describe the application of the aggregating algorithm on the long-short game, resulting in a portfolio with unimpressive results and motivates us to propose modifications to the loss function which are detailed in Section 4. In Section 5 we apply the modified method to the data and discuss the improvements made.

2 Preliminaries

In this section we will introduce the framework of the aggregating algorithm and the games formalising investment with expert advice, as discussed by Vovk and Watkins [1998]. In this paper we mainly use the framework of the long-short game. The way we approach it requires a generalisation, which is covered in Section 2.4. We then discuss two extensions of the aggregating algorithm, sleeping experts and discunting, mostly following Kalnishkan [2022].

2.1 Aggregating Algorithm

We will introduce the aggregating algorithm following Vovk and Watkins [1998].

A *game* a $\mathfrak{G}$ is a triple $\langle \Omega, \Gamma, \lambda \rangle$ consisting of an *outcome space* Ω, a *prediction space* Γ, and a *loss function* $\lambda : \Omega \times \Gamma \to [0, +\infty]$.

The outcomes $\omega_1, \omega_2, \ldots$ occur in succession. A prediction strategy $\mathcal{S}$ outputs a prediction γ_t before seeing each outcome ω_t and suffers loss $\lambda(\omega_t, \gamma_t)$ after the outcome ω_t is revealed. The performance of the strategy over T steps is measured by the *cumulative loss* $\mathrm{Loss}_T(\mathcal{S}) = \sum_{t=1}^{T} \lambda(\omega_t, \gamma_t)$. In the investment scenarios, the semantics is slightly different. The value γ_t represents the decision taken on step t and $\lambda(\omega_t, \gamma_t)$ quantifies the consequences of γ_t facing the developments represented by ω_t. An investor is not aiming to make γ_t "close" to the values of ω_t in any sense, but still wants to minimise the cumulative loss. We will retain the prediction terminology though.

The aggregating algorithm (AA) is a way of making a prediction based on the predictions provided by a pool of *experts* (prediction strategies) Θ, where $\gamma_t(\theta) \in \Gamma$ denotes the prediction of expert $\theta \in \Theta$ at trial t. The AA treats experts as black boxes but has access to their predictions $\gamma_t(\theta)$ before making its own prediction γ_t.

The AA takes the following parameters: a learning rate $\eta > 0$ and an initial distribution on experts $P_0(d\theta)$, which quantifies the initial trust in each expert. We will denote the prediction strategy using AA with parameters η and P_0 by $\mathrm{AA}(\eta, P_0)$.

The AA maintains weights P_t on experts Θ. After each trial t, the experts' weights are updated as follows:

$$P_t(d\theta) = e^{-\eta\lambda(\omega_t, \gamma_t(\theta))} P_{t-1}(d\theta) \ .$$

Therefore the larger the expert's loss the greater the reduction of its weight. To define the AA we first will introduce the aggregating pseudoalgorithm (APA), which at trial t produces a *generalised prediction* (a function $g : \Omega \to (-\infty, +\infty])$ based on the normalised weights as follows:

$$g_t(\omega) = -\frac{1}{\eta}\ln \int_\Theta e^{-\eta\lambda(\omega_t,\gamma_t(\theta))} P^*_{t-1}(d\theta) \ ,$$

where P^*_{t-1} denotes the normalised weights $P^*_{t-1} = P_{t-1}(d\theta)/P_{t-1}(\Theta)$. One can define the cumulative loss of APA as $\mathrm{Loss}(\mathrm{APA}(\eta, P_0)) = \sum_{t=1}^{T} g_t(\omega_t)$. The following lemma can be proven by induction.

Lemma 2.1. *For any learning rate $\eta > 0$, initial distribution P_0, and $T = 1, 2, ...$ we get*

$$\mathrm{Loss}_T(\mathrm{APA}(\eta, P_0)) = -\frac{1}{\eta} \ln \int_\Theta e^{-\eta \mathrm{Loss}_T(\theta)} P_0(d\theta).$$

for all $\omega_1, \omega_2, \ldots, \omega_T$.

To obtain the AA from the APA, we need to find a permitted prediction $\Sigma(g_t)$, where the *substitution function* Σ maps a generalised prediction $g : \Omega \to (-\infty, \infty]$ to a prediction $\Sigma(g) \in \Gamma$ while keeping the loss as low as possible. Let $\mathrm{GA}(\eta)$ be the set of all generalised actions that can be produced by the APA with learning rate η:

$$\mathrm{GA}(\eta) = \left\{ g : \Omega \to \mathbb{R} \mid \exists P \, \forall \omega : g(\omega) = -\frac{1}{\eta} \ln \int_\Gamma e^{-\eta\lambda(\omega,\gamma)} P(d\gamma) \right\} \ ,$$

where P ranges over all distributions[1] on Γ. For a generalised action g, let

$$C(g) = \inf_{\gamma\in\Gamma} \sup_{\omega\in\Omega} \lambda(\gamma, \omega)/g(\omega) \ .$$

Under mild continuity assumptions (namely, if Γ is a compact topological space and $\lambda(\omega, \cdot)$ is continuous in the second argument) the value of $C(g)$ is *achieved* on some $\gamma \in \Gamma$ and we can replace g by γ such that $\lambda(\omega, \gamma) \le C(g)g(\omega)$ for every $\omega \in \Omega$.

The mixability constant C_η is defined as

$$C_\eta = \sup_{g\in \mathrm{GA}(\eta)} C(g) \ .$$

There is a substitution function Σ mapping generalised predictions g to Γ satisfying:

$$\forall g \in \mathrm{GA}(\eta) \, \forall \omega \in \Omega : \lambda(\omega, \Sigma(g)) \le C_\eta g(\omega) \ . \tag{1}$$

Substitution functions satisfying condition (1) are the ones allowed to be used in the AA. Condition (1) and Lemma 2.1 imply

$$\begin{aligned} \mathrm{Loss}_T(\mathrm{AA}(\eta, P_0) &\le C_\eta \mathrm{Loss}_T(\mathrm{APA}(\eta, P_0)) \\ &= -\frac{C_\eta}{\eta} \ln \int_\Theta e^{-\eta \mathrm{Loss}_T(\theta)} P_0(d\theta) \ . \end{aligned} \tag{2}$$

A game is said to be *η-mixable* if $C_\eta = 1$ and *mixable* if it is η-mixable for some $\eta > 0$. For mixable games the learner following the AA can perform almost as well as any expert from a finite pool, as the following lemma shows.

[1] We assume that all $\gamma_t(\cdot) : \Theta \to \Gamma$ are Borel mappings of topological spaces and induce measures on Γ.

Lemma 2.2. *For a finite pool of experts* Θ*,*

$$\mathrm{Loss}_T(\mathrm{AA}(\eta, P_0)) \leq C_\eta \mathrm{Loss}_T(\theta) + \frac{C_\eta}{\eta} \ln \frac{1}{P_0(\theta)} \tag{3}$$

for every expert $\theta \in \Theta$ *and time* $T = 1, 2, \ldots$ *Moreover, if the game is* η*-mixable, then*

$$\mathrm{Loss}_T(\mathrm{AA}(\eta, P_0)) \leq \mathrm{Loss}_T(\theta) + \frac{1}{\eta} \ln \frac{1}{P_0(\theta)} \quad . \tag{4}$$

Indeed, for a finite pool of experts, the integral in (2) turns into a sum of non-negative terms and the sum can be bounded from below by each of its terms.

Bounds (3) provided by the AA are optimal in their class for the uniform initial distribution Vovk [1998]. If an algorithm provides a guarantee of this type, the AA with some η can do the same or better; hence the significance of the AA.

Taking η such that $C_\eta = 1$ minimises the first term on the right-hand side of (3); this is important because this term may be growing with T (and would normally grow as a linear rate). Thus ηs making the game mixable are good choice for practice as long as they exist. Out of such ηs the maximum value should be chosen because in minimises the second term on the right-hand side of (4).

Algorithm 1 is the AA for the case of finitely many experts, which is the main case for this paper. Here N is the number of experts and Θ is identified with $\{1, 2, \ldots, n\}$. In this context, one can think of Σ as a function $\Gamma^N \times \mathbb{P}_{N-1} \to \Gamma$ (where $\mathbb{P}_{N-1}$ is the $(N-1)$-simplex) mapping arrays of experts' predictions and distributions on them to predictions.

Algorithm 1 The aggregating algorithm for finitely many experts

Parameters: Learning rate $\eta > 0$ and
initial experts' weights $p_0(n)$, $n = 1, 2, \ldots, N$
for $t = 1, 2, \ldots$ **do**
  read experts' predictions $\gamma_t(n) \in \Gamma$, $n = 1, 2, \ldots, N$
  normalise the weights: $p^*_{t-1}(n) = p_{t-1}(n) / \sum_{n=1}^N p_{t-1}(n)$
  produce the generalised prediction

$$g_t(\omega) = -\frac{1}{\eta} \ln \sum_{n=1}^N p^*_{t-1}(n) e^{-\eta\lambda(\omega, \gamma_t(n))}$$

  calculate and output $\gamma_t = \Sigma(g_t)$
  read $\omega_t \in \Omega$
  update the weights $p_t(n) = p_{t-1}(n) e^{-\eta\lambda(\omega_t, \gamma_t)}$
end for

2.2 Cover's Game

Cover's game formalises a basic investment scenario. We include Cover's game in our discussion for completeness.

