

Online Portfolio Hedging with the Weak Aggregating Algorithm

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

David Lindsay

DAVID@ALGOLABS.COM

Siân Lindsay

SIAN@ALGOLABS.COM

AlgoLabs, Bracknell, United Kingdom

Editor: Ulf Johansson, Henrik Boström, Khuong An Nguyen, Zhiyuan Luo and Lars Carlsson

Abstract

In this paper we apply the Weak Aggregating Algorithm to find optimal risk management strategies for financial Market Makers (MMs). Here risk is caused by the market exposure. It is effectively represented by the MM's overall net *position*, which is the aggregation of all the *buy* and *sell* trades carried out by the MM's clients at a given point in time. So-called *hedging* strategies are used by MMs to manage their risk and reduce market exposure. In essence, the MM actively places trades in order to reduce its overall net position, keeping it within some predefined bounds and as neutral (or flat) as possible. A flatter net position allows the MM to counter any unfavourable price movements which could otherwise incur a significant loss. We apply the Weak Aggregating Algorithm (WAA) to hedging strategies, which are treated as the experts. We combine their hedging decisions with the goal of reducing portfolio risk and maximising profitability, whilst also attempting to smooth out significant drawdowns. We develop a variation of the WAA using discounting and evaluate the WAA on a subset of real life client risk data in three commonly traded Foreign Exchange (FX) currency symbols: EUR/USD, EUR/GBP and GBP/USD. The results show how varying loss parameters and application of discount factors can enable the WAA to give combinations of hedging strategies that can significantly improve profitability and reduce drawdowns as compared to the benchmark of not hedging.

Keywords: Prediction with Expert Advice, Online Learning, Weak Aggregating Algorithm, Foreign Exchange, Currency Trading, Risk Management, Hedging

1. Introduction

Financial Market Makers (MMs) face a challenging online learning problem when it comes to managing a portfolio of risk. Here we define risk as it relates to market exposure - effectively represented by a MM's overall *position*. In simple terms, position is the aggregation of all the *buy* and *sell* trades carried out by the MM's clients at a given point in time and evolves due to underlying asset price fluctuations and changes in client trading activity. Profit and Loss (also known as *PnL*) is a function of position and asset price movement. Position is described as being *long* when the summation of client sell trades exceeds those of client buy trades. A *short* position indicates a higher summation of client buy trades over sell trades. Roughly speaking, when a position is long and the price goes up, the PnL will increase,

and vice versa when a position is short and the price goes down, the PnL will decrease. *Drawdown* is a metric commonly used to measure the volatility of the PnL and refers to how much the PnL retraces from the highest PnL achieved, defined as:

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

MMs would ideally like to keep this drawdown time series as small as possible, a useful summary is to track the so called ‘maximum drawdown’ which computes the maximum amount of PnL given away over time:

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

An effective risk management strategy is one that ensures that position is kept within some predefined bounds and is as flat as possible. If a position is allowed to build up with no limit, becoming either too long or too short, then any unfavourable asset price movement will result in the MM incurring greater losses than if its position had been flatter or neutral. Position is actively maintained (or *hedged*) when the MM places buy or sell trades (*hedges*) when its position is respectively long or short. Strategies which indicate how much is hedged and when hedge trades are to be placed are known as *hedging strategies*. It follows that a hedging strategy that causes a reduction in position will also impact PnL, yet the nature of this impact will vary based on when and how much is being hedged. Hedging too much will more likely reduce drawdown but also reduce any profit. Hedge too little and risk market exposure and being at the mercy of disadvantageous price movements leading to large drawdowns. We discuss a commonly used hedging strategy known as the Cylinder Hedging Model in Section 2.

In this paper we focus on finding optimal hedging strategies by making use of on-line prediction with expert advice, namely the Weak Aggregating Algorithm or WAA (Cesa-Bianchi and Lugosi, 2006; Kalnishkan and Vyugin, 2008).

The problem of finding an optimal hedging strategy follows naturally from that of portfolio selection, which has itself seen extensive application of both the WAA and the Aggregating Algorithm (AA). The problem was first introduced by Cover and Ordentlich (1996) and later developed by Vovk and Watkins (1998) to consider more realistic trading scenarios. In Al-Baghdadi et al. (2020) the game was further developed to improve the application of the AA to find profitable trading strategies based on the past observations of a pool of investment strategies. The WAA has also been applied to the problem of finding a universal portfolio by Zhang and Yang (2017) and further developed by He and Yang (2020) and Yang et al. (2020).

As well as presenting a hedging framework for the WAA, we also introduce a method for applying discounted loss to the WAA. This allows for a learner to more effectively adapt to changes in market conditions in periods of high volatility. In order to test the efficacy of the approaches presented here we conduct experimental trials on real world market and MM client data on three major currency pairs, EUR/USD, GBP/USD and EUR/GBP.

2. Cylinder Hedging Model

Here we will describe the Cylinder Hedging Model, an algorithm that provides a hedging strategy based on the position of assets within a portfolio. The most fundamental cylinder model has two main parameters: (1) a pair of long and short limits (typically specified in US dollars) and (2) a desired hedge fraction specifying how much to hedge if one of these limits is breached. The limits define a so-called “cylinder” of risk, aiming to prevent the underlying position from growing too large, i.e., if the long limit is breached, the MM would place hedge sell trades to reduce the overall net position according to the hedge fraction (and vice versa short limit breaches). In our application of the cylinder model, the position is recorded at set intervals and a hedge fraction is placed for the duration of the interval. This is a natural hedging model for a MM where the portfolio being hedged is dictated by the flow of client trades. In Figures 1 and 2 we explain this by using a ‘toy’ example of a cylinder model with symmetrical long and short limits of 50 USD and -50 USD respectively (see red dotted lines), and a hedge fraction of 50% over 15 trial epochs. Figure 1 shows that at Trial 0, the client position (in blue) is at 100 USD. This breaches the long cylinder limit which in turn triggers the cylinder model to create an offsetting 50 USD hedge position as indicated by the orange line. This results in an overall net position (green line) of 50 USD. Throughout each trial this basic algorithm is repeated, resulting in the MM’s overall net position staying within the predefined ± 50 USD cylinder limits. Figure 2 shows the PnL values that result from trading this basic cylinder model over the 15 trials for the client (in blue), hedge (in orange) and overall net (in green). The raw client PnL ranges from -20 to $+55$ USD, yet after hedging, the volatility in the PnL is reduced between -10 and 45 USD.




