Reinforcement Learning for Systematic FX Trading

This paper was downloaded from TechRxiv (<https://www.techrxiv.org>).

LICENSE

CC BY 4.0

SUBMISSION DATE / POSTED DATE

09-10-2021 / 14-10-2021

CITATION

Borrageiro, Gabriel; Firoozye, Nick; Barucca, Paolo (2021): Reinforcement Learning for Systematic FX Trading. TechRxiv. Preprint.
https://www.techrxiv.org/articles/preprint/Reinforcement_Learning_for_Systematic_FX_Trading/16778932

DOI

[10.36227/techrxiv.16778932.v1](https://doi.org/10.36227/techrxiv.16778932.v1)

Reinforcement Learning for Systematic FX Trading

Gabriel Borrageiro

GABRIEL.BORRAGEIRO.20@UCL.AC.UK

*Department of Computer Science
University College London
Gower Street, London, WC1E 6BT, UK*

Nick Firoozye

N.FIROOZY@UCL.AC.UK

*Department of Computer Science
University College London
Gower Street, London, WC1E 6BT, UK*

Paolo Barucca

P.BARUCCA@UCL.AC.UK

*Department of Computer Science
University College London
Gower Street, London, WC1E 6BT, UK*

Abstract

We conduct a detailed experiment on major cash fx pairs, accurately accounting for transaction and funding costs. These sources of profit and loss, including the price trends that occur in the currency markets, are made available to our recurrent reinforcement learner via a quadratic utility, which learns to target a position directly. We improve upon earlier work, by casting the problem of learning to target a risk position, in an online learning context. This online learning occurs sequentially in time, but also in the form of transfer learning. We transfer the output of radial basis function hidden processing units, whose means, covariances and overall size are determined by Gaussian mixture models, to the recurrent reinforcement learner and baseline momentum trader. Thus the intrinsic nature of the feature space is learnt and made available to the upstream models. The recurrent reinforcement learning trader achieves an annualised portfolio information ratio of 0.52 with compound return of 9.3%, net of execution and funding cost, over a 7 year test set. This is despite forcing the model to trade at the close of the trading day 5pm EST, when trading costs are statistically the most expensive. These results are comparable with the momentum baseline trader, reflecting the low interest differential environment since the 2008 financial crisis, and very obvious currency trends since then. The recurrent reinforcement learner does nevertheless maintain an important advantage, in that the model's weights can be adapted to reflect the different sources of profit and loss variation. This is demonstrated visually by a USDRUB trading agent, who learns to target different positions, that reflect trading in the absence or presence of cost.

Keywords: policy gradients, recurrent reinforcement learning, online learning, transfer learning, financial time series

1. Introduction

Forecasting in financial time series is commonly done through supervised learning approaches. For example, Tsay and Chen (2019) apply a number of parametric and non-parametric methods, including nonlinear state-space models, Bayesian and classical approaches, to nonlinear time series analysis. These forecasts are then used to take financial risk or positions, for either speculative, investment or hedging purposes. In earlier work,

Bengio (1997) applied learning algorithms to decision making with financial time series. He notes at the time that the traditional approach in this domain, is to train a model using a prediction criterion, such as minimizing mean-square prediction error, or maximizing the likelihood of a conditional model of the dependent variable. He finds that with noisy time series, that better results can be obtained when the model is directly trained in order to maximize the financial criterion of interest, here gains and losses (including those due to transactions) incurred during trading. In this spirit, we extend the earlier work of Moody and Wu (1997) and Gold (2003), where direct, recurrent reinforcement learning agents are put to work in financial trading strategies. Rather than optimising for an intermediate performance measure such as maximal forecast accuracy or minimal forecast error, which is still the traditional approach in this domain, we maximise a more direct performance measure such as quadratic economic utility. An advantage of the approach is that we are able to use the risk adjusted returns of the trading strategy, as well as the execution and funding cost, to influence learning of the model and update model parameters accordingly. The approach is detailed in section 3.1. We conduct experiments using foreign exchange data, see section 4. A background to the foreign exchange market is provided in section 2.2. We baseline the results of our recurrent reinforcement learner against a radial basis function network momentum trader and a model-free carry trader, whose design is discussed in section 3.3.

Summarising the main goals and results of this paper, we extend the earlier work of Moody and Wu (1997), Moody and Saffell (1998), Moody et al. (1998) and Gold (2003), where direct, recurrent reinforcement learning agents are put to work in financial trading strategies. Whereas the focus of Moody and Wu (1997) was on the use of the differential Sharpe ratio as a performance measure, we adopt the simpler quadratic utility of Sharpe (2007). This ameliorates the undesirable property of the Sharpe ratio in that it penalises a model that produces returns larger than $\frac{\sigma_t}{\mu_t}$, that is, the ratio of the standard deviation of returns to its expectation (Moody et al., 1998). Perhaps for this reason, along with the use of relatively weak features and shared backtest hyper-parameters, Gold (2003) obtained mixed results when experimenting with cash currency pairs. In contrast, our experiment with the major cash currency pairs, sees our recurrent reinforcement learning trader achieve an annualised portfolio information ratio of 0.52 with compound return of 9.3%, net of execution and funding cost, over a 7 year test set. This is despite forcing the model to trade at the close of the trading day 5pm EST, when trading costs are statistically the most expensive. Aside from the different utility function, we put these improved experiment results down to a combination of several factors. Firstly, we make use of more powerful feature engineering, in the shape of radial basis function hidden processing units, whose means, covariances and overall size are determined by an unsupervised learning procedure for finite Gaussian mixture models (Figueiredo and Jain, 2002). This is a form of continual learning, specifically transfer learning (Yang et al., 2020), where the knowledge of the mixture model is transferred to upstream models. Secondly, when optimising our utility function with respect to the recurrent reinforcement learner’s parameters, we do so sequentially online during the test set, and also make use of an extended Kalman filter optimisation procedure (Haykin, 2001), whereas the earlier work appeared to use less powerful, offline batch gradient ascent methods. Such offline methods cope less well with financial time series, which exhibit nonstationary behaviour. As discussed in Borrageiro et al. (2021), Merton (1976) modelled the dynamics of financial assets as a jump-diffusion process, which is now commonly used in financial econometrics. The jump-diffusion process implies that financial time series should observe small changes over time, so-called continuous changes, as well as occasional jumps. A sensible approach for coping with nonstationarity, is to allow models to continuously learn.