Cover's game describes investment into a market of M assets. The outcome space Ω describes the behaviour of the market with the non-negative price relative vector $\omega = (\omega[0], ..., \omega[M-1]) \in \Omega = [0, \infty)^M$, where $\omega_t[m]$ represents the ratio of the value of asset m at trial t to the value at trial $t-1$. If $S_t[m]$ denotes the price of asset m at time t, then $\omega_t[m] = S_{t+1}[m]/S_t[m]$. An investment in this market is represented by the m-dimensional portfolio vector γ, where $\gamma[m]$ denotes the proportion of the investor's wealth invested in asset m. In Cover's game we assume that all wealth is invested on every step and no short positions or trading on credit is allowed; in other terms, $\gamma[m] \geq 0$ for $m = 0, 1, \ldots, M-1$, and $\sum_{m=0}^{M-1} \gamma[m] = 1$. The prediction space Γ is the $(M-1)$-simplex $\mathbb{P}_{M-1}$. One can say that the investor partitions the wealth between M assets.

If an investor makes an investment γ and then outcome ω occurs, the investor's wealth changes by a factor of $\langle \omega, \gamma \rangle = \sum_{m=1}^{M-1} \omega[m]\gamma[m]$. In order to link this with the additive framework of prediction games, we define the loss function by $\lambda(\omega, \gamma) := -\ln\langle \omega, \gamma \rangle$. If the investor starts from wealth of $W_0 = 1$ and follows an investment strategy $\mathcal{S}$, then the wealth after step T equals

$$W_T = \prod_{t=1}^{T} \langle \omega_t, \gamma_t \rangle = e^{-\mathrm{Loss}_T(\mathcal{S})} \ .$$

Let us discuss mixability of Cover's game. Note that the game does not quite follow our framework as $\lambda(\omega, \gamma)$ may take negative values. Indeed, the scalar product $\langle \omega, \gamma \rangle$ may well exceed 1 leading to negative loss (in this situation the investor actually earns money). Following Vovk and Watkins [1998], consider a modified loss function

$$\tilde{\lambda}(\omega, \gamma) = \lambda(\omega, \gamma) - \inf_{\delta} \lambda(\omega, \delta) \ . \tag{5}$$

The new loss function is non-negative by construction. The following lemma describes the properties of the game.

Lemma 2.3 (Vovk and Watkins [1998]). *For every $\eta \leq 1$, $C_\eta = 1$. Moreover, for every $\eta \leq 1$ and every $g \in \mathrm{GA}(\eta)$, $C(g) = 1$. The only prediction attaining $C(g) = 1$ is the average*

$$\gamma^* := \int_\Gamma \gamma P(d\gamma) \ , \tag{6}$$

where P is a probability distribution in Γ generating g:

$$g(\omega) = -\frac{1}{\eta} \ln \int_\Gamma e^{-\eta\lambda(\omega,\gamma)} P(d\gamma) \ .$$

Remark 1. *Let $\tilde{\lambda}(\omega, \gamma) = \lambda(\omega, \gamma) + f(\omega)$ for some fucntion of ω. Then the weights $p_t(n)$ in Algorithm 1 calculated for these two functions are proportional and the normalised weights $p_t^*(n)$ are identical. Indeed, all experts suffer loss according to the same ω and $f(\omega)$ drops in normalisation.*

Remark 2. *Now let $C_\eta = 1$ in (1) and a substitution rule Σ (considered as a function of predictions $\gamma_1, \ldots, \gamma_N$ and a distribution on them) works w.r.t. the loss function $\tilde{\lambda}$. Then the same rule satisfies (1) for the original λ, which may take negative values. Indeed, $f(\omega)$ cancels out on both the sides.*

The remarks imply that we can apply the AA for Cover's game with $\eta \in (0, 1]$ using the weighted average as the substitution rule.

Remark 3. *When $C_\eta = 1$ and (4) holds for $\tilde{\lambda}$, the terms $f(\omega_t)$ on both the sides of (4) cancel out so we get the bound for cumulative losses in terms of the original loss function λ.*

In the case of $\eta = 1$ it is easy to prove that the average of experts' predictions has the desired properties directly. Linearity of the scalar product implies

$$g(\omega) = -\ln \int_\Gamma e^{-\lambda(\omega,\gamma)} P(d\gamma) = -\ln \int_\Gamma \langle \omega, \gamma \rangle P(d\gamma) = \\ -\ln \left\langle \omega, \int_\Gamma \gamma P(d\gamma) \right\rangle \ .$$

For larger values of η we do not get mixability.

Lemma 2.4 (Vovk and Watkins [1998]). *When $\eta > 1$, $C_\eta = \eta$.*

We see that $\eta = 1$ is the optimal choice of the parameter and the substitution rule

$$\gamma_t = \sum_{n=1}^{N} p^*_{t-1}(n) \gamma_t(n) \tag{7}$$

should be used with Cover's game in the finite case.

2.3 Long-Short Game

The long-short game is a modification of Cover's game aimed at a more general and more realistic trading scenario. A trader is usually allowed to open positions, both long and short, within certain limits based on their deposit and money they had earned previously. The limits are aimed to minimise the chances of bankruptcy so that the intermediary providing access to the market could avoid handling the consequences of the trader's default.

In a *bounded long-short game* with the *prudence parameter* $a > 0$ an investment decision is represented by a vector $\gamma \in \mathbb{R}^M$ such that

$$\|\gamma\|_1 = |\gamma[0]| + \ldots + |\gamma[M-1]| \le a \ ; \tag{8}$$

in other terms, $\Gamma \subseteq \mathbb{R}^M$ is a ball w.r.t. the $\|\cdot\|_1$-norm. The intuitive interpretation of γ is as follows. Suppose that before step t the trader has wealth $W_{t-1} > 0$. Then on step t the trader opens positions of size $W_{t-1}\gamma_t[m]$, $m = 0, 1, 2, \ldots, M-1$ (long or short depending on the sign of $\gamma_t[m]$) in assets $0, 1, \ldots, M-1$. The sum of the sizes of the positions is bounded by $W_t a$.

It is more convenient to represent outcomes by a vector of returns here, so $\omega_t[m] = (S_t[m] - S_{t-1}[m])/S_{t-1}[m] = S_t[m]/S_{t-1}[m] - 1 \ge -1$. Thus on the position in asset m the trader gets the profit of $W_{t-1}\omega_t[m]\gamma_t[m]$ and the overall trader's wealth changes according to

$$W_t = W_{t-1}(1 + \langle \omega_t, \gamma_t \rangle) \ .$$

We let

$$\lambda(\omega, \gamma) = -\ln(1 + \langle \omega, \gamma \rangle) \ .$$

Note that for some values of ω the expression $1 + \langle \omega, \gamma \rangle$ can go below zero; the trader then goes bankrupt and the expression $-\ln(1 + \langle \omega, \gamma \rangle)$ is undefined. In a bounded game we assume this never happens because all ωs satisfy

$$\|\omega\|_\infty = \max_{m=0,1,\ldots,M-1} |\omega_m| \leq \frac{1}{a} \ . \tag{9}$$

Thus the outcome space in the a-bounded game is the intersection of $[-1, +\infty)^M$ with the $\|\cdot\|_\infty$ ball.

For the analysis of mixability, we need to modify loss function by (5). The following lemma holds for every a-bounded game with $\tilde{\lambda}$.

Lemma 2.5 (Vovk and Watkins [1998]). *For any a-bounded game, $a > 0$, and for every $\eta \leq 1$, $C_\eta = 1$. Moreover, for every $\eta \leq 1$ and every $g \in GA(\eta)$, $C(g) = 1$. The only prediction attaining $C(g) = 1$ is the average (6), where as before P is a probability distribution in Γ generating g. When $\eta > 1$, $C_\eta > 1$.*

As for Cover's game, Remarks 1–3 imply that for $\eta = 1$ we can apply the AA with the original loss function using the weighted average as the substitution rule.

2.4 General Long-Short Game

Although Vovk and Watkins [1998] restrict their attention to the bounded long-short game, its practical applications suffer from an important problem. While bound (8) on the norm of γ is realistic (and can be linked to the restrictions imposed by the market access provider), bound (9) on the norm of ω cannot be guaranteed; the market is not under our control in any way. In this section we will discuss the general long-short game.

Consider the general long-short game with $\Gamma = \mathbb{R}^m$, $\Omega = [-1, +\infty)^M$ and the loss function given by

$$\lambda_{\mathrm{LS}}(\omega, \gamma) = \begin{cases} -\ln(1 + \langle \omega, \gamma \rangle), & \text{if } 1 + \langle \omega, \gamma \rangle > 0 \\ +\infty & \text{otherwise} \end{cases}$$
$$= -\ln(\max(1 + \langle \omega, \gamma \rangle, 0)) \ ,$$

where $\ln 0 = -\infty$. Our earlier analysis suggests that we should apply the AA using $\eta = 1$ and the weighted average of expert' predictions as the substitution rule. We formulate the case for finitely many experts (i.e., with a finite pool of experts Θ of size N) as Algorithm 2.

Let us formulate its properties.

Lemma 2.6. *Suppose that the outcomes ω_t and experts' predictions $\gamma_t(n)$ are such that $1 + \langle \omega_t, \gamma_t(n) \rangle \geq 0$ for all $t = 1, 2, \ldots, T$, $n = 1, 2, \ldots, N$. Then the predictions $\gamma_t = \sum_{n=1}^N p^*_{t-1}(n) \gamma_t(n)$ satisfy $1 + \langle \omega_t, \gamma_t \rangle \geq 0$ and (4) holds with $\eta = 1$:*

$$\mathrm{Loss}_T(\mathrm{AA}(1, P_0)) \leq \mathrm{Loss}_T(n) + \ln \frac{1}{p_0(n)} \tag{10}$$

for every expert $n = 1, 2, \ldots, N$.