Fig. 1: Cylinder Model Position




Fig. 2: Cylinder Model PnL

One natural extension often used to increase profitability of the cylinder model is to skew the limits based on the view of the market direction. The motivation for this is illustrated in Table 1. If the price of the underlying asset is going up, the MM may want to skew the long cylinder limit upwards to prevent the model triggering any hedges, thus permitting the MM to “ride the trend” in its existing long position and generate even further profit.

By making the cylinder limits dynamic and asymmetric, MMs can hedge more selectively based on their net position and the overall market price movements.

Table 1: MM Optimal Hedge Decision

Market Condition	MM Position Long (Client Position Short)	MM Position Short (Client Position Long)
Price Increase	MM PnL Increase Skew long cylinder upwards → less likely to hedge If MM position breaches skewed long limit → place sell hedge	MM PnL Decrease If MM position breaches long cylinder limit → place sell hedge
Price Decrease	MM PnL Decrease If MM position breaches long cylinder limit → place buy hedge	MM PnL Increase Skew short cylinder downwards → less likely to hedge If MM position breaches skewed short limit → place buy hedge

In this paper we use dynamic cylinder models as our experts. To appropriately skew each models' cylinder limits, we will make use of the moving averages of the underlying asset prices as our market directional indicators. Moving averages are computed over various time windows and are commonly used in technical analysis. Thus when the current asset price is higher than its moving average counterpart, we forecast the market to be rising and hence the long cylinder limit can be positively skewed. Likewise, if the current asset price is lower than its moving average counterpart, then the market is likely to trend downwards and the short cylinder limit can be skewed negatively. This dynamic adjustment of the cylinder limits can allow the model to both profit from and hedge against market corrections. Algorithm 1 provides us with the Pseudocode for the cylinder hedging model.

Figures 3 and 4 illustrate how the moving average is used to compute such dynamic skewed cylinder limits. Figures 5 and 6 show two ways of assessing the performance of the respective cylinder hedging models, tracking PnL and drawdowns. In these examples cylinder limits have been computed over the first 10,000 hourly epochs of a real-life EUR/USD trading dataset (described later in Section 4.1). Figure 3 shows the raw underlying price of the asset (EUR/USD) via the blue line. The moving average of the asset price computed over a window of 140 hours is shown by the red line. This line is smoother and lags behind as it is being computed. For each time epoch, we compute the market directional indicator, which is when the price line is above or below the slower moving average line. This indicator skews the cylinder limits up or down respectively. The results of using the indicator in Figure 3 to skew the cylinder limits are shown in Figure 4, where a long limit of 10 million USD and a short limit of -20 million USD are both being skewed up and down by 70% and 80% respectively throughout. Figures 5 and 6 highlight the benefits to the cylinder hedging model in using these skewed limits. The overall client PnL (blue line)

whilst making a profit does suffer from large drawdowns as shown in Figure 6, notably around trials 3000 and 8000. By observing the effect of hedging (orange line) it is clear that the cylinder hedging model has significantly reduced these losses, increasing the final net PnL (green line) by 51% from 538,024 USD to 813,484 USD.




Fig. 3: Price of Underlying Asset




Fig. 4: Client Position with Skewed Limits




Fig. 5: Client, Hedge and Net PnL




Fig. 6: Client, Hedge and Net Drawdown

We can achieve different risk reward profiles by varying our main parameters to the cylinder model, namely the limits and the hedge fraction, but also by varying the window size of the moving average signal to best capture the main shifts in market price. Evidently the task of making accurate long term price movement forecasts is extremely difficult [Fama \(1970\)](#) and whilst moving averages are clearly lagging indicators, our results will later show that using these to dynamically adjust cylinder limits can give good results over different time epochs in this study's historical dataset.

We have made use of moving average market indicators to skew the cylinder limits in this study, yet there are significant challenges presented by the data flows involved. Both the market and underlying client flow evolve through time. Market trends can persist or change rapidly, such that too large a moving average window size could result in signals not skewing limits in time. Conversely, setting the window size too small could feasibly result in erratic oscillations failing to capitalise on any trend. In addition, the nature of client order flow is complex: new clients join and leave, each client has different capitals to risk, differing trading time horizons and risk appetites, all of which can result in client positions growing and shrinking accordingly. If the cylinder limits are set too large / too small this could result in hedges being triggered too little / too much. We can derive rough estimates of what appropriate cylinder limits and moving averages to use based on historic client and price data, but this offers limited guarantee for future success. In this paper we aim to use the Weak Aggregating Algorithm to combine the hedge predictions from a pool of dynamic skewed cylinder model experts, each with differing parameters (i.e. limits, hedge fractions and moving average windows), to give an optimal hedging model that attempts to maximise PnL and minimise drawdowns.

Algorithm 1 Cylinder Hedging Model

Parameters: long/short Limit, Hedge fraction and Skew: $L_l, L_s, H_l, H_s, S_l, S_s$
 Directional indicators $Id_t, t = 1, 2, \dots$

```

for  $t = 1, 2, \dots$  do
  if  $Position_t^C > L_l + (L_l \times S_l \times Id_t)$  then
    | Hedge Fraction $_t \leftarrow H_l$ 
  end
  if  $Position_t^C < L_s + (L_s \times S_s \times Id_t)$  then
    | Hedge Fraction $_t \leftarrow H_s$ 
  end
  else
    | Hedge Fraction $_t \leftarrow 0$ 
  end
end

```

3. Weak Aggregating Algorithm

In this section, we will describe the framework of prediction with expert advice, explain the exponential average forecaster of [Cesa-Bianchi and Lugosi \(2006, Section 2\)](#) (we will refer to it as the Weak Aggregating Algorithm after [Kalnishkan and Vyugin \(2008\)](#), who developed and studied it independently), and introduce a modification of it for the discounted loss.