2. Background

In this section, we introduce the policy gradient form of reinforcement learning and how it has been put to work empirically in quantitative finance, in particular with automated trading strategies. We finish the section with a short review of more recent work.

2.1 Policy Gradient Reinforcement Learning

Williams (1992a) was one of the first to introduce policy gradient methods in a reinforcement learning context. Whereas the majority of reinforcement learning algorithms tend to focus on action value estimation, learning the value of and selecting actions based on their estimated action values, policy gradient methods learn a parameterised policy that can select actions without the use of a value function. Williams (1992a) introduced his *reinforce* algorithm

$$\Delta\theta_{ij} = \eta_{ij}(r - b_{ij}) \ln(\partial\pi_i/\partial\theta_{ij}),$$

where θ_{ij} is the model weight going from the j 'th input to the i 'th output and θ_i is the weight vector for the i 'th hidden processing unit of a network of such units, whose goal it is to adapt in such a way as to maximise the scalar reward r . For the moment, we exclude the dependence on time of the weight update so as to make the notation clearer. Furthermore, η_{ij} is a learning rate, typically applied with gradient ascent, b_{ij} is a reinforcement baseline, conditionally independent of the model outputs y_i , given the network parameters θ and inputs $\mathbf{x}_i$. $\ln(\partial\pi_i/\partial\theta_{ij})$ is known as the characteristic eligibility of θ_{ij} , where $\pi_i(y_i = c, \theta_i, \mathbf{x}_i)$, is a probability mass function determining the value of y_i as a function of the parameters of the unit and its input. Baseline subtraction $r - b_{ij}$ plays an important role in reducing the variance of gradient estimators and Sugiyama (2015) shows that the optimal baseline is given as

$$b^* = \frac{\mathbb{E}_{p(r|\theta)}[r_t \|\sum_{t=1}^T \nabla \ln \pi(a_t|s_t, \theta)\|^2]}{\mathbb{E}_{p(r|\theta)}[\|\sum_{t=1}^T \nabla \ln \pi(a_t|s_t, \theta)\|^2]},$$

where the policy function $\pi(a_t|s_t, \theta)$ denotes the probability of taking action a_t at time t given state s_t , parameterised by θ . The main result of Williams (1992a)'s paper is that for any *reinforce* algorithm, the inner product of $\mathbb{E}[\Delta\theta|\theta]$ and $\nabla\mathbb{E}[r|\theta]$ is non-negative and if $\eta_{ij} > 0$, then this inner product is zero iff $\nabla\mathbb{E}[r|\theta] = 0$. As η_{ij} is independent of i and j , then $\mathbb{E}[\Delta\theta|\theta] = \eta\nabla\mathbb{E}[r|\theta]$. This result relates $\nabla\mathbb{E}[r|\theta]$, the gradient in weight space of the performance measure $\mathbb{E}[r|\theta]$, to $\mathbb{E}[\Delta\theta|\theta]$, the average update vector in weight space. Thus for any *reinforce* algorithm, the average update vector in weight space lies in a direction for which this performance measure is increasing and the quantity $(r - b_{ij}) \ln(\partial\pi_i/\partial\theta_{ij})$ represents an unbiased estimate of $\partial\mathbb{E}[r|\theta]/\partial\theta_{ij}$.

Sutton and Barto (2018) demonstrate an *actor-critic* version of a policy gradient model, where the actor is a reference to the learned policy and the critic refers to the learned value function, usually a state-value function. Denote the scalar performance measure $J(\theta)$, then the gradient ascent update takes the form

$$\theta_{t+1} = \theta_t + \eta\nabla J(\theta).$$

In the one-step actor-critic policy gradient algorithm, one inputs a differentiable policy parameterisation $\pi(a|s, \theta)$ and a differentiable state-value function parameterisation $\hat{v}(s, \mathbf{w})$. One then draws an action

$$a_t \sim \pi(\cdot|s_t, \theta),$$

taking action a_t and observing a transition to state s_{t+1} with reward r_{t+1} . Define

$$\delta_t = r_{t+1} + \gamma\hat{v}(s_{t+1}, \mathbf{w}_t) - \hat{v}(s_t, \mathbf{w}_t),$$

where $0 \ll \gamma \leq 1$ is discount factor. The critic's weight vector is updated as follows

$$\mathbf{w}_t = \mathbf{w}_{t-1} + \eta_{\mathbf{w}} \delta_t \nabla_{\mathbf{w}} \hat{v}(s_t, \mathbf{w}_t),$$

and finally, the actor's weight vector is updated as

$$\boldsymbol{\theta}_t = \boldsymbol{\theta}_{t-1} + \eta_{\boldsymbol{\theta}} \delta_t \nabla \ln \pi(a_t | s_t, \boldsymbol{\theta}).$$