Algorithm 2 The aggregating algorithm for the general long-short game

Parameters: initial distribution P_0 over experts, $p_0(n)$, $n = 1, 2, \ldots, N$
for $t = 1, 2, \ldots$ **do**
  read experts' predictions $\gamma_t(n) \in \Gamma$, $n = 1, 2, \ldots, N$
  normalise the weights: $p^*_{t-1}(n) = p_{t-1}(n) / \sum_{n=1}^N p_{t-1}(n)$
  calculate and output $\gamma_t = \sum_{n=1}^N p^*_{t-1}(n)\gamma_t(n)$
  read $\omega_t \in \Omega$
  update the weights $p_t(n) = p_{t-1}(n)e^{-\lambda \mathrm{LS}(\omega_t, \gamma_t)}$
end for

Note that conditions $1 + \langle \omega_t, \gamma_t(n) \rangle \geq 0$ should hold on the actual realised sequence rather than for all possible ω.

The proof is along the same lines as the general analysis of the AA described in the previous sections. We will give it here for completeness.

Proof. Let us show by induction that

$$e^{-\mathrm{Loss}_s(AA)} \geq \sum_{n=1}^N p_0(n) e^{-\mathrm{Loss}_s(n)} \tag{11}$$

(as a matter of fact, this will hold as an equality). Dropping all terms from the sum on the right-hand side except for one and taking the logarithm of this inequality yields (10).

Let (11) hold for $s = t - 1$. If $1 + \langle \omega_t, \gamma_t(n) \rangle \geq 0$ for all n, then

$$e^{-\lambda \mathrm{LS}(\omega_t, \gamma_t)} = \max(1 + \langle \omega_t, \gamma_t \rangle, 0) \geq \sum_{n=1}^N p^*_{t-1}(n) \max(1 + \langle \omega_t, \gamma_t(n) \rangle, 0) = \sum_{n=1}^N p^*_{t-1}(n) e^{-\lambda \mathrm{LS}(\omega_t, \gamma_t(n))}$$

holds (as an equality) by the linearity of the scalar product.

Multiplying this inequality by (11) with $s = t - 1$ and observing that

$$p^*_{t-1}(n) = \frac{p_{t-1}(n)}{\sum_{i=1}^N p_{t-1}(i)} = \frac{p_0(n) e^{-\mathrm{Loss}_{t-1}(n)}}{\sum_{i=1}^N p_0(i) e^{-\mathrm{Loss}_{t-1}(i)}}$$

completes the inductive step. □

Let, however, $1 + \langle \omega_t, \gamma_t(n) \rangle < 0$ for some n with weights $p^*_{t-1}(n) > 0$ (but let there be values n such that $p^*_{t-1}(n) > 0$ and $1 + \langle \omega_t, \gamma_t(n) \rangle > 0$). In this situation the inequality

$$\max\left(1 + \langle \omega_t, \gamma_t \rangle, 0\right) \geq \sum_{n=1}^N p^*_{t-1}(n) \max(1 + \langle \omega_t, \gamma_t(n) \rangle, 0) \enspace ,$$

where $\gamma_t = \sum_{n=1}^N p^*_{t-1}(n)\gamma_t(n)$, does not hold. Indeed, consider the case $1 + \langle \omega_t, \gamma_t \rangle \leq 0$. In this case we get 0 on the left-hand side and a positive expression on the right. In the opposite case $1 + \langle \omega_t, \gamma_t \rangle > 0$, we get

$$1 + \langle \omega_t, \gamma_t \rangle = 1 + \left\langle \omega_t, \sum_{n=1}^N p^*_{t-1}(n)\gamma_t(n) \right\rangle = \sum_{n=1}^N p^*_{t-1}(n)(1 + \langle \omega_t, \gamma_t(n) \rangle)$$

on the left-hand side and the sum of terms that are the same or greater (by our assumption some are actually greater) on the right-hand side.

The former case is quite hopeless: the strategy following AA goes bankrupt and suffers infinite loss, while some of our experts still have finite. The later case is not. Consider

$$c_t = \frac{\sum_{n=1}^N p^*_{t-1}(n) \max(1 + \langle \omega_t, \gamma_t(n) \rangle, 0)}{1 + \left\langle \omega_t, \sum_{n=1}^N p^*_{t-1}(n) \gamma_t(n) \right\rangle} \ .$$

We have $c_t > 1$, but arguably not by a lot. Although in principle c_t can be arbitrarily high, it is reasonable to expect that even if $1 + \langle \omega_t, \gamma_t(n) \rangle$ is below 0, then not by a lot. Expert n that suffers a bankruptcy is unlikely to have performed very well previously and so its weight $p^*_{t-1}(n)$ should also be small, especially if the pool of experts is large.

We then get

$$c_t e^{-\lambda_{\mathrm{LS}}(\omega_t, \gamma_t)} = c_t \max(1 + \langle \omega_t, \gamma_t \rangle, 0) \geq \sum_{n=1}^N p^*_{t-1}(n) \max(1 + \langle \omega_t, \gamma_t(n) \rangle, 0) = \sum_{n=1}^N p^*_{t-1}(n) e^{-\lambda_{\mathrm{LS}}(\omega_t, \gamma_t(n))}$$

and the term $\ln c_t$ finds its way into the right-hand side of (10), which turns into

$$\mathrm{Loss}_T(\mathrm{AA}(1, P_0)) \leq \mathrm{Loss}_T(n) + \ln \frac{1}{p_0(n)} + {\sum}' \ln c_t \ , \qquad (12)$$

where $\sum'$ is taken over steps when there are bankrupt experts, as long as the strategy following the AA is not bankrupt. Note that an expert going bankrupt worsens the bound for *all* other experts. On the other hand, each expert can contribute to $\sum'$ only once: after it goes bankrupt, its weight is zeroed.

2.5 Specialist Experts

In this section we will introduce specialist following Kalnishkan [2022] after Chernov and Vovk [2009].

A specialist expert is an expert that may refrain from making a prediction on any given trial. In the event an expert makes a prediction the expert is said to be "awake" and if not we say the expert "sleeps". This is a particularly useful approach to take when it comes to finding optimal investment strategies. It is common for investors to have periods without a position in the market, especially in the case of investors with shorter investment time horizons. How can we interpret this? First, the investor's behaviour may be understood as making zero predictions. However, this is not the only possibility. In investment, "doing nothing" often means making a passive investment into an asset perceived as riskless and reliable. This asset is often an index portfolio tapping into the wisdom of the crowd. We thus may think of an expert making no predictions as investing into a kind of index portfolio bases on the behaviour of fellow experts.

Specialist experts provide a natural way to implement the later approach. We assume that a sleeping expert joins the crowd making the same prediction as the learner and therefore suffering the same loss.

Recall Lemma 2.2 with a bound on the loss the aggregating algorithm guarantees. If an expert is sleeping and makes the same prediction as the learner on some turn, the loss terms in (4) would cancel out leaving us with the sums over times when experts are awake to define our learners bound on loss. However, this idea is still hypothetical: to work out the learners prediction one needs to know the experts prediction to begin with. We can solve this by recalling how the learner makes a prediction in AA. The prediction γ_t is chosen to satisfy

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

where ω ranges over Ω. Assuming that the loss of a sleeping expert be equal to that of the learner we get

$$e^{-\eta\lambda(\omega,\gamma_t)} \leq \sum_{n:\ E_n \text{ is awake}} p_{t-1}(n)e^{-\eta\lambda(\omega,\gamma_t(n))} + \sum_{n:\ E_n \text{ is sleeps}} p_{t-1}(n)e^{-\eta\lambda(\omega,\gamma_t)}. \quad (13)$$

We can then subtract the second sum from each side geting

$$e^{-\eta\lambda(\omega,\Sigma(g_t))} \leq \frac{1}{z_t} \sum_{n:\ E_n \text{ is awake}} p_{t-1}(n)e^{-\eta\lambda(\omega,\gamma_t(n))},$$

where

$$z_t = \sum_{n:\ E_n \text{ is awake}} p_{t-1}.$$

2.6 Discounted Loss

Discounting loss is a well-established practise in both on-line and reinforcement learning. It may be used to account for inflation over time or to reflect changes in market conditions. In this section we will explore the use of discounting and how it can be implemented within the aggregating algorithm. We will follow the approach of Chernov and Zhdanov [2010] and Kalnishkan [2022].

Take coefficients $\alpha_1, \alpha_2, ... \in [0,1]$ and define the cumulative discounted loss for a learner as

$$\widetilde{\mathrm{Loss}}_T(\mathrm{AA}(\eta, P_0)) = \sum_{t=1}^{T} \lambda(\omega_t, \gamma_t) \left(\prod_{s=t}^{T-1} \alpha_s \right) = \alpha_{T-1}\widetilde{\mathrm{Loss}}_{T-1}(\mathrm{AA}(\eta, P_0)) + \lambda(\omega_T, \gamma_T) \ .$$

We define the discounted loss of an expert in the same way. In the case where all α_i are equal, $\alpha_1 = \alpha_2 = ... = \alpha$ and $\lambda(\gamma_t, \omega_t)$ comes into the formula with the discounting coefficient α^{t-1}.

Algorithm 3 Specialist Experts learning protocol

Parameters: Learning rate $\eta > 0$ and
initial experts' weights $p_0(n)$, $n = 1, 2, \ldots, N$
for $t = 1, 2, \ldots$ **do**
  read awake experts' predictions $\gamma_t(n) \in \Gamma$, $n = 1, 2, \ldots, N$
  normalise awake experts weights:

$$p^*_{t-1}(n) = p_{t-1}(n) / \sum_{n:\ E_n \text{ is awake}} p_{t-1}(n)$$

  produce the generalised prediction:

$$g_t(\omega) = -\frac{1}{\eta} \ln \sum_{n:\ E_n \text{ is awake}} p^*_{t-1}(n) e^{-\eta\lambda(\omega, \gamma_t(n))}$$

  calculate and output $\gamma_t = \Sigma(g_t)$
  read $\omega_t \in \Omega$
  update the awake weights $p_t(n) = p_{t-1}(n) e^{-\eta\lambda(\omega_t, \gamma_t)}$
  update the sleeping weights $p_t(n) = p_{t-1}(n) e^{-\eta\lambda(\omega, \Sigma(g_t))}$
end for

Let us calculate experts' weight as follows

$$p^\theta_{t-1} \propto p_0(\theta) e^{-\eta\alpha_{t-1}\widetilde{\text{Loss}}_{t-1}(\theta)}$$

Kalnishkan [2022] shows by induction that the following bound on the learner's loss holds

$$\widetilde{\text{Loss}}_T(\text{AA}(\eta, P_0)) \le \widetilde{\text{Loss}}_T(\theta) + \frac{1}{\eta} \ln \frac{1}{P_0(\theta)}. \qquad (14)$$

3 Application of Long-Short Game

In this section we apply the aggregating algorithm to the long-short game, merging the investment strategies of 100 unique clients.

