Proceedings of Machine Learning Research 128:1–20, 2020 Conformal and Probabilistic Prediction and Applications

Practical Investment with the Long-Short Game

Najim Al-baghdadi najim.al-baghdadi.2019@live.rhul.ac.uk
Department of Computer Science, Royal Holloway, University of London, Egham, United Kingdom
Yuri Kalnishkan yuri.kalnishkan@rhul.ac.uk
Department of Computer Science, Royal Holloway, University of London Egham, United Kingdom and Laboratory of Advanced Combinatorics and Network Applications, Moscow Institute of Physics and Technology, Institutsky per., 9, Dolgoprudny, 41701, Russia
David Lindsay david@algolabs.com
Siân Lindsay sian@algolabs.com
AlgoLabs, Bracknell, United Kingdom

Editor: Alexander Gammerman, Vladimir Vovk, Zhiyuan Luo, Evgueni Smirnov and Giovanni Cherubin

Abstract

In this paper we apply the aggregating algorithm, an on-line prediction with expert advice algorithm, to real-world foreign exchange trading data with the aim of finding investment strategies with optimal returns. We consider the Long-Short game first introduced in Vovk and Watkins (1998) and it's implementation, including the derivation of expert predictions from model trading data. Furthermore, we propose modifications to improve the practical performance of the game with respect to well-known portfolio performance indicators.

Keywords: Prediction with Expert Advice, Online Learning, Aggregating Algorithm, Portfolio Optimisation, Long-Short Game, Foreign Exchange, Currency Trading

1. Introduction

Since modern portfolio theory was first introduced in Markowitz (1952) the problem of portfolio selection has become increasingly prominent. We approach this problem using the framework of on-line learning, where an investor makes investment decisions based on the observations of a pool of investors' portfolios.

A well-known formalisation of this is Cover's game (section 2.2), where an investor partitions the available money between the assets. In Cover and Ordentlich (1996), a universal investment strategy is constructed for Cover's game: it performs nearly as well as any constant rebalanced portfolio. This strategy is a special case of the more general aggregating algorithm, which is capable of combining a finite or infinite pool of on-line 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 first introduced by Littlestone et al. (1989). The aggregating algorithm was later developed in Vovk (1990) and Vovk (1998) to include the more general concept of a loss function and prediction and outcome spaces. The intuitive idea of the aggregating algorithm is that given a series of predictions from a pool of experts, a learner observing the loss of each experts' prediction over time can adjust their trust in each expert to make future predictions. The problem of portfolio selection is a natural special case of a prediction with expert advice problem.

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. In Vovk and Watkins (1998) a modification of Cover's game, namely, the Long-Short game (see Section 2.3 of this paper) was introduced. It admits long positions exceeding the trader's capital (within specified limits) and short positions. The major downside of this 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 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 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 of investment with expert advice, as discussed by Vovk and Watkins (1998).

2.1. Aggregating Algorithm

In this section we 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* $\lambda : \Omega \times \Gamma \to (-\infty, +\infty]$.

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) provides a learner with a strategy to make a prediction of some future outcome 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 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 $\mathrm{AA}(\eta, P_0)$.

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 Pseudo-Algorithm (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 $P^*_{t-1} = P_{t-1}(d\theta)/P_{t-1}(d\Theta)$. 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 1 *For any learning rate $\eta > 0$, initial distribution P_0, and $T = 1, 2, \ldots$ 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 action $\Sigma(g_t)$, 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:

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

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

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

The mixability constant C_η is defined as

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

We can then find a substitution function Σ mapping generalised predictions g to Γ satisfying:

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

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

$$\mathrm{Loss}_T(\mathrm{AA}(\eta, P_0) \leq -\frac{C_\eta}{\eta} \ln \int_\Theta e^{-\eta \mathrm{Loss}_T(\theta)} P_0(d\theta) \ .$$

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.