2.1.1 POLICY GRADIENT METHODS IN FINANCIAL TRADING

Moody et al. (1998) propose to train trading systems and portfolios by optimising objective functions that directly measure trading and investment performance. Rather than basing a trading system on forecasts, or training via a supervised learning algorithm using labelled trading data, they train their systems using a direct, recurrent reinforcement learning algorithm, which is an example of the policy gradient method. The 'direct' part refers to the fact that the model tries to target a position directly and the model's weights are adapted such that the performance measure is maximised. The performance function that they primarily consider is the differential Sharpe ratio. Denoting the annualised Sharpe ratio (Sharpe, 1966) as

$$sr_t = 252^{0.5} \times \frac{r_k - r_f}{s_k},$$

where r_k is the return of the k 'th strategy, with standard deviation s_k and r_f is the risk free rate. The differential Sharpe ratio is defined as

$$v_t = \frac{dsr_t}{d\tau} = \frac{b_{t-1} \Delta a_t - 0.5 a_{t-1} \Delta b_t}{(a_{t-1} - a_{t-1}^2)^{3/2}}, \quad (1)$$

where the quantities a_t and b_t are exponentially weighted estimates of the first and second moments of r_t

$$\begin{aligned} a_t &= a_{t-1} + \tau \Delta a_t = a_{t-1} + \tau(r_t - a_{t-1}) \\ b_t &= b_{t-1} + \tau \Delta b_t = b_{t-1} + \tau(r_t^2 - b_{t-1}), \end{aligned}$$

with $\tau \in (0, 1]$. They consider a batch gradient ascent update

$$\Delta \boldsymbol{\theta}_T = \eta \frac{dsr_T}{d\boldsymbol{\theta}},$$

where

$$\begin{aligned} \frac{dsr_T}{d\boldsymbol{\theta}} &= \sum_{t=1}^T \left\{ \frac{dsr}{da_T} \frac{da_T}{dr_t} + \frac{ds}{db_T} \frac{db_T}{dr_t} \right\} \left\{ \frac{dr_t}{df_t} \frac{df_t}{d\boldsymbol{\theta}} + \frac{dr_t}{df_{t-1}} \frac{df_{t-1}}{d\boldsymbol{\theta}} \right\} \\ &= \frac{1}{T} \sum_{t=1}^T \frac{dv_T}{dr_t} \left\{ \frac{dr_t}{df_t} \frac{df_t}{d\boldsymbol{\theta}} + \frac{dr_t}{df_{t-1}} \frac{df_{t-1}}{d\boldsymbol{\theta}} \right\} \\ &= \frac{1}{T} \sum_{t=1}^T \left\{ \frac{b_T - a_T r_t}{(b_T - a_T^2)^{3/2}} \right\} \left\{ \frac{dr_t}{df_t} \frac{df_t}{d\boldsymbol{\theta}} + \frac{dr_t}{df_{t-1}} \frac{df_{t-1}}{d\boldsymbol{\theta}} \right\}, \end{aligned}$$

and the reward

$$r_t = \Delta p_t f_{t-1} - \delta_t |\Delta f_t|,$$

depends on the change in reference price p_t from which a gross profit and loss (pnl) is computed, δ_t denotes transaction costs and $-1 \leq f_t \leq 1$ is a differentiable position function of the model inputs and parameters $f_t \triangleq f(\mathbf{x}_t, \boldsymbol{\theta}_t)$.

Trading and portfolio management systems require prior decisions as input in order to properly take into account the effect of transactions costs, market impact and taxes. This temporal dependence on system state requires the use of reinforcement versions of standard recurrent learning algorithms. They present empirical results in controlled experiments that demonstrate the efficacy of some of their methods for optimising trading systems and portfolios. For a long/short trader, they find that maximizing the differential Sharpe ratio yields more consistent results than maximizing profits and that both methods outperform a trading system based on forecasts that minimize mean square error. They find that portfolio trading agents trained to maximize the differential Sharpe ratio achieve better risk-adjusted returns than those trained to maximize profit. Finally, they provide simulation results for a S&P 500/TBill asset allocation system that demonstrate the presence of out of sample predictability in the monthly S&P 500 stock index for the 25 year period 1970 through 1994. They base their recurrent reinforcement learning algorithm on the real-time recurrent learning model of Williams and Zipser (1989). In their paper, Williams and Zipser note that their algorithm has the advantage that it does not require a precisely defined training interval, operating while the network runs and the disadvantage that it requires non-local communication in the network being trained and is computationally expensive. An undesirable property of the Sharpe ratio is that it penalises a model that produces returns larger than $\frac{\sigma_t}{\mu_t}$, that is, the ratio of the standard deviation of returns to its expectation.

Gold (2003) extends (Moody et al., 1998)’s work and investigates high frequency currency trading with neural networks trained via recurrent reinforcement learning. He compares the performance of linear networks, with neural networks containing a single hidden layer and examines the impact of shared system hyper-parameters on performance. In general, he concludes that the trading systems may be effective, but that the performance varies widely for different currency markets and this variability cannot be explained by simple statistics of the markets. He also finds that the linear recurrent reinforcement learners outperform the neural recurrent reinforcement learners in this application. On this last point, we suspect that the choice of inputs, that is, past returns of the target, results in features with weak predictive power. As a result, the neural reinforcement learner struggles to make meaningful forecasts. In contrast, the linear recurrent reinforcement learner does a better job of coping with both noisy inputs and outputs, generating biased, yet stable predictions. Gold also used shared hyper-parameters. Many of the currency pairs behave differently in terms of their price action. US dollar crosses tend to be momentum driven. Cross-currencies, such as the Australian dollar versus New Zealand dollar, tend to be mean-reverting in nature. Therefore, sharing hyper-parameters probably negatively impacts ex-poste performance here.