3.1 Data Set

The data we are using is derived from the basket orders of 100 clients using demo trading accounts over a period of 4 months, during September 2019 to January 2020. The data is representative of the behaviour of investors trading in the currency exchange market over a relatively calm period. A basket order allows a group of financial market instruments to be traded simultaneously. Different weighting criteria for different instruments can be used to tailor the basket according to the client's needs. Clients can either trade their baskets manually or use automated models. In this data the clients construct their baskets from the 55 most liquid Foreign Exchange (FX) pairs, as shown in Table 1.

Algorithm 4 The aggregating algorithm with discounting

Parameters: Learning rate $\eta > 0$ and
discounting factors $\alpha_1, \alpha_2, ...$ and
initial experts' weights $p_0(n)$, $n = 1, 2, \ldots, N$
for $t = 1, 2, \ldots$ **do**
  read experts' predictions $\gamma_t(n) \in \Gamma$, $n = 1, 2, \ldots, N$
  normalise the weights:

$$p^*_{t-1}(n) = P_0(n)(p_{t-1}(n))^{\alpha_{t-1}} / \sum_{n=1}^{N} P_0(\theta)(p_{t-1}(n))^{\alpha_{t-1}}$$

  produce the generalised prediction

$$g_t(\omega) = -\frac{1}{\eta} \ln \sum_{n=1}^{N} p^*_{t-1}(n) e^{-\eta\lambda(\omega, \gamma_t(n))}$$

  calculate and output $\gamma_t = \Sigma(g_t)$
  read $\omega_t \in \Omega$
  update the weights $p_t(n) = p_{t-1}^{\alpha_{t-1}}(n) e^{-\eta\lambda(\omega_t, \gamma_t)}$
end for

Table 1: 55 FX (currency) pairs, the symbol format is a pair of currency codes delimited by a "/", where the currency code is in the ISO 4217 format.

AUD/CAD	EUR/AUD	EUR/SGD	HKD/JPY	USD/DKK	CAD/CHF	EUR/JPY	GBP/JPY	NZD/USD	USD/RUB	AUD/USD
AUD/CHF	EUR/CAD	EUR/USD	MXN/JPY	USD/HKD	CAD/JPY	EUR/MXN	GBP/NZD	SGD/JPY	USD/SEK	EUR/HKD
AUD/JPY	EUR/CHF	GBP/AUD	NZD/CAD	USD/JPY	CAD/SGD	EUR/NZD	GBP/SEK	USD/CAD	USD/SGD	GBP/HKD
AUD/NZD	EUR/DKK	GBP/CAD	NZD/CHF	USD/MXN	CHF/JPY	EUR/PLN	GBP/SGD	USD/CHF	USD/TRY	NZD/SGD
AUD/SGD	EUR/GBP	GBP/CHF	NZD/JPY	USD/NOK	CHF/SGD	EUR/SEK	GBP/USD	USD/CNH	USD/ZAR	USD/PLN

Table 2 illustrates raw trade data from the basket orders of four different clients, B1, B2, B3 and B10. We see client B1 has a basket trading 4 different currency pairs (NZD/SGD, GBP/SGD, NZD/CAD and CHF/JPY) where each block of trades all have the same opening timestamp, for example, 24^{th} Oct 2019 at 07:08 and holds that position for 3 days. Later that same day at 19:45 we see client B1 trades the same basket again, so building on their position. "Position" is the summation of the client's trades and at a given point in time is described as being long, flat or short. At 19:45, client B1's position goes long in NZD/SGD by 27,000, and short in NZD/CAD by 41,000. Table 2 further demonstrates that clients have the freedom to trade any combination of currency pairs and with any notional weightings they desire for their baskets (these can be derived manually or using proprietary algorithms). For example, client B3 trades a basket of 5 different currency pairs, whereas client B2 trades larger notional preferring GBP currency crosses. Client B10 solely trades 3 symbols which are all USD crosses.

Before we can apply the AA to this data we must first:

- Normalise client positions into a common currency, in our case we

Table 2: Raw trade data taken from basket orders of four different clients (B1, B2, B3 and B10) on the 24^{th} Oct 2019. Each trade has an open and close timestamp, and corresponding open and close price. Whether the trade was a buy or sell is denoted by a 1 or -1 sign, respectively.

Open Time	Open Price	Client	Amount	Sign	Symbol	Order Id	Close Time	Close Price	Mins In Trade
2019.10.24T07:08:00.000	0.87236	B1	27,000	1	NZD/SGD	B1_87	2019.10.27T22:12:00.000	0.86643	5,224
2019.10.24T07:08:00.000	1.76210	B1	6,000	-1	GBP/SGD	B1_267	2019.10.30T00:20:00.000	1.75285	8,232
2019.10.24T07:08:00.000	0.83763	B1	41,000	-1	NZD/CAD	B1_447	2019.10.30T00:20:00.000	0.83112	8,232
2019.10.24T07:08:00.000	109.76200	B1	8,000	1	CHF/JPY	B1_634	2019.10.27T22:12:00.000	109.30400	5,224
2019.10.24T13:36:00.000	9.61360	B3	7,000	-1	USD/SEK	B3_64	2019.11.04T10:38:00.000	9.64949	15,662
2019.10.24T13:36:00.000	0.93133	B3	30,000	-1	AUD/SGD	B3_230	2019.11.03T22:12:00.000	0.93782	14,916
2019.10.24T13:36:00.000	2.01415	B3	2,000	-1	GBP/NZD	B3_397	2019.10.27T22:27:00.000	2.01888	4,851
2019.10.24T13:36:00.000	1.74062	B3	9,000	-1	EUR/NZD	B3_573	2019.10.28T03:50:00.000	1.74552	5,174
2019.10.24T13:36:00.000	19.08875	B3	1,000	-1	USD/MXN	B3_746	2019.11.05T07:47:00.000	19.13142	16,931
2019.10.24T14:58:00.000	10.07130	B2	451,000	1	GBP/HKD	B2_25	2019.11.01T09:23:00.000	10.09472	11,185
2019.10.24T14:58:00.000	8.70289	B2	59,000	-1	EUR/HKD	B2_84	2019.10.25T11:04:00.000	8.71086	1,206
2019.10.24T14:58:00.000	1.27429	B2	338,000	-1	GBP/CHF	B2_144	2019.10.25T14:54:00.000	1.27440	1,436
2019.10.24T14:58:00.000	12.38930	B2	19,000	-1	GBP/SEK	B2_202	2019.10.25T11:19:00.000	12.40480	1,221
2019.10.24T15:03:00.000	0.99168	B10	167,000	-1	USD/CHF	B10_22	2019.10.27T22:12:00.000	0.99454	4,749
2019.10.24T15:03:00.000	3.85160	B10	95,000	-1	USD/PLN	B10_91	2019.10.27T22:12:00.000	3.85830	4,749
2019.10.24T15:03:00.000	6.72958	B10	749,000	1	USD/DKK	B10_161	2019.10.28T08:11:00.000	6.73906	5,348
2019.10.24T19:45:00.000	0.86990	B1	27,000	1	NZD/SGD	B1_88	2019.10.27T22:12:00.000	0.86643	4,467
2019.10.24T19:45:00.000	1.75211	B1	5,000	-1	GBP/SGD	B1_269	2019.10.30T00:20:00.000	1.75285	7,475
2019.10.24T19:45:00.000	0.83419	B1	41,000	-1	NZD/CAD	B1_449	2019.10.30T00:20:00.000	0.83112	7,475
2019.10.24T19:45:00.000	109.47500	B1	9,000	1	CHF/JPY	B1_636	2019.10.27T22:12:00.000	109.30400	4,467

use USD. We do this because clients trade many different currencies, all whose notional values differ through time. Therefore, to compare the positions and derive a price vector they must be normalised.

- Sample the data at regular time intervals (for this data we chose a resolution of 1 minute) across all clients and currency pairs. This is because whilst clients are at liberty to trade and hold positions for however long they wish, the AA must make a prediction regarding the future behaviour of the market at regular time intervals.

Al-baghdadi et al. [2019] introduced the data staging technique DAPRA (Data Aggregation Partition Reduction Algorithm) which, when applied to data streams pertaining to client trades and market data, allows one to normalise and sample the data as required for this study. DAPRA follows a three step process:

1. **D**ata **A**ggregation, where data from one or more sources of irregularly sampled time series data are combined into a regular sampled time series. Here this will be the client trade and market price data.
2. The aggregated stream is **P**artitioned into regular time intervals. It is also at this stage where derived fields such as position are calculated, using features from both data sets.
3. Finally this aggregated, partitioned stream of data is **R**educed where the aggregated data is grouped based on client and symbol in this study.

Fig. 1 compares the net positions of the first 10 clients in the data set, following DAPRA transformation over the trial period. We can see here that clients clearly have different trading strategies and activity through time, some clients build a position steadily over a period of days and weeks, such as B3 and B4. Whereas others such as B6 and B9, open and close fixed amounts over short periods of time. All clients trade different amounts resulting in net positions between long and short 5 million across all 55 symbols shown in table 1. The positions in Fig. 1 show step changes when trades of basket orders are placed. Returning to our earlier example, we can see the shifts in position related to the trades in Table 2 on the 24^{th} Oct 2019. Hence, we see that client B10's basket order comprises a sell USD/DKK trade which means a position change from flat to short. Importantly we can see great variability in the notional sizes, basket symbol composition and holding periods of the different clients, producing are varied rage of investment strategies.