Lemma 2 *For a finite pool of experts Θ,*

$$\mathrm{Loss}_T(\mathrm{AA}(\eta, P_0)) \leq C_\eta \mathrm{Loss}_T(\theta) + \frac{C_\eta}{\eta} \ln \frac{1}{P_0(\theta)} \ . \tag{2}$$

Moreover, if the game is η-mixable, then

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

Taking η such that $C_\eta = 1$ minimises the first term on the right-hand side of (2); this is important because this term may be growing with T. 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 (3).

Let us formulate the AA for the case of finitely many experts, which is the main case for this paper. Let N be the number of experts. We take $\Theta = \{1, 2, \ldots, n\}$.

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 predictions and distributions on them to predictions.

Parameters: Learning rate $\eta > 0$ and initial experts' weights $p_0(n)$, $n = 1, 2, \dots, N$
for $t = 1, 2, \dots$ **do**
read experts' predictions $\gamma_t(n) \in \Gamma$, $n = 1, 2, \dots, N$
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 $\gamma_t \in \Gamma$
update the weights $p_t(n) = p_{t-1}(n) e^{-\eta\lambda(\omega_t,\gamma_t)}$
end

Algorithm 1: On-line learning protocol

2.2. Cover's Game

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 $m = 0, 1, \dots, M-1$, 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

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

Cover's game is mixable.

Lemma 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*

$$\gamma^* := \int_\Gamma \gamma P(d\gamma) \ , \tag{4}$$

where P is a probability distribution in Γ generating g:

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

Lemma 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

$$\gamma_t = \sum_{n=1}^{N} p^*_{t-1}(n)\gamma_t(n) \tag{5}$$

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

$$\|\gamma\|_1 := |\gamma[0]| + ... + |\gamma[M-1]| \leq a \ ; \tag{6}$$

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 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 $0, 1, \ldots, M-1$.

It is more convenient to represent outcomes by a vector of returns here, so $\omega_t[m] = (S_t[m] - S_{t-1}[m])/S_{t-1}[m] = S_t[m]/S_{t-1}[m] - 1 \geq -1$. 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

$$\|\omega\|_\infty = \max_{m=0,1,\ldots,M-1} |\omega_m| \leq \frac{1}{a} \ . \tag{7}$$

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

The following lemma holds for every a-bounded game:

Lemma 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 (4), where as before P is a probability distribution in Γ generating g. When $\eta > 1$, $C_\eta > 1$.*

Substitution rule (4) (or (5) in the case of finitely many experts) should be used.

For practical applications it is inconvenient to use the a-bounded game because the condition (7) is not guaranteed. One would rather be talking of the general Long-Short game with $\Gamma = \mathbb{R}^m$, $\Omega = [-1, +\infty)^M$ and the loss function given by

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

However, this game does not have good mixability properties.

Lemma 6 *For every $\eta > 0$, there is no finite $C_\eta > 0$ such that (1) holds for the general Long-Short game.*

Proof We will construct a generalised prediction in the form

$$g(\omega) = -\frac{1}{\eta} \ln \sum_{n=1}^{N} p_n e^{-\eta \lambda_{\mathrm{LS}}(\omega, \gamma^n)} ,$$

where $N = M$ (the number of experts equals the number of assets), $\sum_n p_n = 1$ and $p_n > 0$ (one may take $p_n = 1/N$ for the sake of being definite), and $\gamma^n = -e_{n-1}$, i.e., $\gamma^n[n-1] = -1$ for $n = 1, 2, \ldots, N-1$ and $\gamma^n[m] = 0$ for $m \neq n-1$.

We need to show that no γ satisfies $\lambda_{\mathrm{LS}}(\omega, \gamma) \leq C g(\omega)$ for a finite $C > 0$ and all ω. First, take a γ such that $\gamma[n_0] < 0$ for some n_0. Let $\omega[n_0] = -2/\gamma[n_0]$ (or a larger positive number) and $\omega[n] = 0$ for $n \neq n_0$. Then $1 + \langle \omega, \gamma \rangle < 0$ and $\lambda_{\mathrm{LS}}(\omega, \gamma) = +\infty$. However, $\lambda_{\mathrm{LS}}(\omega, \gamma^n) = 0$ for $n \neq n_0$ so that $\sum_{n=1}^{N} p_n e^{-\eta \lambda_{\mathrm{LS}}(\omega, \gamma^n)} \geq 1 - p_{n_0} > 0$ and therefore $g(\omega) < +\infty$ is a finite number.