3.1. Prediction with Expert Advice Framework

Consider the following prediction scenario. On every step $t = 1, 2, \dots$, the learner L produces a *prediction* $\gamma_t \in \Gamma$, where Γ is a known prediction space. The nature produces a *loss function* $\lambda_t : \Gamma \rightarrow \mathbb{R}$ and the learner suffers loss $\ell_t = \lambda_t(\gamma_t)$. We measure the performance

of L by the *cumulative loss* over T steps given by

$$\text{Loss}_T(\mathcal{L}) = \sum_{t=1}^T \ell_t .$$

We want the cumulative loss to be as low as possible.

Now suppose that there are N *experts* \mathcal{E}_n , $n = 1, 2, \dots, N$, making prediction in the same environment as \mathcal{L} so that their predictions are available to \mathcal{L} before it makes its own. We will treat the experts as black boxes and will not be concerned with their internal mechanics. It is an important requirement that their predictions are available to \mathcal{L} before it makes its own and that they will suffer loss according to the same function λ_t . The interaction with experts may be described by Protocol 1.

Protocol 1 Prediction with Expert Advice Protocol

for $t = 1, 2, \dots$ **do**

experts \mathcal{E}_n output predictions $\gamma_t^n \in \Gamma$, $n = 1, 2, \dots, N$
 learner \mathcal{L} outputs a prediction $\gamma_t \in \Gamma$
 nature produces a function $\lambda_t : \Gamma \rightarrow \mathbb{R}$
 experts \mathcal{E}_n suffer losses $\ell_t^n = \lambda_t(\gamma_t^n)$, $n = 1, 2, \dots, N$
 learner \mathcal{L} suffers loss $\ell_t = \lambda_t(\gamma_t)$

end

We want the cumulative loss $\text{Loss}_T(\mathcal{L})$ to be small compared to the minimum of experts' losses $\text{Loss}_T(\mathcal{E}_n) = \sum_{t=1}^T \ell_t^n$. Formally one can think of \mathcal{L} as a *merging strategy*

$$\mathcal{L} : ((\Gamma \times \mathbb{R})^N)^* \times \Gamma^N \rightarrow \Gamma$$

turning an array of experts' predictions and a history of their prediction and losses into its own prediction.

We will aim to impose minimal restrictions on the loss functions λ_t output by the nature; we will not be assuming that the nature can be modelled in a reasonable sense.

Remark 1 *A simple scenario covered by Protocol 1 is the one where $\Gamma = [0, 1]$ and $\lambda_t(\gamma) = |\gamma - \omega_t|$, where ω_t is generated by the nature. Here the aim of a predictor is to output predictions γ_t approximating outcomes ω_t .*

3.2. Weak Aggregating Algorithm

Let Γ be a convex set so that for any $\gamma_1, \gamma_2, \dots, \gamma_N \in \Gamma$ and probabilities p_1, p_2, \dots, p_N ($p_n \geq 0$ for $n = 1, 2, \dots, N$ and $\sum_{n=1}^N p_n = 1$) the convex combination $\gamma = \sum_{n=1}^N p_n \gamma_n$ is defined and belongs to Γ . Then the learner \mathcal{L} can use Algorithm 2, which we will call the Weak Aggregating Algorithm (WAA).

In order to obtain performance bounds for WAA, one needs to assume convexity of loss functions λ_t ; this ensures the inequality $\ell_t \leq \sum_{n=1}^N p_{t-1}^n \ell_t^n$. We will also need losses to be bounded. Let $L \in \mathbb{R}$ be such that

$$\max_{n=1,2,\dots,N} \ell_t^n - \min_{n=1,2,\dots,N} \ell_t^n \leq L \tag{1}$$

Algorithm 2 Weak Aggregating Algorithm

Parameters: Initial distribution q_1, q_2, \dots, q_N , $q_n \geq 0$ for $n = 1, 2, \dots$ and $\sum_{n=1}^N q_n = 1$
 Learning rates $\eta_t > 0$, $t = 1, 2, \dots$

let $L_0^n = 0$, $n = 1, 2, \dots, N$

for $t = 1, 2, \dots$ **do**

 calculate weights $w_{t-1}^n = q_n e^{-\eta_t L_{t-1}^n}$, $n = 1, 2, \dots, N$
 normalise the weights $p_{t-1}^n = w_{t-1}^n / \sum_{i=1}^N w_{t-1}^i$, $n = 1, 2, \dots, N$
 read experts' predictions $\gamma_t^n \in \Gamma$, $n = 1, 2, \dots, N$
 output $\gamma_t = \sum_{n=1}^N p_{t-1}^n \gamma_t^n$
 read experts losses ℓ_t^n , $n = 1, 2, \dots, N$
 update $L_t^n = L_{t-1}^n + \ell_t^n$, $n = 1, 2, \dots, N$

end

for every $t = 1, 2, \dots$. This is guaranteed if, for example, $\sup_{\Gamma} \lambda_t(\gamma) - \inf_{\Gamma} \lambda_t(\gamma) \leq L$ for all $t = 1, 2, \dots$. Sometimes we assume that L is known in advance.

Theorem 2 *Let the learning rates in WAA be $\eta_t = c/\sqrt{t}$ for every $t = 1, 2, \dots$, where $c > 0$. If all loss functions λ_t are convex and L satisfies (1) for $t = 1, 2, \dots$ then*

$$\text{Loss}_T(\mathcal{L}) \leq \text{Loss}_T(\mathcal{E}_n) + \frac{\sqrt{T} \ln(1/q_n)}{c} + \frac{cL^2\sqrt{T}}{4}$$

for all $T = 1, 2, \dots$ and all experts \mathcal{E}_n , $n = 1, 2, \dots, N$.

Corollary 3 *Under the conditions of Theorem 2, if L satisfying (1) is known in advance, one can take $c = 2\sqrt{\ln N}/L$ and ensure for equal weights $q_1 = q_2 = \dots = q_N = 1/N$ the bound*

$$\text{Loss}_T(\mathcal{L}) \leq \text{Loss}_T(\mathcal{E}_n) + L\sqrt{T \ln N}$$

for all $T = 1, 2, \dots$ and all experts \mathcal{E}_n , $n = 1, 2, \dots, N$.

These results improve on both Corollary 14 by Kalnishkan and Vyugin (2008) and Theorem 2.3 by Cesa-Bianchi and Lugosi (2006). An equivalent result was obtained by Chernov (2010). The theorem can be proven along the same lines as the result for discounted loss below.