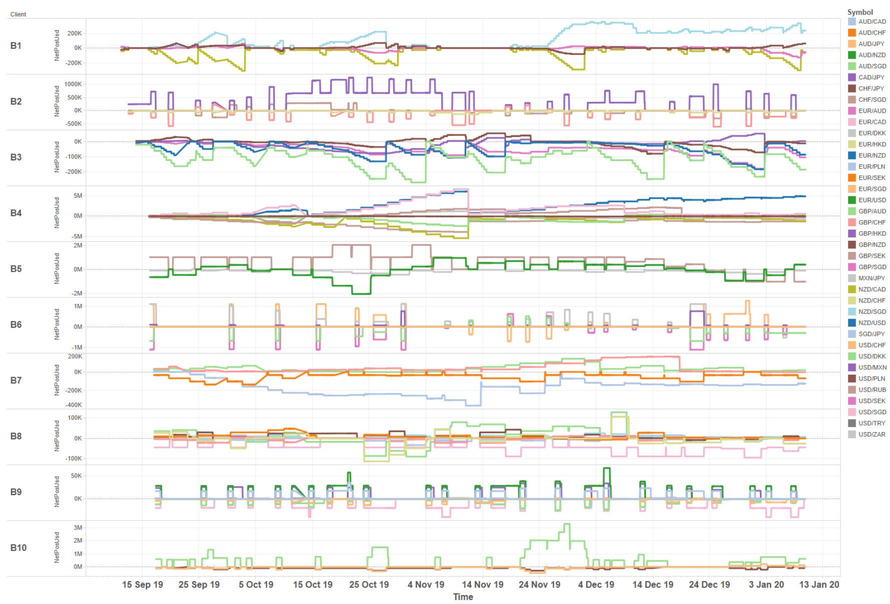

Fig. 1: Net positions of first 10 clients in data set from Sept 2019 to Jan 2020.

Until now we have assumed the existence of the portfolio vector γ, describing an expert's investment decisions. However, as we can see from the example trading data presented in Table 2, it is not clear how to define a clients prediction with that on the long-short game. We must therefore define a method of calculating an investor's portfolio vector from raw trade data that describes their investment decision over each time interval. The portfolio vector γ describes the sizes of investors' positions in relation to their wealth. This requires knowledge of the investors wealth however, as is common we do not have access to the total funds available to an investor. Instead, we can make the assumption that at each trial the investor has invested their total funds across each of the available assets. The DAPRA data set provides us with the normalised positions an investor held in each asset at the start of each interval where we will use $\text{Pos}_t^{\theta}[n]$ to denote the position of investor θ in asset n at time t.

Therefore, a natural method of approximation is to define the portfolio vector $\gamma_t^\theta \in \mathbb{R}^M$ as $\gamma_t^\theta[m] = \text{Pos}_t^\theta[m] / \sum_{m=0}^{M-1} |\text{Pos}_t^\theta[m]|$

The data set is available online [Al-baghdadi and Lindsay, 2020].

3.2 Empirical Results

Here we will compare the portfolio performance of an investor following the investment protocol of the long-short game applied to the 100 clients to giving each client a fixed equal weight. We will make the assumption all assets within are market are arbitrarily indivisible and no transactions costs are present. Naturally, as we are using historical market data it is implicit our trading behaviour has no effect on the market. A practical performance measure of a portfolio is the return on investment (ROI), assuming no transaction cost we can calculate from the wealth of the investor using $\text{ROI} = (W_T - W_0)/W_0 \times 100 = (W_T - 1) \times 100$. This is equal to $e^{-\text{Loss}_T(\text{AA}(\eta, P_0))} \times 100$ where $\text{Loss}_T(\text{AA}(\eta, P_0))$ is the total loss w.r.t the long-short game loss following the AA with learning rate η and initial distribution P_0.

After partitioning our data into 1 minute intervals the result is 123596 trials. In Fig. 2 we can see the number of awake experts at each trial. Here an awake expert is defined as one that has an open position in at least one symbol over a given trial, otherwise they are defined as sleeping. We can see it takes around 60,000 trials for the number of awake experts to stabilise. This is not an uncommon problem to face in real world in cases such as an emerging market maker or the introduction of a new investment instrument. In order to keep the comparison between the various implementations of the aggregating algorithm fair, we will be ignoring the initial 60,000 trials from our experiments, however, the full data set has been provided.

The percentage return on investment (ROI) to the portfolio of the long-short game is 0.6449% as we can see in Fig. 3 this is good growth over the investment period. However, comparing this to assigning all experts equal weights we see an ROI of 0.6447%, this is a minor change in the ROI. In Fig. 4 we plot the excess ROI to the long-short game compared to equal weights, whilst the difference is small we do see a clear indication the AA has the potential superior predictions than simply following each expert equally.

Whilst the long-short game appears to be an improvement over using equal weights the gains are small with the difference in ROI after 63,596 trials being 0.002%.

The resulting ROI close to 1% is disconcerting compared to the results of the experts. The returns of the experts range from -2.33% to 4.62% with the mean of 0.64%. Clearly, AA fails to align itself with the best experts.

In terms of Lemma 2.2 this may be explained as follows. The total losses of the experts range from -0.05 to 0.020. The extra term in (4) with $\eta = 1$ and uniform initial weights of $P_0(\theta) = 1/100$ equals 4.6052. The guarantees of Lemma 2.2 are thus very loose.

In Fig. 5 we can see the weight the AA assigns to each client in the long-short game, at each trial. As clients weights are updated using their loss

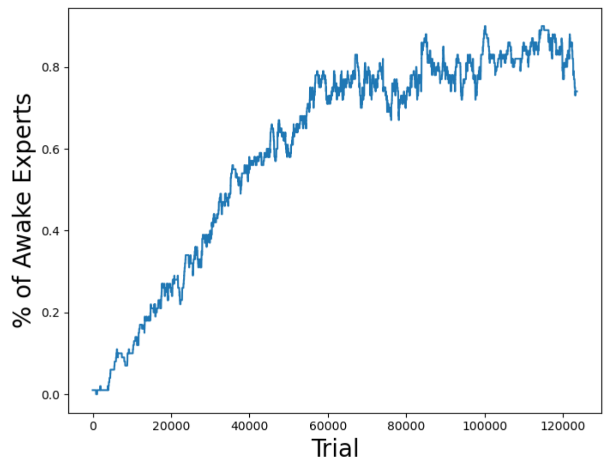

Fig. 2: Percentage of Awake Experts Through Time

at each trail, the weight is a reflection of the ROI. Therefore, showing their are clients in the pool with returns far greater that the achieved by the AA. We see the game does differentiate between the various strategies however, there appears to be insufficient discrimination of weights to allow above average strategies to influence the overall investment decisions of the AA. In this case the models final weighs have a mean of 1.01 with a standard deviation of 0.013 and with maximum and minimum weights of 1.04 and 0.98, this may explain for the limited performance improvements of the long-short game. Therefore, it seems logical the performance of the long-short game may be increased by modifications to the game that increase the discrimination between the weights of each investment strategy.

4 Long-Short Game Modifications

In this section we introduce several modifications in order to improve the practical performance.

The ideas developed here stem from the following intuition. While Cover's game has a natural scaling (the components of γ sum to 1), the long-short game does not. In the example we considered above, AA produces vectors γ_t that lead to very small but positive profit. One can multiply these vectors by a factor of $A > 1$ and the profit will increase. Where the investor earned the profit of $W_{t-1}\langle\omega_t, \gamma_t\rangle$, they will earn $W_{t-1}\langle\omega_t, A\gamma_t\rangle = AW_{t-1}\langle\omega_t, \gamma_t\rangle$.

The downside of this is risk. Larger positions can potentially lead to bankruptcy, which causes problems as explained in Section 2.4. In the spirit of prediction with expert advice, we can analyse the possibility w.r.t. the bankruptcy of experts.

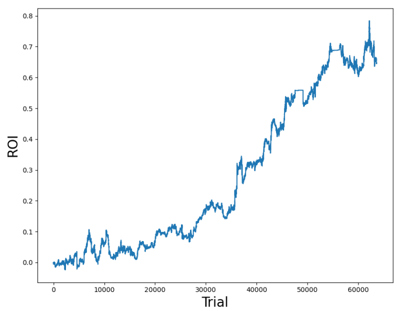

Fig. 3: Long Short Game ROI

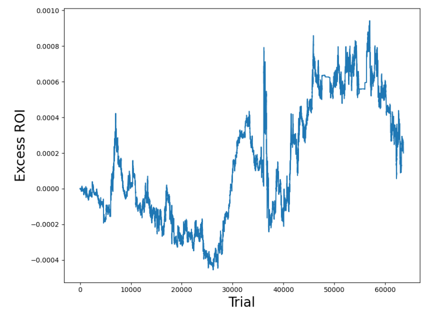

Fig. 4: Long Short Game Excess ROI

We will call a prediction γ satisfying (8) a-bounded.

Definition 1. *A method of merging experts' predictions is* a-conservative *if for all experts' predictions* $\gamma(\theta)$ *if* $\gamma(\theta)$ *is* a*-bounded for all* $\theta \in \theta$*, then the prediction* γ *produced by the method is* a*-bounded.*

Note that in practice being a-bounded is neither necessary nor sufficient for avoiding bankruptcy. An investor may take calculated risk and get away with it.