Now let $\gamma[n] \geq 0$ for all n. Let $\omega[n] = -1$ for all n. Then $1 + \langle \omega, \gamma \rangle \leq 1$ and $\lambda_{\mathrm{LS}}(\omega, \gamma) \geq 0$. But for every n we have $1 + \langle \omega, \gamma^n \rangle = 2$, so that $g(\omega) = -\frac{1}{\eta} \ln e^{-\eta(-\ln 2)} = -\ln 2 < 0$. ■

Still in practice one can apply the aggregating algorithm with $\eta = 1$, $C_\eta = 1$ and the substitution rule given by (5) to the general long-short game. If $1 + \langle \gamma_t, \omega_t \rangle > 0$ for $t = 1, 2, \ldots, T$, i.e., the learner does not get bankrupt along the way, the bounds of Lemma 2 will hold. This can be checked by retracing the proof.

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

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

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.

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 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 trade prices, allows one to normalise and sample the data as required for this study.

Fig. 1 compares the net positions of the first 10 clients in the data set, following DAPRA transformation over the trial period. 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

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

symbol composition and holding periods of the different clients, producing are varied rage of investment strategies.

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 $\mathrm{Pos}_t^\theta[n]$ 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 $\gamma_t^\theta[m] = \mathrm{Pos}_t^\theta[m] / \sum_{m=0}^{M-1} |\mathrm{Pos}_t^\theta[m]|$

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

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

of a portfolio is the return on investment (ROI), assuming no transaction cost we can calculate from the wealth of the investor using $\mathrm{ROI} = (W_T - W_0)/W_0 \times 100 = (W_T - 1) \times 100$. This is equal to $e^{-\mathrm{Loss}_T(\mathrm{AA}(\eta, P_0))} \times 100$ where $\mathrm{Loss}_T(\mathrm{AA}(\eta, P_0))$ 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. The percentage return on investment (ROI) to the portfolio of the Long-Short game is 1.0857% whereas the percentage ROI to the portfolio assigning equal weights to each model is 1.0825%. In Fig. 2 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 123596 trials being 0.003%.

The resulting ROI close to 1% is disconcerting compared to the results of the experts. The returns of the experts range from -2.0235% to 8.8201% with the mean of 1.0908%. Clearly, AA fails to align itself with the best experts.

In terms of Lemma 2 this may be explained as follows. The total losses of the experts range from -0.0845 to 0.0204. The extra term in (3) with $\eta = 1$ and uniform initial weights of $P_0(\theta) = 1/100$ equals 4.6052. The guarantees of Lemma 2 are thus very loose.

In Fig. 2 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 at each trail, the weight is a reflection of the ROI. Therefore, showing their are clients in the pool with returns