2.1.2 MORE RECENT WORK

In terms of more recent work involving policy gradient methods in finance, Tamar et al. (2017) discuss how several authors have developed risk-sensitive policy gradient methods that augment the standard expected cost minimization problem with a measure of variability in cost. These studies have focused on specific risk measures, such as the variance or conditional value at risk. Tamar et al. (2017) extend the policy gradient method to the whole class of coherent risk measures, which is widely accepted in finance and operations research, among other fields. They consider both static and time-consistent dynamic risk measures. For static risk measures, their approach is in the spirit of policy gradient algorithms and combines a standard sampling approach with convex programming. For dynamic risk measures, their approach is actor-critic style and involves explicit approximation of value functions.

Luo et al. (2019) build a novel reinforcement learning framework trader. They adopt an actor-critic algorithm called *deep deterministic policy gradient* to find the optimal policy. Their proposed algorithm has two different convolutional neural networks (Goodfellow et al., 2016) based function approximators and outperforms some baseline methods when experimenting with stock index futures. They discuss further the generalization and implications of the proposed method for business operations.

Zhang et al. (2019) use deep reinforcement learning algorithms such as deep q-learning networks (Mnih et al., 2013), neural policy gradients (Silver et al., 2016) and advantage actor-critic (Mnih et al., 2016), to design trading strategies for continuous futures contracts. They use long short-term memory neural networks (Hochreiter and Schmidhuber, 1997) to train both the actor and critic networks. Both discrete and continuous action spaces are considered and volatility scaling is incorporated to create reward functions which scale trade positions based on market volatility. They test their algorithms on the 50 most liquid futures contracts from 2011 to 2019, sampled monthly and investigate how performance varies across different asset classes including commodities, equity indices, fixed income and FX markets. They compare their algorithms against baseline models such as a long only strategy, the sign of annualised returns and a moving average convergence-divergence signal (Baz et al., 2015) and show that their method outperforms such baseline models, delivering positive profits despite heavy transaction costs. Their experiments show that the proposed algorithms can follow large market trends without changing positions and can also scale down, or hold, through consolidation periods. The data used in their experiment are sampled monthly, most likely due to the high computational cost of training the deep q-learning networks. This limits the practical use of these algorithms to a lower frequency (longer term forecasting) forms of trading, such as fund or asset management.

Ye et al. (2020) address the optimal trade execution problem using limit order books, where a model must learn how best to execute a given block of shares at minimal cost or for maximal return. To this end, they propose a deep reinforcement learning based solution which uses a deterministic policy gradient framework. Experiments on three real market datasets show that the proposed approach significantly outperforms other methods such as a submit and leave policy, a q-learning algorithm (Watkins, 1989) and a hybrid method that combines the Almgren-Chriss model (Almgren and Chriss, 2001) with reinforcement learning.

2.2 Foreign Exchange Trading

In this section, we provide a description of the foreign exchange (fx) market and the mechanics of the fx derivatives which are central to the experimentation that we conduct in section 4. The global foreign exchange market sees transactions in excess of 6 billion US dollars traded daily. Plot 1 shows this breakdown by instrument type and is extracted from the Bank of International Settlements Triennial Central Bank Survey, 2019.

FX transactions implicitly involve two currencies. The dominant or base currency is quoted by convention on the left hand side, and the secondary or counter currency is quoted by convention on the right hand side. For example, GBPUSD denotes the number of US dollars per one unit of UK Stirling. If one were to go long GBPUSD, one is in effect buying (investing in) GBP and selling (borrowing in) USD. As a result, for any overnight positions, the trader will earn the GBP interest rate whilst paying the USD interest rate. The interest rates for specific maturities are determined in the inter-bank currency market and are heavily influenced by the base rates typically set by central banks. By market convention, foreign exchange trades settle two business days ($t+2$) after the trade date, unless otherwise specified. This is commonly referred to as the value date for spot.

Figure 1: average daily global foreign exchange market turnover in in millions of US dollars

Clients fund their positions by rolling them forward via tomorrow/next (tomnext) swaps. Tomnext is a short-term fx transaction where a currency pair is simultaneously bought and sold over two separate business days: those being tomorrow (in one business day) and the following day (two business days from today). The tomnext transaction allows traders to maintain their position without being forced to take physical delivery and is the convention applied by prime brokers to their clients on the inter-bank fx market. In order to determine this funding cost, one needs to compute the forward prices, which are a combination of the spot rates (prices) and forward points. Forwards are an agreement between two counterparties to exchange currencies at a predetermined rate on some future date. Forwards can include outright and swaps. An outright is a spot transaction for a future date, beyond the current trade date. It is a derivative product consisting of a spot transaction combined with a forward spread. The spot portion of the transaction is more volatile than the forward portion, thus most of the price action will occur in that portion of the outright. To arrive at a forward rate at which to deal, forward points are applied to the spot rate. Forward points may be either positive or negative and are a function of the interest rate differential between the two currencies being dealt and the maturity of the trade. Forward points do not represent an expectation of the direction of a currency, but rather the interest rate differential. Forward points are computed and typically shown as follows

$$fpts = s_t(e_2 - e_1) \frac{t}{360\theta}$$

where s is the spot rate, e_2 is the secondary interest rate, e_1 is the dominant interest rate, t is the number of days till maturity and θ is the tick size or pip value for the associated currency pair. Example forward points for GBPUSD are shown in figure 2. $GBP=$ is the Refinitiv information code (ric) for cash GBPUSD and $GBPTN=$ is the ric for tomnext GBPUSD forward points. The tomnext outright points are computed as

$$\begin{aligned} bid_t^{tn} &= bid_t^{spot} + ask_t^{fpts} \theta \\ ask_t^{tn} &= ask_t^{spot} + bid_t^{fpts} \theta. \end{aligned}$$

As an example of rolling a long GBPUSD position forward, the tomnext swap would involve selling GBPUSD at bid_t^{spot} and buying it back at ask_t^{tn} . The cost of this 'roll' is thus $notional \times (bid_t^{spot} - ask_t^{tn})$, where *notional* denotes the size of the position taken by the trader. If a trader was short GBPUSD, then to roll the position forward, she would buy ask_t^{spot} and sell forward bid_t^{tn} , with the funding cost being $notional \times (bid_t^{tn} - ask_t^{spot})$. This funding may be a loss, but also a profit. There are many currency market participants that hold fx deliberately in order to capture the positive interest rate differential between two currency pairs. This is known as the carry trade and is extremely popular with the retail public in Japan, where the Yen interest rates have been historically low relative to other countries, for quite some time.