Definition 2. *A method of merging experts' predictions is* conservative *if for all experts' predictions* $\gamma(\theta)$ *the prediction* γ *produced by the method is such that for all* $\omega \in \Omega$ *if the values* $-\ln(1 + \langle \omega, \gamma(\theta) \rangle)$ *are uniformly bounded from above by a finite number, i.e.,* $-\ln(1+\langle \omega, \gamma(\theta) \rangle) \le C < +\infty$ *for all* $\theta \in \Theta$*, then the value* $-\ln(1+\langle \omega, \gamma \rangle)$ *is finite, i.e.,* $-\ln(1+\langle \omega, \gamma \rangle) < +\infty$*.*

Theorem 4.1. *Any merging algorithm outputting an average of experts' predictions w.r.t. some distribution, i.e.,* $\gamma = \int_\Theta \gamma(\theta) P(d\theta)$*, where* P *is some distribution, is conservative and* a*-conservative for all* $a > 0$*.*

Proof. If $-\ln(1 + \langle \omega, \gamma(\theta) \rangle \le C < +\infty$ for all $\theta \in \Theta$, then $1 + \langle \omega, \gamma(\theta) \rangle \ge 2^{-C} > 0$ and

$$1 + \langle \omega, \gamma \rangle = 1 + \left\langle \omega, \int_\Theta \gamma(\theta) P(d\theta) \right\rangle = \int_\Theta (1 + \langle \omega, \gamma(\theta) \rangle) P(d\theta) \ge \int_\Theta 2^{-C} P(d\theta) > 0$$

if $\|\gamma(\theta)\|_1 \le a$, then

$$\left\| \int_\Theta \gamma(\theta) P(d\theta) \right\|_1 = \sum_{m=0}^{M-1} \left| \int_\Theta \gamma(\theta)[m] P(d\theta) \right|$$

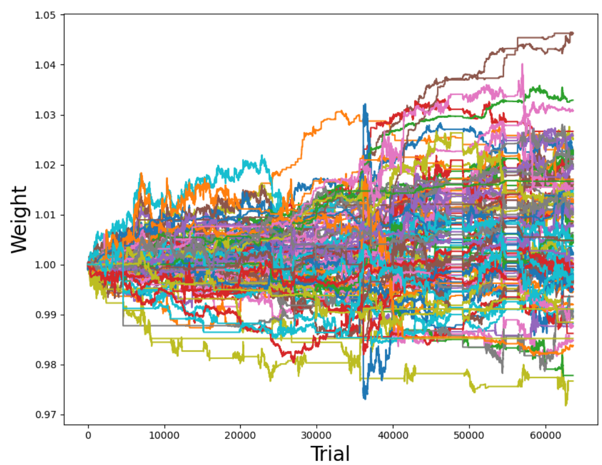

Fig. 5: Long-Short Game Client Weights

$$\leq \sum_{m=0}^{M-1} \int_{\Theta} |\gamma_t(\theta)[m]| P(d\theta) \leq \int_{\Theta} aP(d\theta) = a$$

□

We can see that whenever the experts' predictions satisfy a broker's safety requirements, so do the AA predictions and whenever the experts' predictions do not lead to bankruptcy, neither does the AA.

In this section we will introduce modification of the long-short game keeping this property w.r.t. the original λ_{LS} and improving on the practical performance of the AA.

4.1 Return Scaling

Take a number $\rho > 0$ and consider the loss function

$$\lambda_{\mathrm{LS},\rho} = \begin{cases} -\ln(1+\rho\langle\omega,\gamma\rangle) & \text{if } 1+\rho\langle\omega,\gamma\rangle > 0 \\ +\infty \text{ otherwise} \end{cases} .$$

One can define the a-bounded and general versions of the game with the same Γ and Ω as the original long-short game.

Theorem 4.2. *For any $\rho > 0$, for any a-bounded game, $a > 0$, and for every $\eta \leq 1$ we have $C_\eta = 1$. The only prediction attaining $C(g) = 1$ is the average (6), where as before P is a probability distribution in Γ generating g. When $\eta > 1$, $C_\eta > 1$.*

The proof is the same as for Lemma 2.5, which is proven by Vovk and Watkins [1998]. For the general game with unrestricted sets of predictions and outcomes the analysis of Section 2.4 applies.

We will apply the AA in the following fashion. We will use $\lambda_{\mathrm{LS},\rho}$ in the algorithm for calculating weights and working out the predictions γ_t. Then we will evaluate w.r.t. the original λ_{LS}.

Of course, the λ_{LS}-loss of the resulting algorithm will not satisfy Lemma 2.2. However, Lemma 2.2 will mostly hold for the loss $\lambda_{\mathrm{LS},\rho}$. Note that the $\lambda_{\mathrm{LS},\rho}$-loss of a strategy is the same as the loss of the strategy with all predictions multiplied by ρ. This strategy suffers larger loss and the term $\frac{1}{\eta}\ln\frac{1}{P_0(\theta)}$ will be small in comparison. Thus the algorithm will allow better differentiation of the weights, which *may* result in a better λ_{LS}-loss.

As discussed above, there is danger that the strategy with predictions multiplied by ρ goes bankrupt. However this will not necessarily propagate to the original unscaled loss and the mixture.

Corollary 4.2.1. *The aggregating algorithm applied w.r.t. the loss $\lambda_{\mathrm{LS},\rho}$ is conservative and a-conservative for every $a > 0$.*

Proof. We still average experts predictions with some weights. While the weights may be different to AA, the argument of Theorem 4.1 stays. □

We may be affected by bankruptcy in two way. If a ρ-multiple of an original expert goes bankrupt, we get an extra term in (12) worsening the performance. On the other hand, the future predictions of the strategy disappear from the mixture and the losses will not appear in the comparison. Still we do not have to go bankrupt as per Corollary 4.2.1.

4.2 Downside Loss

The developments of this section are based on the following intuition.

From the practical perspective, the ability of a strategy not to lose money may be more important than the ability to earn money. Consider a strategy that earns little money, but does so very consistently and never looses much. This strategy can then be scaled up and earn more money.

Thus one often wants to minimise the drawdown of a trading strategy. There are various indicators quantifying it; they are discussed in the next section. One cannot apply AA directly to this problem because the notion of a drawdown is not local in time. Still one can try and modify the loss function to penalise financial losses stronger.

Consider the downside loss function modifying the scaled long-short loss:

$$\lambda_{\mathrm{LS,down},\rho}(\omega,\gamma) = \max(-\ln(1+\rho\langle\omega,\gamma\rangle),0) = -\ln(1+\rho\min(\langle\omega,\gamma\rangle,0))$$

This function penalises financial losses but does not reward gains.

The following statement can be made about its mixability properties.

Theorem 4.3. *For $\lambda_{\mathrm{LS},\rho}$ the average (6) attains $C = 1$ for every g, where as before P is a probability distribution in Γ generating g.*

Proof. Consider a distribution P on Θ, and predictions $\gamma(\theta)$. One has

$$e^{-\lambda_{\mathrm{LS,down},\rho}(\omega,\gamma)} = 1+\rho\min(\langle\omega,\gamma\rangle,0)$$

and therefore it is sufficient to prove that

$$1 + \rho \min(\langle \omega, \int_\Theta \gamma(\theta) P(d\theta) \rangle, 0) \geq \int_\Theta (1 + \rho \min(\langle \omega, \gamma(\theta) \rangle, 0) P(d\theta) \ .$$

This follows from the concavity of $\min(x, 0)$ in x and Jensen's inequality.
□

It is important to point out that this loss function is really special. There is $\gamma_0 = 0$ such that $0 = \lambda_{\mathrm{LS,down},\rho}(\omega, \gamma_0) \leq \lambda_{\mathrm{LS,down},\rho}(\omega, \gamma)$ for any ω and any other γ. Technically $C = 0$ and the problem of prediction with expert advice is trivial for this loss function: the learner only needs to predict 0.

Still applying the AA with $\lambda_{\mathrm{LS,down},\rho}$ and the substitution (6) in meaningful and the losses will satisfy Lemma 2.2.

4.3 Combined Loss Function

One can consider the combined loss function parameterised by scalings $\rho_1 \geq 0$ and $\rho_2 \geq 0$ and coefficients $u \geq 0$ and $v \geq 0$ (we assume that $\rho_1 + \rho_2 > 0$ and $u + v > 0$):

$$\begin{aligned} \lambda(\omega, \gamma) &= -\ln(u e^{-\lambda_{\mathrm{LS},\rho}(\omega,\gamma)} + v e^{-\lambda_{\mathrm{LS,down},\rho}(\omega,\gamma)}) \\ &= -\ln((u+v) + u\rho_1 \langle \omega, \gamma \rangle + v\rho_2 \min(\langle \omega, \gamma \rangle, 0)) \\ &= -\ln(u+v) - \ln\left(1 + \frac{u\rho_1}{u+v} \langle \omega, \gamma \rangle + \frac{v\rho_2}{u+v} \min(\langle \omega, \gamma \rangle, 0)\right) \end{aligned}$$

Since this loss function is mixable (as we will see in a moment) the additive term $-\ln(u+v)$ makes no difference and can be ignored. One may think of the combined function as having only two parameters, $u\rho_1/(u+v)$ and $v\rho_2/(u+v)$, but speaking of four parameters may be more convenient.

Theorem 4.4. *For any $\rho_1, \rho_2, u, v \geq 0$ such that that $\rho_1 + \rho_2 > 0$ and $u+v > 0$, for any a-bounded game, $a > 0$, and general game with combined loss $\lambda_{\mathrm{LS},\rho}$ and for $\eta = 1$ we have $C_\eta = 1$. This is attained by the average substitution function (6), where as before P is a probability distribution in Γ generating g.*

Corollary 4.4.1. *For any $\rho_1, \rho_2, u, v \geq 0$ such that $\rho_1 + \rho_2 > 0$ and $u + v > 0$, the aggregating algorithm with the combined loss function in conservative and a-conservative for every $a > 0$.*

5 Experiments

Here we evaluate the performance of the proposed modifications to the long-short game and discuss the advantages of the various investment strategies. We will also be comparing the effect of sleeping experts using the same parameters.