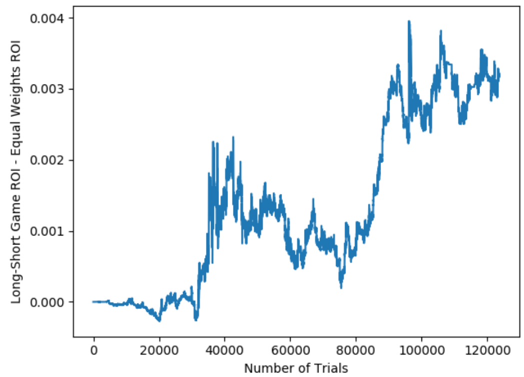

Figure 2: Long-Short Game ROI - Equal weights ROI

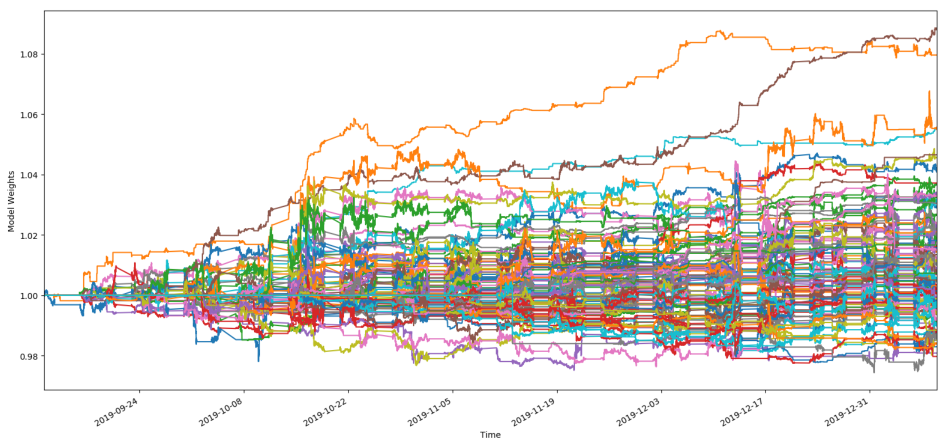

Figure 3: Long-Short Game ROI - Equal Weights ROI

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 0.02 and with maximum and minimum weights of 1.09 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. In the spirit of prediction with expert advice, we can analyse the possibility w.r.t. the bankruptcy of experts.

We will call a prediction γ satisfying (6) a-bounded.

Definition 7 *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 8 *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 9 *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

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

■

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 games.

Theorem 10 *For any $\rho > 0$, for any a-bounded game, $a > 0$, with loss $\lambda_{\mathrm{LS},\rho}$, and for every $\eta \leq 1$ we have $C_\eta = 1$. 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 (4), 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 5, which is proven by Vovk and Watkins (1998). The general unbounded game with the loss $\lambda_{\mathrm{LS},\rho}$ remains non-mixable as an analogue of Lemma 6 holds.

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. However, Lemma 2 will 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 propagate to the mixture.

Corollary 11 *The aggregating algorithm applied w.r.t. the loss $\lambda_{\mathrm{LS},\rho}$ and 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 9 stays. ■

We may be affected by bankruptcy in the following way. If a ρ-multiple of an original expert goes bankrupt, its weights in the algorithm drop to zero. Its future predictions disappear from the mixture and the losses do not appear in the comparison. Still we do not go bankrupt as per Corollary 11.

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 12 *For $\lambda_{\mathrm{LS},\rho}$ with a-bounded predictions and outcomes the average (4) 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 (4) in meaningful and the losses will satisfy Lemma 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 $u+v > 0$):

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

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 13 *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, $a > 0$, with combined loss $\lambda_{\mathrm{LS},\rho}$ and for $\eta = 1$ we have $C_\eta = 1$. This is attained by the average substitution function (4), where as before P is a probability distribution in Γ generating g.*

Corollary 14 *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

In this section we will provide empirical results to compare the performance of the Long-Short game to the modifications proposed in section 4 on the same data as used in section 3.

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:

$$\mathrm{Sharpe}(P) = \frac{R_P - R_f}{\sigma(R_P)} \quad , \tag{8}$$

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

$$\mathrm{Sortino}(p) = \frac{R_P - R_f}{\sigma(R_d)} \quad , \tag{9}$$

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.

5.2. Empirical Results

Here we evaluate the performance of the proposed modifications the Long-Short game and discuss the advantageous of the various investment strategies. As before we are using the data introduced in Section 3 and making the same set of assumptions.