GBP FWD=		GBP FWD RATES								
RIC		Bid/Ask	Contributor	Loc	Src	Deal	Time	High	Low	Date
GBP=	↓	1.3545/1.3550	RADA FOREX	NYC	RADA		23:04	1.3545	1.3550	01OCT21
RIC		Bid/Ask	Contributor	Loc	Src	Deal	Time	High	Low	Date
GBPON=	↑	-0.1660/0.2450	NATIXIS	PAR	NXIR	NXIP	22:02	-0.1660	0.2270	01OCT21
GBPTN=	↓	-0.0040/0.0360	ZUERCHER KB	ZUR	ZKBZ	ZKBZ	9:37	0.0000	0.0100	02OCT21
GBPSN=	↑	0.0000/0.0100	LBBW	STG	LBBW	LBWF	9:36	0.0000	0.0100	02OCT21
GBPSW=	↓	0.05/0.07	LBBW	STG	LBBW	LBWF	9:38	0.09	0.03	02OCT21
GBP2W=	↓	0.18/0.26	ZUERCHER KB	ZUR	ZKBZ	ZKBZ	9:37	0.19	0.06	02OCT21
GBP3W=	↓	0.28/0.39	ZUERCHER KB	ZUR	ZKBZ	ZKBZ	9:37	0.31	0.36	02OCT21
GBP1M=	↑	0.54/0.59	LBBW	STG	LBBW	LBWF	9:35	0.58	0.58	02OCT21
GBP2M=	↑	0.45/0.55	LBBW	STG	LBBW	LBWF	9:35	0.45	0.37	02OCT21
GBP3M=	↑	3.81/3.97	LBBW	STG	LBBW	LBWF	9:36	3.81	3.85	02OCT21
GBP4M=	↑	2.92/3.12	LBBW	STG	LBBW	LBWF	9:35	2.92	3.10	02OCT21
GBP5M=	↑	1.10/1.35	LBBW	STG	LBBW	LBWF	9:35	1.10	1.06	02OCT21
GBP6M=	↑	-0.81/-0.51	LBBW	STG	LBBW	LBWF	9:35	-0.78	-0.62	02OCT21
GBP7M=	↓	-3.40/-2.95	LBBW	STG	LBBW	LBWF	9:35	-3.28	-3.05	02OCT21
GBP8M=	↓	-8.64/-3.46	ZUERCHER KB	ZUR	ZKBZ	ZKBZ	9:37	-7.35	-7.09	02OCT21
GBP9M=	↓	-11.06/-3.62	ZUERCHER KB	ZUR	ZKBZ	ZKBZ	9:37	-10.37	-9.80	02OCT21
GBP10M=	↓	-16.21/-2.96	ZUERCHER KB	ZUR	ZKBZ	ZKBZ	9:37	-14.45	-13.70	02OCT21
GBP11M=	↑	-18.82/-17.85	LBBW	STG	LBBW	LBWF	9:35	-18.63	-18.12	02OCT21
GBP1Y=	↑	-22.30/-21.20	LBBW	STG	LBBW	LBWF	9:35	-22.27	-21.77	02OCT21
GBP15M=	↑	-35.05/-31.50	NXIS	PAR	NXIP	NXIP	7:25	-33.00	-32.55	02OCT21
GBP18M=	↑	-44.58/-39.75	NXIS	PAR	NXIP	NXIP	7:25	-44.58	-40.23	02OCT21
GBP21M=	↑	-51.30/-45.68	NXIS	PAR	NXIP	NXIP	7:25	-51.30	-45.68	02OCT21
GBP2Y=	↓	-68.85/-36.76	STANDCHART	NYC	SCBN	SCON	23:10	-59.93	-44.93	01OCT21
GBP30M=	↓	-59.48/-36.98	BOA	NYC	BAFX	BANY	22:56	-59.48	-36.98	01OCT21
GBP3Y=	↓	-36.30/-28.30	COMMERZBANK	FFT	CBFW	CBFT	7:10	-36.10	-28.30	02OCT21
GBP4Y=	↑	27.70/37.70	COMMERZBANK	FFT	CBFW	CBFT	7:40	27.70	37.50	02OCT21
GBP5Y=	↑	115.00/125.00	COMMERZBANK	FFT	CBFW	CBFT	7:40	115.00	124.60	02OCT21
GBP6Y=	↑	214.00/234.00	COMMERZBANK	FFT	CBFW	CBFT	7:00	214.00	233.00	02OCT21
GBP7Y=	↑	325.00/345.00	COMMERZBANK	FFT	CBFW	CBFT	7:30	325.00	344.00	02OCT21
GBP8Y=	↑	426.00/456.00	COMMERZBANK	FFT	CBFW	CBFT	7:00	426.00	456.00	02OCT21
GBP9Y=	↑	531.00/561.00	COMMERZBANK	FFT	CBFW	CBFT	7:00	531.00	561.00	02OCT21
GBP10Y=	↓	483.00/771.00	ERSTE BANK	VIE		ERST	7:01	638.00	673.00	02OCT21

Figure 2: Refinitiv GBPUSD forward rates

3. Experiment Methods

In this section, we describe how our recurrent reinforcement learner targets a position directly. In addition, we also describe the baseline models that we compare and contrast experiment results against.