3.3. Discounted Loss

In this section, we will introduce WAA for discounted loss. Discounting losses with time in the context of prediction with expert advice was first considered by Chernov and Zhdanov (2010); see also a concise overview of discounting applied to the Aggregating Algorithm by Kalnishkan (2022, Section 9).

The following argument for discounting can be given. The learner may want to align itself with the experts that have performed well lately, rather than over the whole course of history. Distant past may be irrelevant, especially in the context of economics and finance. Discounting offers a convenient framework for discarding the past.

There is also a purely numerical reason. The cumulative losses L_t^n calculated by the WAA may grow quite large with time and, correspondingly, the weights w_t^n very close to zero. This problem can be ameliorated by shifting the losses, but to some extent it is unavoidable if there are experts performing very differently.

Suppose that we are given coefficients $\alpha_1, \alpha_2, \dots \in (0, 1]$. Let the cumulative discounted loss for a learner \mathcal{L} be given by

$$\widetilde{\text{Loss}}_T(\mathcal{L}) = \sum_{t=1}^T \lambda(\gamma_t) \left(\prod_{s=t}^{T-1} \alpha_s \right) = \alpha_{T-1} \widetilde{\text{Loss}}_{T-1}(\mathcal{L}) + \lambda(\gamma_T) ;$$

the discounted loss $\widetilde{\text{Loss}}_T(\mathcal{E}_n)$ of an expert \mathcal{E}_n is defined in a similar way.

Algorithm 3 is identical to the Weak Aggregating Algorithm except that it uses discounted losses for L_t^n . We will refer to it as Weak Aggregating Algorithm with discounting (WAAd).

Algorithm 3 Weak Aggregating Algorithm with Discounting

Parameters: Initial distribution q_1, q_2, \dots, q_N , $q_n \geq 0$ for $n = 1, 2, \dots$ and $\sum_{n=1}^N q_n = 1$.
 Discounting factors $\alpha_1, \alpha_2, \dots \in (0, 1]$.
 Learning rates $\eta_t > 0$, $t = 1, 2, \dots$

let $L_0^n = 0$, $n = 1, 2, \dots, N$

for $t = 1, 2, \dots$ **do**

calculate weights $w_{t-1}^n = q_n e^{-\eta_t \alpha_{t-1} L_{t-1}^n}$, $n = 1, 2, \dots, N$
 normalise the weights $p_{t-1}^n = w_{t-1}^n / \sum_{i=1}^N w_{t-1}^i$, $n = 1, 2, \dots, N$
 read experts' predictions $\gamma_t^n \in \Gamma$, $n = 1, 2, \dots, N$
 output $\gamma_t = \sum_{n=1}^N p_{t-1}^n \gamma_t^n$
 read experts losses ℓ_t^n , $n = 1, 2, \dots, N$
 update $L_t^n = \alpha_{t-1} L_{t-1}^n + \ell_t^n$, $n = 1, 2, \dots, N$

end

Theorem 4 *Let the learning rates in WAAd be positive and non-decreasing, $\eta_{t-1} \geq \eta_t > 0$, $t = 1, 2, \dots$. If all loss functions λ_t are convex and L satisfies (1) for $t = 1, 2, \dots$ then*

$$\text{Loss}_T(\mathcal{L}) \leq \text{Loss}_T(\mathcal{E}_n) + \frac{\ln(1/q_n)}{\eta_T} + \frac{L^2}{8} \sum_{t=1}^T \eta_t \prod_{s=t}^T \alpha_s \quad (2)$$

for all $T = 1, 2, \dots$ and all experts \mathcal{E}_n , $n = 1, 2, \dots, N$.

A proof sketch is given in Appendix A.

Corollary 5 *Under the conditions of Theorem 2, if L satisfying (1) is known in advance and all discounting factors are equal and less than 1, $0 < \alpha_1 = \alpha_2 = \dots = \alpha < 1$, one can take*

$$\eta_t = \eta = \frac{2\sqrt{2(1-\alpha)\ln N}}{L}$$

and ensure for equal weights $q_1 = q_2 = \dots = q_N = 1/N$ the bound

$$\text{Loss}_T(\mathcal{L}) \leq \text{Loss}_T(\mathcal{E}_n) + L\sqrt{\frac{\ln N}{2(1-\alpha)}} \quad (3)$$

for all $T = 1, 2, \dots$ and all experts \mathcal{E}_n , $n = 1, 2, \dots, N$.

Remark 6 *The discounted loss would not normally grow with time as regular loss usually does. If $0 \leq \lambda_t(\gamma) \leq L$ and the discounting factors are constant, then*

$$\widetilde{\text{Loss}}_T(\mathcal{L}) \leq \frac{L}{1-\alpha} .$$

Our bound will then be meaningful only if the extra term in (3) is much less than this trivial bound, i.e.,

$$L\sqrt{\frac{\ln N}{2(1-\alpha)}} \ll \frac{L}{1-\alpha} .$$

This happens if α is close to 1 and the difference $1-\alpha$ is small compared to $2/\ln N$.

3.4. WAA for Hedging

In this section, we discuss how hedging can be considered within the framework of prediction with expert advice. The aim of the learner is to find the optimal hedge fraction for a MM hedging the risk associated with client positions. The pool of experts here are a set of cylinder models with different input parameters producing different hedge fractions.

Here we consider the case of the client position in a single asset, and a hedging decision is represented by $\gamma \in [-1, 0]$, where $\gamma_t = -1$ implies hedging out the entire client position and $\gamma_t = 0$ corresponds to a decision not to hedge over trial t at all.

It is natural to define loss in terms of the the MM's PnL resulting from facilitating client orders. Note, as PnL represents the MM's gain, we need to take its inverse when defining the loss. We can therefore take the loss at time t to be $\lambda(\gamma_t) = -\text{PnL}_t\gamma_t$. The cumulative loss over T trials is the negation of PnL over these trials, i.e., the amount MM has lost.

Since $\text{PnL} \in \mathbb{R}$ is linear in the hedge fraction, the loss functions λ_t are convex. In order to establish a bound on loss, we need to see whether we can define limits on the value of the MM's client PnL. As discussed, client PnL is a function of client position and the price of the underlying asset. Let us first consider the case when the MM's client position is long. Here PnL is bounded by the value of the net position, because in the worst possible scenario the value of the underlying can fall to zero and the loss in PnL is equal to the value of the net position. In the case where the MM takes a short position there is no simple bound on the possible loss. However, based on the interval of predictions and historical market volatility we can make assumptions on the maximum rise in value of the underlying asset and therefore the bound on loss.