Fig. 4 shows us the ROI of an investor following the Long-Short (LS) game and Downside Long-Short (DLS) game without combining loss functions, using a scaling constant ρ from

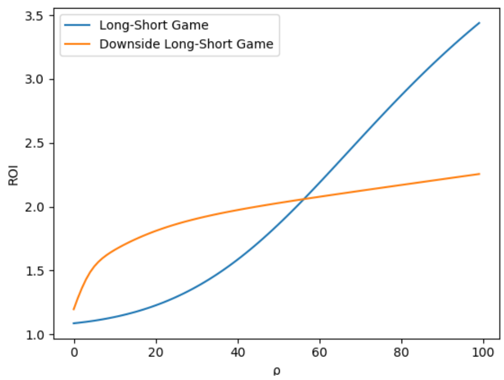

Figure 4: ROI for increasing ρ

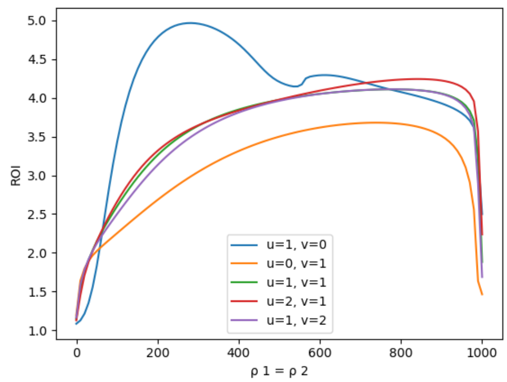

Figure 5: ROI for increasing $\rho_1 = \rho_2$

0 to 100. In the standard games, without the use of scaling we can see the DLS strategy outperforms the LS strategy with a return of 1.2% compared to 1.09%. The DLS portfolio continues to outperform the LS game, with the initial difference between the two games increasing, until a ρ value of 57, where the ROI of the two games meet. In Fig. 5, we extend this to combined loss functions setting $\rho_1 = \rho_2$ from 0 to 1000 (sampling every 10 steps). We see the ROI of the standard LS game continue to rise to a maximum of 4.96% at a scaling constant of 330 and this is the highest ROI of any AA strategy tested.

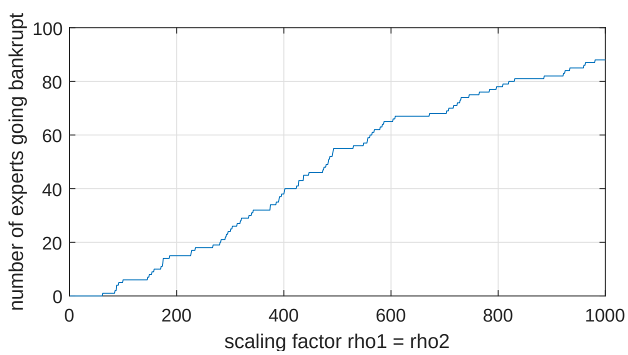

Figure 6: Number of bankrupt clients as return scaling constant increases.

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 (3). This is offset the growth of the number of bankrupt experts. In Fig. 6 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.

However, evaluating a portfolio solely based on ROI dose not provide us with complete understanding of the risk associated with the investment strategy. In Fig. 7 we have the daily Sharpe ratio of the standard games of ρ values up to 100. Again, we observe in the standard games the DLS produces a better result with a Sharpe ratio of 0.32 compared to 0.25, suggesting an investor following the DLS trading strategy takes less risk per unit of return. Both the LS and DLS games Sharpe ratio increased with an increasing ρ with maximum values of 0.41 and 0.52 respectively; however, both games converge to a ratio of around 0.4 as $\rho \to 100$. In Fig. 8 looking at the Sharpe ratio for combined loss we see a maximum Sharpe ratio of 0.55 at $\rho_1 = \rho_2 = 40$, where $u = 2$ and $v = 1$.