3.1 Targeting A Position With Direct Recurrent Reinforcement

Sharpe (2007) discusses asset allocation as a function of expected utility maximisation, where the utility function may be more complex than that associated with mean-variance analysis. Denote the expected utility for a single portfolio constituent as

$$v_t = \mu_t - \frac{\lambda}{2} \sigma_t^2, \quad (2)$$

where we omit the portfolio index to simplify the notation and where the expected return and variance of returns may be estimated in an online fashion with exponential decay

$$\begin{aligned}\mu_t &= \tau\mu_{t-1} + (1 - \tau)r_t \\ \sigma_t^2 &= \tau\sigma_{t-1}^2 + (1 - \tau)(r_t - \mu_t)^2.\end{aligned}$$

The risk appetite constant $\lambda > 0$, can be set as a function of an investor's desired risk adjusted return, as demonstrated by Grinold and Kahn (2019). Define the annualised information ratio as the risk adjusted differential reward measure, where the difference is taken with respect to a benchmark or baseline strategy

$$ir_k = 252^{0.5} \times \frac{\mu_k - \mu_b}{\sigma_k}.$$

The similarity to the Sharpe ratio is apparent. Substituting the non-annualised information ratio into the quadratic utility and differentiating with respect to the risk, we obtain a suitable value for the risk appetite parameter:

$$\begin{aligned}ir_t &= \frac{\mu_t}{\sigma_t} \\ v_t &= ir_t\sigma_t - \frac{\lambda}{2}\sigma_t^2 \\ \frac{dv_t}{d\sigma_t} &= ir_t - \lambda\sigma_t = 0 \\ \lambda &= \frac{ir_t}{\sigma_t}.\end{aligned}$$

The net returns whose expectation and variance we seek to learn, are decomposed as

$$r_t = \Delta p_t f_{t-1} - \delta_t |\Delta f_t| + \kappa_t, \quad (3)$$

where Δp_t is the change in reference price, typically a mid price

$$\Delta p_t = 0.5 \times (bid_t + ask_t - bid_{t-1} - ask_{t-1}),$$

δ_t represents the execution cost for a price taker

$$\delta_t = \max[0.5 \times (ask_t - bid_t), 0],$$

κ_t is the profit or loss of rolling the overnight fx position, the so called 'carry' (see section 2.2) and f_t is the desired position learnt by the recurrent reinforcement learner

$$f_t = \tanh(\theta_t^T \mathbf{x}_t). \quad (4)$$

The model is maximally short when $f_t = -1$ and maximally long when $f_t = 1$. The recurrency occurs in the input feature space where the previous position is fed to the model input

$$\mathbf{x}_t = [1, \phi_1(\mathbf{u}_t), \dots, \phi_m(\mathbf{u}_t), f_{t-1}]^T \in \mathbb{R}^{m+2},$$

and $\phi_j(\cdot)$ denotes a radial basis function hidden processing unit which takes as input a feature vector $\mathbf{u}_t$, see section 3.2. The goal of our recurrent reinforcement learner is to maximise the utility in equation 2, by targeting a position in equation 4. To do this, we apply an online stochastic gradient ascent update

$$\theta_t = \theta_{t-1} + \eta \nabla v_t \equiv \Delta \theta_t + \eta \frac{dv_t}{d\theta_t}.$$

Instead of a static learning rate η , we considered the Adam optimiser of Kingma and Ba (2017), where in effect an adaptive learning rate is applied. This adaptive learning rate is a function of the gradient expectation and variance. The weight update then takes the form

$$\begin{aligned}\mathbf{m}_t &= \beta_1 \mathbf{m}_{t-1} + (1 - \beta_1) \nabla v_t \\ \mathbf{v}_t &= \beta_2 \mathbf{v}_{t-1} + (1 - \beta_2) (\nabla v_t)^2 \\ \boldsymbol{\theta}_t &= \boldsymbol{\theta}_{t-1} + \eta \frac{\hat{\mathbf{m}}_t}{\hat{\mathbf{v}}_t^{0.5} + \epsilon},\end{aligned}$$

with $\hat{\mathbf{m}}_t = \mathbf{m}_t / (1 - \beta_1)$ and $\hat{\mathbf{v}}_t = \mathbf{v}_t / (1 - \beta_2)$ denoting bias corrected versions of the expected gradient and gradient variance respectively. In earlier work Bottou (2010) had considered approximating the Hessian of the performance measure with respect to the model weights, as a function of gradient only information. In practice, we found that Adam took many iterations of model fitting in the training set in order to get the weights to be large enough, such that a meaningfully sized position is taken via function 4. If the weights are too small, then the average position taken by the recurrent reinforcement learner will be small as well. Therefore, we finally settled on an extended Kalman filter (Williams (1992b), Haykin (2001)) style update, albeit modified for reinforcement learning in this context

$$\begin{aligned}z &= 1 + \nabla v_t^T \mathbf{P}_{t-1} \nabla v_t / \tau \\ \mathbf{k} &= \mathbf{P}_{t-1} \nabla v_t / (z\tau) \\ \boldsymbol{\theta}_t &= \boldsymbol{\theta}_{t-1} + \mathbf{k} \\ \mathbf{P}_t &= \mathbf{P}_{t-1} / \tau - \mathbf{k} \mathbf{k}^T z.\end{aligned}$$

Above, τ is an exponential decay factor, $\mathbf{P}_t$ is an approximation to $\nabla^2 v_t$, the inverse Hessian of the utility function v_t with respect to the model weights $\boldsymbol{\theta}_t$.