5.1 Portfolio Performance Evaluation

To evaluate the performance we will be using well established portfolio risk measures. As the games we are studying use returns over each trial to update the weight assigned to each expert investor we will naturally use ROI of each learner's portfolio as a measure of success. However, it is typical to not only evaluate a portfolio base on return alone but rather the risk-reward of the portfolio. The Sharpe ratio [Sharpe, 1966] of a portfolio P is a measure of the amount of return an investor receives per unit of risk defined as:

$$\text{Sharpe}(P) = \frac{R_P - R_f}{\sigma(R_P)} \ , \tag{15}$$

where R_p denotes the return to the portfolio p and R_f the return of the risk-free asset. This allows us to compare the risk of each of the learners portfolios, using the standard deviation of the returns to the portfolio as a measure of volatility.

As we discuses in Section 4.2, one may be specifically interested in reducing the financial losses. The Sortino ratio [Sortino and Price, 1994] is a measure of return per unit of downside risk defined as

$$\text{Sortino}(p) = \frac{R_P - R_f}{\sigma(R_d)} \ , \tag{16}$$

where R_d denotes the downside returns to the portfolio P being the returns recorded less than some target return. In the following we will assume the a return to the risk-free asset of 0% and a target return of 0%, for the propose of performance comparison.

We will also be looking at the *Drawdown* of investors portfolios, is a metric commonly used to measure the volatility of the ROI and refers to how much the ROI retraces from the highest ROI achieved, defined as:

$$\text{Drawdown}(T) = \min_{t \in [0,T]} (0, \text{ROI(T)} - \text{ROI}(t))$$

we ideally want to keep drawdown as small as possible, a useful summary is to track the so called 'maximum drawdown' which computes the maximum amount of ROI given away over time:

$$\text{Max Drawdown}(T) = \min_{t \in [0,T]} \text{Drawdown}(t)$$

5.2 Empirical Results

In Fig. 6 we can see the ROI of the learners portfolio using various combined loss coefficients and setting $\rho_1 = \rho_2$ increasing from a value of 0-1000 (sampling every 100 intervals). The first thing to note is that for all games there is an improvement in the ROI up until some point where we begin to see a drop off in the results.

The learner following the investment strategy using a loss function of $\rho = 900, u = 1$ and $v = 0$ has an ROI of 3.83%, which is a 5.9 fold increase compared to the ROI of 0.65% resuting from the portfolio with a loss function with parameters of $\rho = 1, u = 1$ and $v = 0$. However, whilst a combined loss of u =1 and v = 0 produces the highest ROI and sharpe ratio, its sortino ratio is much lower when compared to games applying downside loss. This suggested that whilst games without downside loss can provide high ROI per unit of volatility we can expect to see much higher drawdowns in these portfolios when compared to those using downside loss. We can see an example of this represented in Fig. 14 comparing the portfolios of two investment strategies in this case with the implementation of specialist experts. We observe that while the combined loss coefficients of $u = 1$ and $v = 0$ produces a final ROI higher than the strategy using coefficients $u = 2$ and $v = 1$, the drawdowns are far greater seen in Fig. 15.

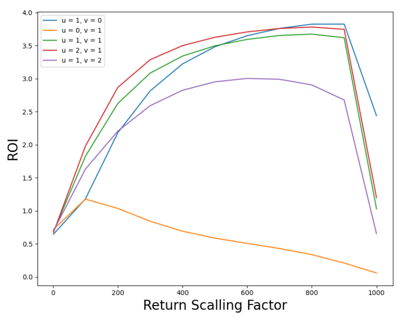

Fig. 6: ROI for increasing ρ

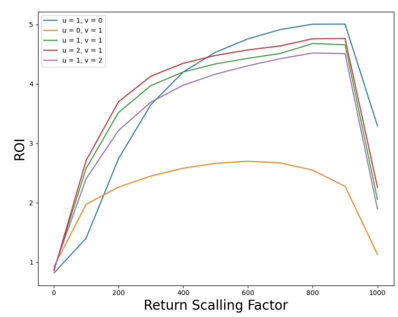

Fig. 7: ROI with Sleeping Experts

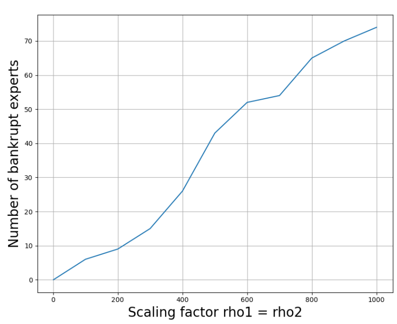

Fig. 8: Bankrupt clients as ρ increases

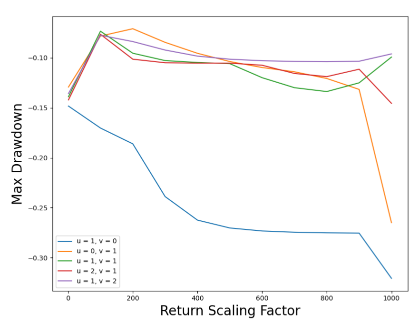

Fig. 9: Max Drawdown as ρ increases

We see that the performance of the algorithms at first improves with the growth of ρ and then starts falling. According to Section 3 and 4 the improvement in performance is caused by better differentiation in the experts' weights and increase in the significance of 4. This is offset by the growth of the number of bankrupt experts. In Fig. 8 we plot the number of clients with bankrupt trading strategies as the return scaling constant applied to the experts loss increases. We see the number of bankruptcies steadily grows as $\rho \to \infty$ as we would expect as larger losses force client weights to zero. This may account for the sharp drop in portfolio performance as $\rho_1 = \rho_2$ reach 1000. This is a clear representation of why we must increase the return scaling constant with caution as too large a value will grantee a portfolio less than optimal performance.

In Fig. 9 we can see the maximum drawdown for games without the use of specialist experts or discounting. Firstly, we see clearly that in the case of $u = 1$ and $v = 0$ as we increase the return scaling constant we increase the maximum drawdown of the portfolio. This must be look at along side the Sharpe ratio that is increasing. This suggest that while we increase the return per unit of risk of the learners portfolio we also increase the volatility of the portfolio. However, we can see for games where $v > 1$ the correlation dose not hold. This implies that for investors looking to reduce drawdowns they can achieve this using combined loss coefficients where $u \leq v$,

Comparing Fig. 6 with Fig. 7 we can see that while the pattern of the results largely remains the same, there is a significant increase in the ROI of the portfolios following the sleeping experts decision making method. Taking the parameters that result in the highest ROI for non sleeping experts of $\rho = 900, u = 1$ and $v = 0$ and a ROI of 3.83%, the application of sleeping experts increases this to an ROI of 5% a 1.31 fold increase. We also see a similar increases in the sharpe ratios of learners portfolios comparing the same two portofolios we see an increase of 20% from 0.52% to 0.62%. However, whilst for the same game we saw an increase in the sortino ratio of 0.08%, overall sortino ratios decreased. The optimal sortino ratio without the use of sleeping experts was achieved using parameters of $\rho = 700, u = 2$ and $v = 1$ of 7.48. With the use of sleeping experts the highest sortino ratio was given using $\rho = 300, u = 2$ and $v = 1$ of 6.09, a decrease of 18%.

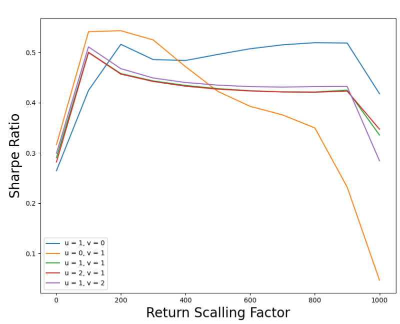

Fig. 10: Sharpe Ratio for increasing ρ

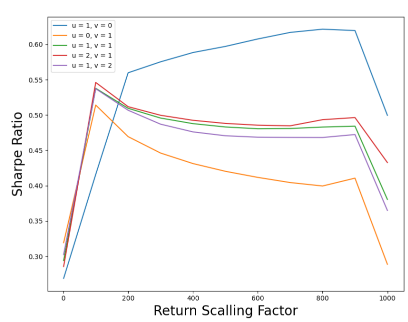

Fig. 11: Sharpe Ratio with Sleeping Experts

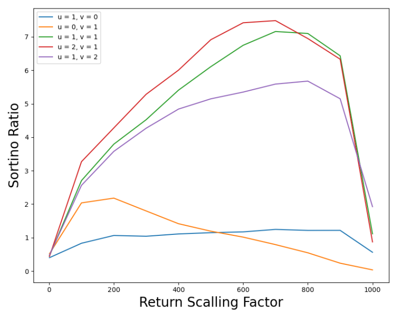

Fig. 12: Sortino Ratio for increasing ρ

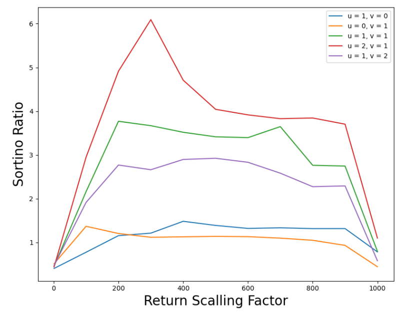

Fig. 13: Sortino Ratio with Sleeping Experts

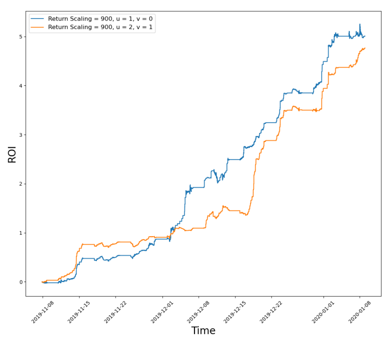

Fig. 14: ROI of Learner Portfolio

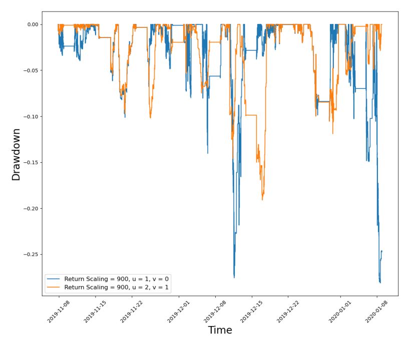

Fig. 15: Drawdown of Learner Portfolio

Here we presented the results of discounting taking at $\alpha = 0.995$. Examining the ROI of the learner applying discounting whilst we see an increase correlated to the return scaling constant, the returns are lower and the growth is not as significant as other implementations of the AA. However, we do see a notable change in the pattern of ROI taking the combined loss coefficients $u = 0$ and $v = 1$, with stronger performance than without the use of discounting. Moreover, we see a notable improvement in the sortino ratio that one may assume to be a result of an increase in the effectiveness of downside loss. In Fig19 we can see the ROI with the parameters $u = 1$, $v = 1$, $\rho = 1000$ and $\alpha = 0.995$, this is a noteworthy example as we can see whilst the ROI is less than previous games the learner is still actively investing in the market with minimal drawdowns. This suggested that if an investor prioritises reducing drawdowns in their portfolio at the expense of long-term ROI the use of discounting may be an optimal strategy for a learner.