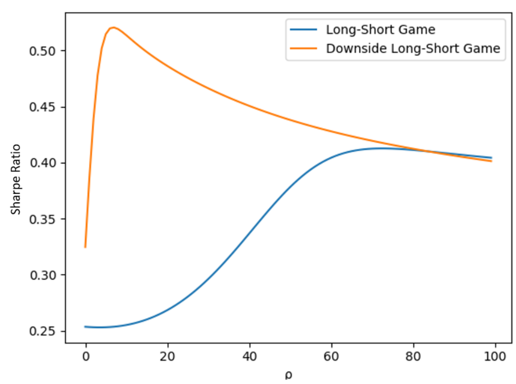

Figure 7: Sharpe ratio for increasing ρ

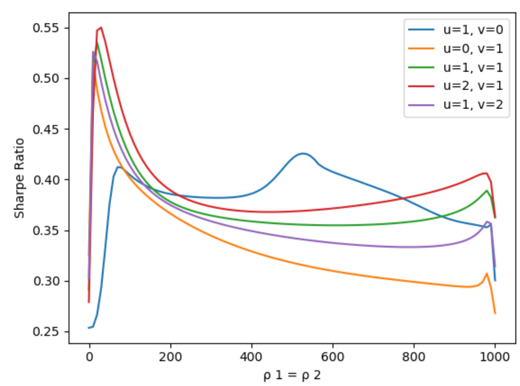

Figure 8: Sharpe ratio increasing $\rho_1 = \rho_2$

The picture of the daily Sortino ratio on the other hand is noticeably different between games, with the standard LS achieving a ratio of 0.45 and the standard DLS 0.61, suggesting an investor following the DLS investment strategy earns a higher reward per unit of downside risk. Here we see in games without combined loss the LS Sortino Ration fails to compete with the DLS Sortino ratio. However, we can clearly see the benefit of combined loss in Fig. 10 where all game with $v > 0$ far outperform those with $v = 0$.

In Fig. 11–14 we compare the ROI between the LS and DLS games (without combined loss) skipping some initial training time where the weight of each client are adjusted. In Fig. 11 we can see what appears to be a relatively steady growth over time suggesting after an initial training period the performance of a strategy following the DLS steadily improves upon that of the LS.

In Fig. 12 we see a similar pattern repeated with a ρ of 30 but with the difference in ROI being less. However, again referring to Fig. 7 and Fig. 9 we see that the Sharpe and Sortino ratio of the DLS portfolio far exceed that of the LS suggesting that whilst the ROI is converging the DLS portfolio is preferable due to the fact the investor receives a greater return per unit of risk. In Fig. 13 and Fig. 14 we see the pattern change with the ROI of the LS game showing periods of greater ROI than the DLS. However, it is noteworthy the increased ROI does not appear to be stable growth; this is reflected in the Sharpe and