We decompose the gradient of the utility function with respect to the recurrent reinforcement learner's parameters as follows

$$\begin{aligned}\nabla v_t &= \frac{dv_t}{dr_t} \left\{ \frac{dr_t}{df_t} \frac{df_t}{d\boldsymbol{\theta}_t} + \frac{dr_t}{df_{t-1}} \frac{df_{t-1}}{d\boldsymbol{\theta}_{t-1}} \right\} \\ &= \frac{dv_t}{dr_t} \left\{ \frac{dr_t}{df_t} \left\{ \frac{\partial f_t}{\partial \boldsymbol{\theta}_t} + \frac{\partial f_t}{\partial f_{t-1}} \frac{\partial f_{t-1}}{\partial \boldsymbol{\theta}_{t-1}} \right\} + \frac{dr_t}{df_{t-1}} \left\{ \frac{\partial f_{t-1}}{\partial \boldsymbol{\theta}_{t-1}} + \frac{\partial f_{t-1}}{\partial f_{t-2}} \frac{\partial f_{t-2}}{\partial \boldsymbol{\theta}_{t-2}} \right\} \right\}.\end{aligned}\tag{5}$$

The constituent derivatives for the left half of equation 5 are:

$$\begin{aligned}\frac{dv_t}{dr_t} &= (1 - \eta)[1 - \lambda(r_t - \mu_t)] \\ \frac{dr_t}{df_t} &= -\delta_t \times \text{sign}(\Delta f_t) + \kappa_t \times \text{sign}(f_t) \\ \frac{df_t}{d\boldsymbol{\theta}_t} &= \mathbf{x}_t[1 - \tanh^2(\boldsymbol{\theta}_t^T \mathbf{x}_t)] + \boldsymbol{\theta}_{t,m+2}[1 - \tanh^2(\boldsymbol{\theta}_t^T \mathbf{x}_t)] \times \mathbf{x}_{t-1}[1 - \tanh^2(\boldsymbol{\theta}_{t-1}^T \mathbf{x}_{t-1})].\end{aligned}$$

3.2 Radial Basis Function Networks

In Borraigeiro et al. (2021), the authors show how online transfer learning in the form of radial basis function networks, where the hidden processing unit means and covariances are determined through clustering algorithms, provides a residual benefit in forecasting non-stationary time series, beyond the sequential fitting and prediction of recursive models. The feature engineering that we make use of in this paper, uses clusters formed of Gaussian mixture models, where the network size is determined by the unsupervised learning

procedure of finite mixture models described by Figueiredo and Jain (2002). We briefly describe the key ingredients of this meta-algorithm here. For greater detail, see Borraigeiro et al. (2021).

The radial basis function network is a network of $m > 0$ Gaussian basis functions

$$\phi_j(\mathbf{u}) = \exp\left(-\frac{1}{2}(\mathbf{u} - \boldsymbol{\mu}_j)^T \boldsymbol{\Sigma}_j^{-1}(\mathbf{u} - \boldsymbol{\mu}_j)\right).$$

Here we learn the j' th mean $\boldsymbol{\mu}_j$ and covariance $\boldsymbol{\Sigma}_j$ through a Gaussian mixture model fitting procedure. Denote the probability density function of a k component mixture as

$$p(\mathbf{u}|\boldsymbol{\theta}) = \sum_{j=1}^k \pi_j p(\mathbf{u}|\boldsymbol{\theta}_j) = \sum_{j=1}^k \pi_j \mathcal{N}(\mathbf{u}|\boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j),$$

where

$$\mathcal{N}(\mathbf{u}|\boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{1}{(2\pi)^{d/2} |\boldsymbol{\Sigma}|^{1/2}} \exp\left[-\frac{1}{2}(\mathbf{u} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\mathbf{u} - \boldsymbol{\mu})\right], \quad (6)$$

and the mixing weights satisfy $0 \leq \pi_j \leq 1$, $\sum_{j=1}^k \pi_j = 1$. The maximum likelihood estimate

$$\boldsymbol{\theta}_{ML} = \arg \max_{\boldsymbol{\theta}} \ln p(\mathbf{u}|\boldsymbol{\theta}),$$

and the Bayesian maximum a posteriori criterion

$$\boldsymbol{\theta}_{MAP} = \arg \max_{\boldsymbol{\theta}} \ln p(\mathbf{u}|\boldsymbol{\theta}) + \ln p(\boldsymbol{\theta}),$$

cannot be found analytically. The standard way of estimating $\boldsymbol{\theta}_{ML}$ or $\boldsymbol{\theta}_{MAP}$ is the expectation-maximisation algorithm (Dempster et al., 1977). This iterative procedure is based on the interpretation of $\mathbf{u}$ as incomplete data. The missing part for finite mixtures is the set of labels $\mathcal{Z} = \mathbf{z}_0, \dots, \mathbf{z}_n$ which accompany the training data $\mathbf{u}_0, \dots, \mathbf{u}_n$, indicating which component produced each training vector. Following Murphy (2012), let us define the complete data log likelihood to be

$$\ell_c(\boldsymbol{\theta}) = \sum_{i=1}^n \ln p(\mathbf{u}_i, \mathbf{z}_i|\boldsymbol{\theta}),$$

which cannot be computed, since $\mathbf{z}_i$ is unknown. Thus, let us define an auxiliary function

$$\mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}_{t-1}) = \mathbb{E}[\ell_c(\boldsymbol{\theta})|\mathbf{u}, \boldsymbol{\theta}_{t-1}],$$

where t is the current time step. The expectation is taken with respect to the old parameters $\boldsymbol{\theta}_{t-1}$ and the observed data $\mathbf{u}$. Denote as $r_{ic} = p(z_i = c|\mathbf{u}_i, \boldsymbol{\theta}_{t-1})$, the responsibility that cluster c takes for datum i . The expectation step has the following form

$$r_{ic} = \frac{\pi_c p(\mathbf{u}_i|\boldsymbol{\theta}_{c,t-1})}{\sum_{j=1}^k \pi_j p(\mathbf{u}_i|\boldsymbol{\theta}_{j,t-1})}.$$