Following [Al-Baghdadi et al. \(2020\)](#), we introduce downside and combined losses. We aim to increase the penalty for losses and reduce the reward for gains. In the context of optimisation of investment portfolios, this has been shown to reduce the drawdown of

a portfolio; combined loss has proven particularly effective. Here we will take a similar approach considering the loss function with the coefficients $u \geq 0$ and $v \geq 0$:

$$\lambda(\gamma) = - \left(\frac{u}{u+v} \text{PnL}\gamma + \frac{v}{u+v} \min(\text{PnL}\gamma, 0) \right)$$

This has the aim of allowing the learner to focus on finding hedging strategies that minimise drawdown to provide smoother returns. Since $-\min(-x, 0)$ is convex in x , the loss function remains convex.

4. Experiments

To empirically evaluate the effectiveness of the WAAd at finding optimal parameters for the cylinder hedging model, we have conducted experiments with a real-world Foreign Exchange (FX) dataset drawn from the data flows of trading three currency pairs during February 2014 to June 2017.

4.1. Data Set

The data set we will be using is real-world currency exchange data, based on the trading behaviour of individuals opening positions with a FX MM. The dataset focuses on the net position of the following three currency pairs - EUR/USD, GBP/USD and EUR/GBP, over a 41 month period (Feb 2014 - June 2017) represented in hourly epochs. Figures 7 through 15 show the price of each of the currency pairs over this period in addition to the client position and resulting PnL. Note that client position refers to the MM’s position resulting from client orders, and similarly client PnL is the MM’s PnL resulting from the net client position. In order to feed expert predictions to the WAAd, the data was partitioned into regularised time intervals using a technique known as DAPRA, as outlined in [Al-baghdadi et al. \(2019\)](#). A copy of the data and WAAd implementation can be found at [Al-baghdadi \(2022\)](#).






Fig. 7: EUR/USD Price

Fig. 8: EUR/USD Position

Fig. 9: EUR/USD PnL

For each currency pair, 100 unique cylinder model parameter combinations were chosen, resulting in the Net PnL’s shown in Figures 16 to 18. These are the expert predictions used by the WAAd to generate a learner prediction. It is important to note that for each currency pair there are experts that both improve and worsen the MM’s overall PnL and drawdown.




Fig. 10: GBP/USD Price




Fig. 11: GBP/USD Position

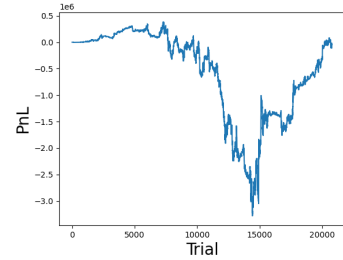


Fig. 12: GBP/USD PnL




Fig. 13: EUR/GBP Price

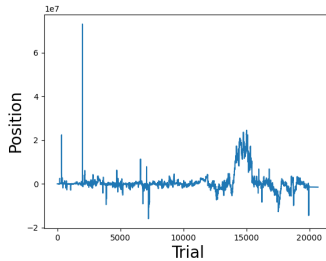


Fig. 14: EUR/GBP Position




Fig. 15: EUR/GBP PnL

4.2. Numerical Analysis

We will now analyse the performance of the WAAd on each of the currency pairs. To do this we use the Calmar Ratio, which is a well defined risk metric used to evaluate the performance of a portfolio, first introduced in [Young \(1991\)](#) and defined as follows:

$$\text{Calmar}(T) = \frac{\text{Return}[0, T]}{|\text{Max Drawdown}(T)|} \quad (4)$$

For each currency pair, a scatter plot of Net PnL against maximum drawdown is used to illustrate how the Calmar ratio of the un-hedged and expert portfolios compares to that of the WAAd's. [Tables 2 to 4](#) provide the maximum drawdown, total PnL and the mean and standard deviation of each epoch's PnL values. Each table shows the results of three useful benchmarks to evaluate the performance of each WAAd result: (1) un-hedged (just client results), (2) the best and (3) the worst of the experts used in the study. We wish to maximise the Calmar, maximum drawdown, PnL and mean PnL measures, whilst minimise the standard deviation of the PnL.

[Figure 19](#) shows a plot of maximum drawdown against final PnL of un-hedged, expert and learner portfolios, all for EUR/USD. We have used different shapes to represent each of the combined loss parameters for the WAAd, and different colours to represent the various discount factors. Models with the highest Calmar ratio and therefore optimal, will by definition be located in the top right of the plot. We can see the worst hedging strategy produced by the WAAd is with a combined loss of $u = 1$ and $v = 0$ and no discounting.

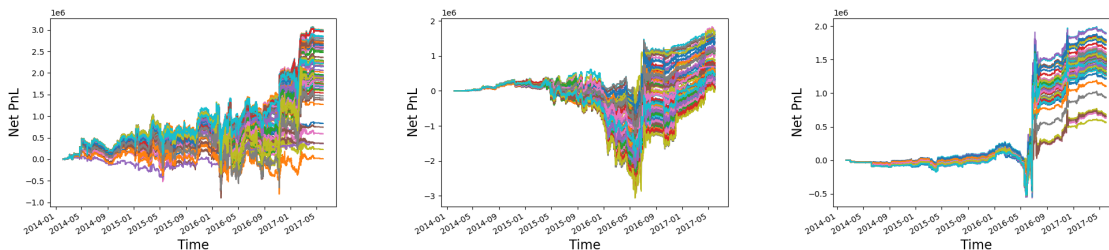


Fig. 16: EUR/USD Net PnL Fig. 17: GBP/USD Net PnL Fig. 18: EUR/GBP Net PnL

Following this strategy results in a Calmar ratio of 0.58 and a decrease in PnL of 55% only reducing max drawdown by 6% when compared to the un-hedged portfolio.




Fig. 19: EUR/USD Expert and WAA PnL against Max Drawdown.

Discount Key Red: 0%, Cyan: 2.5%, Purple: 5%, Black: 7.5%, Pink: 10%, Orange: 20%

If we refer to Table 2 we can see the WAA learner with the highest Calmar ratio of 3.63, taking combined loss coefficients of $u = 0$ and $v = 1$ applying a discount factor of 10%. This results in a 9.7% increase in PnL and a 62.6% decrease in drawdown when compared to that of the un-hedged portfolio.