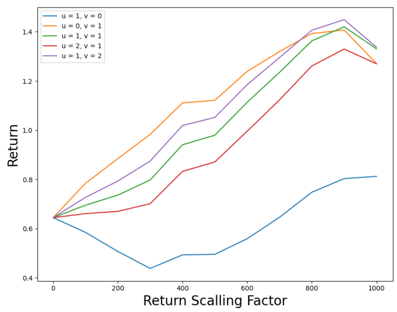

Fig. 16: ROI with Discounting

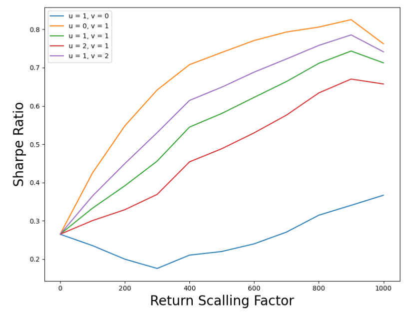

Fig. 17: Sharpe Ratio with Discounting

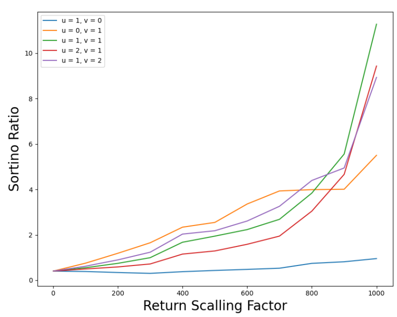

Fig. 18: Sortino Ratio with Discounting

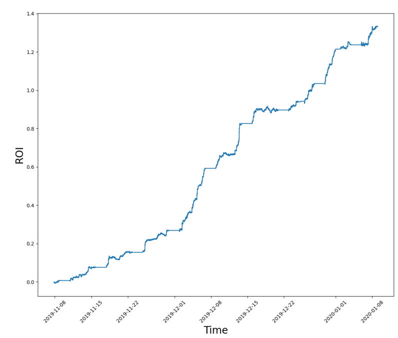

Fig. 19: ROI of Learner with Discounting

In Table 3 we can see the parameters and results of some the AA implementations with optimal results. We see that for less risk averse investors looking for high returns, the use of specialist experts can achieve this. Where in the case of the highest ROI of any AA it is in fact higher than the ROI of the best expert. However, for more risk averse investors the use of downside loss and discounting can certainly reduce the overall volatility of the portfolio. We must also consider the use of specialist experts in combination with downside loss as this has shown to be a good compromise for investors.

Table 3: Parameters and results of optimal solutions

	u	v	Return Scalling Factor	Specialist Experts	Alpha	ROI	Sharpe	Sortino	Max Drawdown
Equal Weights	-	-	-	-	-	0.64%	0.27	0.40	-0.15
Worst Expert	-	-	-	-	-	-2.34%	-0.23	-0.15	-3.04
Best Expert	-	-	-	-	-	4.52%	0.42	4.03	-0.22
Highest ROI	1	0	900	YES	1	**5%**	0.62	1.31	-0.28
Highest Sharpe and Best Max Drawdown	0	1	900	NO	0.995	1.4%	**0.83**	4.01	**-0.02**
Highest Sortino	1	1	1000	NO	0.995	1.33%	0.71	**11.28**	-0.035

6 Conclusion

We have shown the practical limitations of the long-short game and have introduced modifications with clear performance benefits in our experimental results.

This study has presented a novel time series data set that describes client trades in the Foreign Exchange market (Al-baghdadi and Lindsay [2020]), and has used this data to introduce a method of deriving expert predictions from client positions. Return scaling of the long-short game has been introduced, aimed to address the practical issue of insufficient discrimination of expert weights and has been shown to provide significant performance improvements. We have also presented the downside long-short game with the motivation of reducing the downside risk of the investment strategy of the AA, which has proven to be effective using the Sortino ratio as a measure of downside risk. We have used combined loss functions to produce optimal performance of AA portfolios both maximising returns and reducing risk. Finally, we have shown the use of specialist experts can improve the overall performance of a learners portfolio. Whereas, discounting can significantly reduce the volatility of an investment portfolio when used with downside loss at the expense of ROI.

In this paper we have shown that by adjusting the relationship between return scaling and downside loss, an investor can improve both their overall profitability as well as reduce their drawdown. These contributions made to the Long-Short game give investors the ability to vary their decision making process of merging trading strategies, based on their risk/reward appetite.

Data Availability Statement

The datasets generated during and analysed during the current study are available in the Kaggle Aggregating Algorithm Long-Short Game Dataset repository, Al-baghdadi and Lindsay [2020]. In order to reproduce the results found in this paper, code for the Aggregating Algorithm and loss functions can be found on the Git repository Al-baghdadi [2022].

Conflicts of Interests

This study was funded by Algorithmic Laboratories Ltd and Royal Holloway, University of London. The authors have no competing interests to declare that are relevant to the content of this article.

Acknowledgments

We would like to give special thanks to the team at Algorithmic Laboratories Ltd (AlgoLabs) for the support and guidance on this project.

References

Najim Al-baghdadi. Aggregating algorithm git repository. https://github.com/Najim-Al/PhD_AA, 2022.

Najim Al-baghdadi and David Lindsay. Aggregating algorithm long-short game dataset. https://www.kaggle.com/najimal/aggregating-algorithm-longshort-game-dataset, 2020.

Najim Al-baghdadi, Wojciech Wisniewski, David Lindsay, Sian Lindsay, Yuri Kalnishkan, and Chris Watkins. Structuring time series data to gain insight into agent behaviour. In *Proceedings of the 3rd International Workshop on Big Data for Financial News and Data*. IEEE, 2019.

A. Chernov and V. Vovk. Prediction with expert evaluators' advice. In *Algorithmic Learning Theory, ALT 2009, Proceedings*, volume 5809 of *LNCS*, pages 8–22. Springer, 2009.

A. Chernov and F. Zhdanov. Prediction with expert advice under discounted loss. In *Proccedings of ALT 2010*, volume LNAI 6331, pages 255–269. Springer, 2010. See also arXiv:1005.1918 [cs.LG].

T. M. Cover and E. Ordentlich. Universal portfolios with side information. *IEEE Transactions on Information Theory*, 42(2):348–363, 1996.

László Györfi, Frederic Udina, and Harro Walk. Experiments on universal portfolio selection usingdata from real markets. https://api.semanticscholar.org/CorpusID:3133273, 2008.

David P. Helmbold, Robert E. Schapire, Yoram Singer, and Manfred K. Warmuth. On-line portfolio selection using multiplicative updates. *Mathematical Finance*, 8(4):325–347, 1998.

Yuri Kalnishkan. Prediction with expert advice for a finite number of experts: A practical introduction. *Pattern Recognition*, 126, June 2022. ISSN 0031-3203. doi: 10.1016/j.patcog.2022.108557.

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

Harry Markowitz. Portfolio selection. *The Journal of Finance*, 7(1):77–91, 1952.

William F Sharpe. Mutual fund performance. *The Journal of business*, 39(1):119–138, 1966.

Frank A Sortino and Lee N Price. Performance measurement in a downside risk framework. *the Journal of Investing*, 3(3):59–64, 1994.

Vladimir Vovk. Aggregating strategies. In M. Fulk and John Case, editors, *Proceedings of the Third Annual Workshop on Computational Learning Theory*, pages 371–383. Morgan Kaufmann, 1990.

Vladimir Vovk. A game of prediction with expert advice. *Journal of Computer and System Sciences*, 56(2):153–173, 1998.

Volodya Vovk and Chris Watkins. Universal portfolio selection. In *Proceedings of the eleventh annual conference on Computational learning theory*, pages 12–23. ACM, 1998.

V V'yugin and V Trunov. Universal algorithmic trading. *Journal of Investment Strategies*, 2(1):13, 2012.

VV V'yugin, IA Stel'makh, and VG Trunov. Adaptive algorithm of tracking the best experts trajectory. *Journal of Communications Technology and Electronics*, 62(12):1434–1447, 2017.

Xingyu Yang, Hong Lin, Yong Zhang, et al. Boosting exponential gradient strategy for online portfolio selection: An aggregating experts' advice method. *Computational Economics*, 55(1):231–251, 2020.

Yong Zhang and Xingyu Yang. Online portfolio selection strategy based on combining experts' advice. *Computational Economics*, 50(1):141–159, 2017.