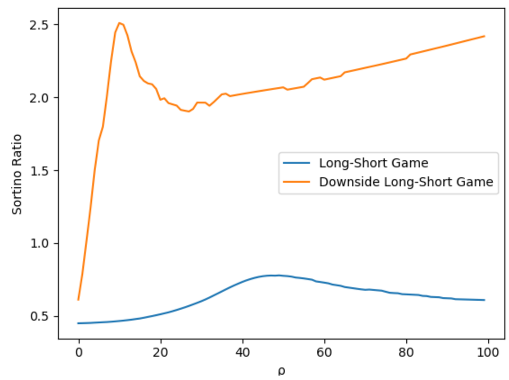

Figure 9: Sortino ratio for increasing ρ

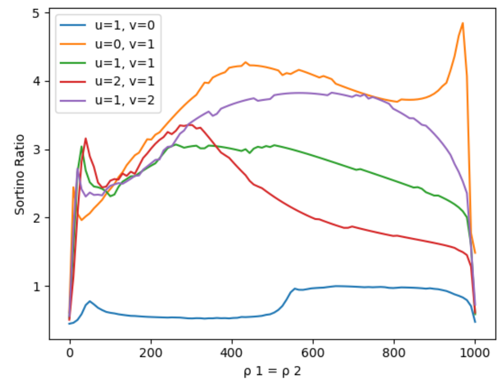

Figure 10: Sortino ratio increasing $\rho_1 = \rho_2$

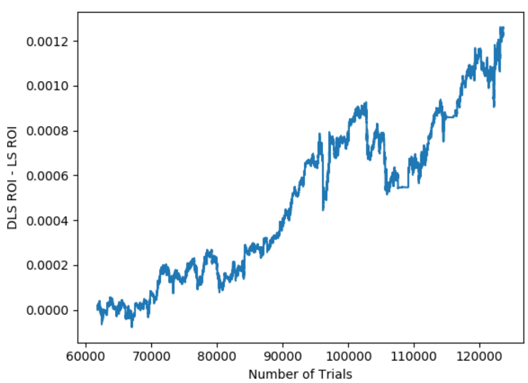

Figure 11: DLS ROI - LS ROI, $\rho = 1$

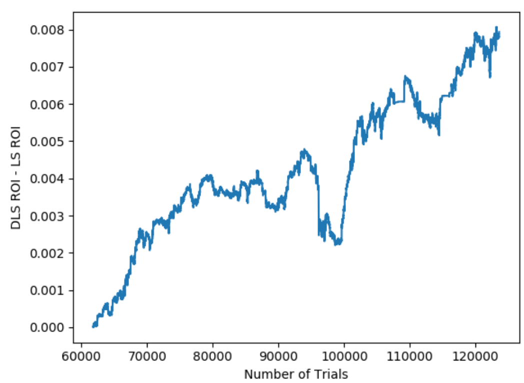

Figure 12: DLS ROI - LS ROI, $\rho = 30$

Sortino ratios for increasing scaling constants failing to outperform the DLS portfolios in the games without combined loss.

It is clear from these results there is no one set of parameters that is clearly optimal, but that the introduction of the DLS game, return scaling and combined loss allows for investment strategies that outperform the LS game. We observe that investor looking to maximise returns will favour aggregating strategies where $u > v$ and returns are scaled far above their true value. Whereas, for more risk averse investors return scaling is still effective however they may look to AA with higher v values especially where looking to reduce downside risk.

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.

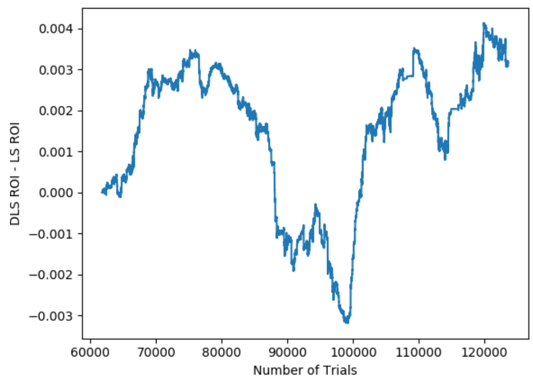

Figure 13: DLS ROI - LS ROI, $\rho = 60$

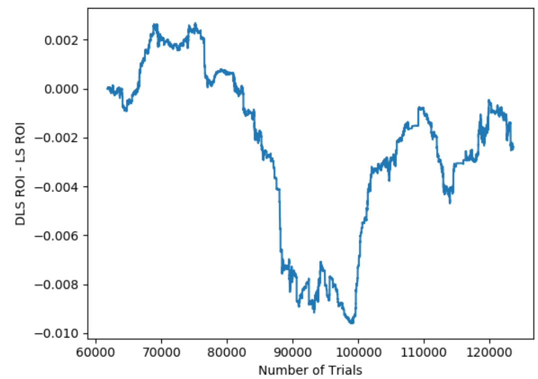

Figure 14: DLS ROI - LS ROI, $\rho = 90$

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. Finally, we have used combined loss functions to produce optimal performance of AA portfolios both maximising returns and reducing risk.

Acknowledgments

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

References

Najim Al-baghdadi and David Lindsay. Aggregating algorithm long-short game dataset. 5 2020. URL https://www.kaggle.com/najimal/aggregating-algorithm-longshort-game-dataset.

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, 10 2019.

Thomas M Cover and Erik 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. 2008. URL https://api.semanticscholar.org/CorpusID:3133273.

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.

Nick Littlestone, Manfred K Warmuth, et al. *The weighted majority algorithm*. University of California, Santa Cruz, Computer Research Laboratory, 1989.

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.