Table 2: EUR/USD WAA Table of results

Hedging Models			Calmar	Max Drawdown (10^6)	PnL (10^6)	Mean PnL	PnL Standard Deviation
u	v	Discount %					
Un-hedged			1.22	-1.50	1.90	92	15,988
Best Model			3.63	-0.58	2.10	100	8,448
Worst Model			0.58	-1.40	0.84	40	10,731
1	0	0	0.58	-1.40	0.84	40	10,731
1	0	2.5	2.13	-0.83	1.80	85	10,741
1	0	5	1.64	-1.00	1.60	79	10,856
1	0	7.5	1.6	-1.00	1.70	80	10,897
1	0	10	1.61	-1.10	1.70	83	10,918
1	0	15	1.7	-1.10	1.90	90	10,958
1	0	20	1.8	-1.10	2.00	96	10,997
1	1	0	0.96	-0.46	0.44	21	4,065
1	1	2.5	3.12	-0.49	1.50	74	7,337
1	1	5	3.17	-0.57	1.80	87	8,658
1	1	7.5	2.9	-0.67	1.90	93	9,319
1	1	10	2.53	-0.78	2.00	95	9,698
1	1	15	2.2	-0.92	2.00	97	10,130
1	1	20	2.16	-0.99	2.10	102	10,381
0	1	0	0.85	-0.49	0.41	20	4,029
0	1	2.5	2.57	-0.41	1.10	51	5,646
0	1	5	3.15	-0.50	1.60	75	6,831
0	1	7.5	3.51	-0.54	1.90	91	7,783
0	1	10	3.63	-0.58	2.10	100	8,448
0	1	15	3.36	-0.65	2.20	106	9,248
0	1	20	2.81	-0.80	2.20	108	9,710
2	1	0	2.27	-0.46	1.10	51	4,701
2	1	2.5	3.16	-0.53	1.70	81	8,396
2	1	5	2.65	-0.68	1.80	86	9,404
2	1	7.5	2.33	-0.80	1.90	89	9,856
2	1	10	2.1	-0.89	1.90	90	10,113
2	1	15	1.98	-0.99	2.00	94	10,411
2	1	20	2.01	-1.00	2.10	100	10,591
1	2	0	0.7	-0.53	0.37	18	4,032
1	2	2.5	3.02	-0.46	1.40	66	6,558
1	2	5	3.32	-0.53	1.80	85	7,965
1	2	7.5	3.34	-0.59	2.00	94	8,785
1	2	10	3.1	-0.66	2.00	98	9,279
1	2	15	2.49	-0.84	2.10	100	9,841
1	2	20	2.33	-0.93	2.20	104	10,164

As we observed with EUR/USD, all trials of the WAA hedging GBP/USD reduce the MM's drawdown and in this case also increase PnL significantly. The model with the highest Calmar ratio of 2.86, is that with combined loss coefficients of $u = 0$ and $v = 1$ and a discount factor of 2.5%. When comparing this to the MM's un-hedged portfolio, we observe that the PnL is increased from a loss of 42,908 USD to a profit of 1,123,464 USD. This is possible due to the significant drawdowns in the client PnL over the initial 1500 trials and is reflected in the reduction of the maximum drawdown from 3,663,188 USD to 392,370 USD.

However, it is important to note that without applying discounting to experts loss and taking combined loss coefficients of $u = 1$ and $v = 1$ or $u = 2$ and $v = 1$, we can in fact find solutions with higher PnL at the expense of increased maximum drawdown.




Fig. 20: GBP/USD Expert and WAA PnL against Max Drawdown.
Discount Key Red:0%, Cyan: 2.5%, Purple: 5%,
 Black: 7.5%:, Pink: 10%, Orange: 20%

EUR/GBP is an interesting experiment - as is shown in Figure 15 the majority of the PnL is a result of a surge in the underlying currency pair value. Figure 21 shows that the un-hedged portfolio is relatively high, yet has suffered from significant maximum drawdown. In this case combined loss coefficients of $u = 1$ and $v = 1$ provide the optimal Calmar ratio, with a decrease in maximum drawdown value of 42% and an increase in PnL of 6% compared to the un-hedged portfolio.

It may not appear immediately obvious as to why the application of discounted loss on EUR/USD is apparently more effective at improving the Calmar ratio than its use on GBP/USD and EUR/GBP. However some intuition can be gleaned if one compares the results of the WAA to the data outlined in section 4.1. When examining the accumulation of the MM's PnL in EUR/USD there is no clear period of profit or loss. One possible explanation for this may be found by referring to the price of the currency pair over the experiment. Excluding the initial quarter of the trial shows there is no dominant trend in the price. The result of this is that there is no clear optimal hedging strategy for any sustained period throughout the experiment. If we compare this to the client PnL of the GBP/USD currency pair, we see there is a clear downward trend over the initial three quarters, followed by an upward trend for the remainder of the experiment with the MM netting a small loss. Naturally, if there is a clear trend in the direction of PnL we can expect

Table 3: GBP/USD WAA Table of results

Hedging Models			Calmar	Max Drawdown (10^6)	PnL (10^6)	Mean PnL	PnL Standard Deviation
u	v	Discount %					
Un-hedged			-0.01	-3.70	-0.04	-2	12,366
Best Model			2.86	-0.39	1.10	54	4,354
Worst Model			0.11	-1.80	0.19	9	6,922
1	0	0	0.11	-1.80	0.19	9	6,922
1	0	2.5	0.49	-1.70	0.82	39	7,945
1	0	5	0.4	-1.80	0.70	34	8,025
1	0	7.5	0.35	-1.80	0.64	31	8,042
1	0	10	0.32	-1.90	0.60	29	8,050
1	0	15	0.29	-2.00	0.57	28	8,058
1	0	20	0.29	-2.00	0.58	28	8,064
1	1	0	2.12	-0.60	1.30	61	5,531
1	1	2.5	1.24	-0.81	1.00	48	5,749
1	1	5	0.61	-1.20	0.75	36	6,543
1	1	7.5	0.48	-1.40	0.69	33	6,963
1	1	10	0.42	-1.60	0.67	32	7,204
1	1	15	0.37	-1.70	0.64	31	7,467
1	1	20	0.35	-1.80	0.64	31	7,610
0	1	0	1.78	-0.60	1.10	51	5,846
0	1	2.5	2.86	-0.39	1.10	54	4,354
0	1	5	1.13	-0.82	0.93	44	5,075
0	1	7.5	0.7	-1.10	0.76	36	5,670
0	1	10	0.54	-1.30	0.67	32	6,102
0	1	15	0.4	-1.50	0.60	29	6,650
0	1	20	0.36	-1.60	0.59	29	6,981
2	1	0	1.84	-0.70	1.30	62	5,618
2	1	2.5	0.78	-1.10	0.85	41	6,474
2	1	5	0.52	-1.40	0.73	35	7,099
2	1	7.5	0.43	-1.60	0.68	33	7,373
2	1	10	0.39	-1.70	0.65	31	7,522
2	1	15	0.35	-1.80	0.62	30	7,686
2	1	20	0.33	-1.90	0.62	30	7,776
1	2	0	1.85	-0.60	1.10	53	5,611
1	2	2.5	1.77	-0.61	1.10	52	5,118
1	2	5	0.76	-1.10	0.82	39	5,983
1	2	7.5	0.55	-1.30	0.70	34	6,513
1	2	10	0.46	-1.50	0.67	32	6,842
1	2	15	0.39	-1.70	0.64	31	7,215
1	2	20	0.36	-1.80	0.64	31	7,420

that a pool of models will consistently outperform the majority over this period. Therefore, the benefit of using discounted loss is diminished in such scenarios.

5. Conclusion

Building on previous work, we have shown that the Weak Aggregating Algorithm (WAA) can be used to combine the predictions from a pool of cylinder hedging models to improve key performance metrics - namely the overall profit (PnL) - whilst simultaneously not compromising on the smoothness of returns by minimising drawdowns. In this study we have further introduced a method for applying discounted loss to the WAA (WAA_d). Using a real-world Foreign Exchange (FX) trading dataset we demonstrate empirical efficacy of our




Fig. 21: EUR/GBP Expert and WAA PnL against Max Drawdown.

Discount Key Red:0%, Cyan: 2.5%, Purple: 5%,
 Black: 7.5%, Pink: 10%, Orange: 20%

approach for three major currency pairs, varying the parameters of the WAA namely the use of combined loss and discount factor. Whilst our analysis was carried out for individual FX currency pairs, a possible extension to this work would allow combinations of pairs into a *single currency risk*, allowing one to focus on individual EUR, USD and GBP positions that build up across the triad of currency pairs (i.e. EUR/USD, GBP/USD, EUR/GBP). Due to the nature of these triangular relationships there is enforced correlation between the price movements. These correlations strengthen and weaken in time and are typically fed in to compute a Value At Risk (VaR) measure, introduced in 1994 by J. P. Morgan [Guldimann \(1995\)](#).

We have observed that discount loss can be varied in the WAA to achieve an improved Calmar ratio and its effectiveness is a product of market conditions and client behaviour. It is worth remembering that the underlying data is a complex combination of two major factors: (1) the trading activity of the MMs clients, and (2) the price movements of the symbols being traded. The client net position is the aggregate risk across a *dynamic* portfolio of trader activity: new traders join, some traders leave over time, each client presents with differing risk appetites, time horizons and capital that they can trade with. Market conditions like volatility of the underlying prices and correlation between different symbols also vary profoundly through time due to various macro economic factors. In future work we

Table 4: EUR/GBP WAA Table of results

Hedging Models			Calmar	Max Drawdown (10^5)	PnL (10^6)	Mean PnL	PnL Standard Deviation
u	v	Discount %					
Un-hedged			2.11	-7.30	1.50	76	9,194
Best Model			3.68	-4.20	1.60	75	7,273
Worst Model			2.1	-6.00	1.20	60	7,017
1	0	0	3.08	-5.00	1.50	75	7,434
1	0	2.5	2.84	-5.70	1.60	78	7,923
1	0	5	2.71	-5.80	1.60	76	7,985
1	0	7.5	2.63	-5.80	1.50	74	8,009
1	0	10	2.56	-5.90	1.50	73	8,023
1	0	15	2.44	-6.00	1.50	71	8,039
1	0	20	2.36	-6.10	1.40	69	8,049
1	1	0	3.68	-4.20	1.60	75	7,273
1	1	2.5	2.65	-5.00	1.30	64	6,417
1	1	5	2.57	-5.40	1.40	68	6,967
1	1	7.5	2.54	-5.60	1.40	69	7,343
1	1	10	2.48	-5.70	1.40	69	7,520
1	1	15	2.35	-5.90	1.40	67	7,689
1	1	20	2.28	-6.00	1.40	66	7,776
0	1	0	2.87	-4.90	1.40	68	6,882
0	1	2.5	2.87	-4.40	1.30	61	5,700
0	1	5	2.35	-5.10	1.20	58	5,789
0	1	7.5	2.18	-5.40	1.20	57	6,080
0	1	10	2.15	-5.60	1.20	58	6,323
0	1	15	2.12	-5.80	1.20	60	6,719
0	1	20	2.1	-6.00	1.20	60	7,017
2	1	0	3.41	-4.50	1.50	74	7,354
2	1	2.5	2.86	-5.30	1.50	73	7,116
2	1	5	2.79	-5.50	1.50	75	7,585
2	1	7.5	2.66	-5.70	1.50	73	7,719
2	1	10	2.56	-5.80	1.50	72	7,787
2	1	15	2.41	-5.90	1.40	69	7,861
2	1	20	2.32	-6.00	1.40	68	7,903
1	2	0	3.37	-4.40	1.50	72	7,133
1	2	2.5	2.7	-4.80	1.30	63	6,126
1	2	5	2.26	-5.30	1.20	58	6,255
1	2	7.5	2.28	-5.60	1.30	61	6,696
1	2	10	2.3	-5.70	1.30	63	7,024
1	2	15	2.26	-5.90	1.30	64	7,388
1	2	20	2.22	-6.00	1.30	64	7,568

hope to investigate the use of dynamic discounting based on the observations in fundamental changes of trading and market activity to capitalise when “regime change” is detected.

Acknowledgments

The authors would like to acknowledge the support of Algorithmic Laboratories Ltd (AlgoLabs) and Royal Holloway, University of London for establishing and developing this research.

References

- Najim Al-baghdadi. Waa cylinder hedging model git repository. 5 2022. URL https://github.com/Najim-Al/PhD_AA.
- 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.
- Najim Al-Baghdadi, David Lindsay, Yuri Kalnishkan, and Sian Lindsay. Practical investment with the long-short game. In Alexander Gammernan, Vladimir Vovk, Zhiyuan Luo, Evgeni Smirnov, and Giovanni Cherubin, editors, *Proceedings of the Ninth Symposium on Conformal and Probabilistic Prediction and Applications*, volume 128 of *Proceedings of Machine Learning Research*, pages 209–228. PMLR, 09–11 Sep 2020. URL <https://proceedings.mlr.press/v128/al-baghdadi20a.html>.
- N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- A. Chernov. On Theorem 2.3 in “Prediction, learning, and games” by Cesa-Bianchi and Lugosi. *CoRR*, abs/1011.5668, 2010.
- A. Chernov and F. Zhdanov. Prediction with expert advice under discounted loss. In *Proceedings of ALT 2010*, volume LNAI 6331, pages 255–269. Springer, 2010. See also arXiv:1005.1918 [cs.LG].
- Thomas M Cover and Erik Ordentlich. Universal portfolios with side information. *IEEE Transactions on Information Theory*, 42(2):348–363, 1996.
- Eugene F Fama. Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25(2):383–417, 1970.
- T. Guldumann. *Riskmetrics TM*, 1995.
- Jin’an He and Xingyu Yang. Universal portfolio selection strategy by aggregating online expert advice. *Optimization and Engineering*, pages 1–25, 2020.
- Y. Kalnishkan. Prediction with expert advice for a finite number of experts: A practical introduction. *Pattern Recognition*, page 108557, 2022.
- Yuri Kalnishkan and Michael V Vyugin. The weak aggregating algorithm and weak mixability. *Journal of Computer and System Sciences*, 74(8):1228–1244, 2008.
- Volodya Vovk and Chris Watkins. Universal portfolio selection. In *Proceedings of the eleventh annual conference on Computational learning theory*, pages 12–23. ACM, 1998.
- Xingyu Yang, Jin’an He, Jiayi Xian, Hong Lin, and Yong Zhang. Aggregating expert advice strategy for online portfolio selection with side information. *Soft Computing*, 24(3):2067–2081, 2020.

Terry W Young. Calmar ratio: A smoother tool. *Futures*, 20(1):40, 1991.

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

Appendix A. Proof Sketch

Proof [of Theorem 4]

Without loss of generality one can assume that $\ell_t^n \in [0, L]$ for all $n = 1, 2, \dots, N$ and $t = 1, 2, \dots$. Indeed, let us change λ_s to $\lambda_s + C_s$ for some constants $C_s \in \mathbb{R}$, $s = 1, 2, \dots$. The discounted cumulative losses L_t^n then shift by some values independent of n . In the expressions for p_{t-1}^n , the shifts cancel out and the value of p_{t-1}^n will be unaffected. Therefore the predictions γ_t will not change and ℓ_t will change by C_t in line with ℓ_t^n . In inequality (2), which we need to prove, the values C_s cancel out.

Lemma A.1 by [Cesa-Bianchi and Lugosi \(2006\)](#) implies that

$$\ell_t \leq \sum_{n=1}^N p_n^{t-1} \ell_t^n \leq -\frac{1}{\eta_t} \ln \sum_{n=1}^N p_n^{t-1} e^{-\eta_t \ell_t^n} + \frac{L^2}{8} \eta_t, \quad (5)$$

$t = 1, 2, \dots, T$. Multiplying the inequality for time t by $\prod_{s=t}^{T-1} \alpha_s$ and adding them together yields

$$L_t \leq -\sum_{t=1}^T \frac{1}{\eta_t} \prod_{s=t}^{T-1} \alpha_s \ln \sum_{n=1}^N p_{t-1}^n e^{-\eta_t \ell_t^n} + \frac{L^2}{8} \sum_{t=1}^T \eta_t \prod_{s=t}^{T-1} \alpha_s. \quad (6)$$

Let us analyse the logarithm in this inequality. Substituting the expression for p_{t-1}^n yields

$$\begin{aligned} -\frac{1}{\eta_t} \ln \sum_{n=1}^N p_{t-1}^n e^{-\eta_t \ell_t^n} &= -\frac{1}{\eta_t} \ln \sum_{n=1}^N \frac{q_n e^{-\eta_t \alpha_{t-1} L_{t-1}^n}}{\sum_{m=1}^N q_m e^{-\eta_t \alpha_{t-1} L_{t-1}^m}} e^{-\eta_t \ell_t^n} \\ &= -\frac{1}{\eta_t} \ln \sum_{n=1}^N q_n e^{-\eta_t (\alpha_{t-1} L_{t-1}^n + \ell_t^n)} + \frac{1}{\eta_t} \ln \sum_{n=1}^N q_n e^{-\eta_t \alpha_{t-1} L_{t-1}^n}. \end{aligned}$$

In the first term, we get $\alpha_{t-1} L_{t-1}^n + \ell_t^n = L_t^n$. The second term can be upper bounded using Jensen's inequality

$$\begin{aligned} \frac{1}{\eta_t} \ln \sum_{n=1}^N q_n e^{-\eta_t \alpha_{t-1} L_{t-1}^n} &= \frac{1}{\eta_t} \ln \sum_{n=1}^N q_n e^{-\eta_{t-1} (\eta_t \alpha_{t-1} / \eta_{t-1}) L_{t-1}^n} \\ &\leq \frac{1}{\eta_t} \ln \left(\sum_{n=1}^N q_n e^{-\eta_{t-1} L_{t-1}^n} \right)^{\eta_t \alpha_{t-1} / \eta_{t-1}} \\ &= \frac{\alpha_{t-1}}{\eta_{t-1}} \sum_{n=1}^N q_n e^{-\eta_{t-1} L_{t-1}^n} \end{aligned}$$

as long as $\eta_t / \eta_{t-1} \leq 1$.

When we add over t in (6), the sum telescopes and we get (2). ■