The maximisation step optimises the auxiliary function $\mathcal{Q}$ with respect to $\boldsymbol{\theta}$

$$\boldsymbol{\theta}_t = \arg \max_{\boldsymbol{\theta}} \mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}_{t-1}).$$

The c' th mixing weight is estimated as

$$\pi_c = \frac{1}{n} \sum_{i=1}^n r_{ic} = \frac{r_c}{n}.$$

The parameter set $\theta_c = \{\mu_c, \Sigma_c\}$ is then

$$\begin{aligned}\mu_c &= \frac{\sum_{i=1}^n r_{ic} \mathbf{u}_i}{r_c} \\ \Sigma_c &= \frac{\sum_{i=1}^n r_{ic} (\mathbf{u}_i - \mu_c)(\mathbf{u}_i - \mu_c)^T}{r_c}.\end{aligned}$$

As discussed by Figueiredo and Jain (2002), expectation-maximisation is highly dependent on initialisation. They highlight several strategies to ameliorate this problem, such as multiple random starts, with final selection based on the highest maximum likelihood of the mixture, or kmeans based initialisation. However, in mixture models, the distinction between model-class selection and model estimation is unclear. For example, a 3 component mixture in which one of the mixing probabilities is zero, is indistinguishable for a 2 component mixture. They propose an unsupervised algorithm for learning a finite mixture model from multivariate data. Their approach is based on the philosophy of minimum message length encoding (Wallace and Dowe, 1999), where one aims to build a short code that facilitates a good data generation model. Their algorithm is capable of selecting the number of components and unlike the standard expectation-maximization algorithm, does not require careful initialization. The proposed method also avoids another drawback of expectation-maximization for mixture fitting: the possibility of convergence toward a singular estimate at the boundary of the parameter space. Denote the optimal mixture parameter set

$$\theta^* = \arg \min_{\theta} \ell_{FJ}(\theta, \mathbf{u}),$$

where

$$\ell_{FJ}(\theta, \mathbf{u}) = \frac{n}{2} \sum_{j=1}^k \ln \left(\frac{n\pi_k}{12} \right) + \frac{k}{2} \ln \left(\frac{n}{12} \right) + \frac{k(n+1)}{2} - \ln p(\mathbf{u}|\theta).$$

This leads to a modified maximisation step

$$\begin{aligned}\pi_c &= \frac{\max \left\{ 0, \left(\sum_{i=1}^n r_{ic} \right) - \frac{n}{2} \right\}}{\sum_{j=1}^k \max \left\{ 0, \left(\sum_{i=1}^n r_{ij} \right) - \frac{n}{2} \right\}} \\ &\quad \text{for } c = 1, 2, \dots, k.\end{aligned}$$

The maximisation step is identical to expectation-maximisation, except that the c' th parameter set θ_c is only estimated when $\pi_c > 0$ and θ_c is discarded from θ^* when $\pi_c = 0$. A distinctive feature of the modified maximisation step is that it leads to component annihilation. This prevents the algorithm from approaching the boundary of the parameter space. In other words, if one of the mixtures is not supported by the data, it is annihilated.

3.3 Baseline Models

In order to assess the comparative strength of the model of section 3.1, we employ two baseline models. The first model is a momentum trader, which uses the sign of the next step ahead return forecast as a target position. This supervised learner is based on exponentially

weighted recursive least squares, where the updates below are applied sequentially online:

$$\begin{aligned} z &= 1 + \mathbf{x}_t^T \mathbf{P}_{t-1} \mathbf{x}_t / \tau \\ \mathbf{k} &= \mathbf{P}_{t-1} \mathbf{x}_t / (z\tau) \\ \mathbf{w}_t &= \mathbf{w}_{t-1} + \mathbf{k}(y_t - \mathbf{w}_{t-1}^T \mathbf{x}_t) \\ \mathbf{P}_t &= \mathbf{P}_{t-1} / \tau - \mathbf{k} \mathbf{k}^T z \\ \hat{y}_t &= \mathbf{w}_t^T \mathbf{x}_t. \end{aligned}$$

The precision matrix $\mathbf{P}_0$ may be initialised to the identity matrix scaled by the inverse of the Ridge penalty, $\mathbf{I}_d \alpha^{-1}$ and the initial weights $\mathbf{w}_0$ are typically initialised to the zero vector. The discount factor τ , is typically close to but less than 1. This momentum trader is formulated as a radial basis function network in the feature space, where the features are mapped to the response, next step ahead mid to mid returns, via an exponentially weighted recursive least squares model. The particular model form is experimented with by Borrageiro et al. (2021) in a multi-step horizon forecasting context.

Our second baseline is the carry trader, which hopes to earn the positive differential overnight fx rate. Denoting the long and short carry as

$$\begin{aligned} \kappa_t^{long} &= bid_t^{spot} - ask_t^{tn} \\ \kappa_t^{short} &= bid_t^{tn} - ask_t^{spot}, \end{aligned}$$

where the superscript *spot* denotes the cash price and the superscript *tn* denotes the tomorrow/next price, the position of the carry trader is

$$f_t^{carry} = \begin{cases} sign(\kappa_t^{long} - \kappa_t^{short}), & \text{if } \kappa_t^{long} > 0 \text{ or } \kappa_t^{short} > 0 \\ 0 & \text{otherwise.} \end{cases}$$

In words, the carry trader goes long the base currency if the base currency has an overnight interest rate that is higher than the counter currency. Equally, the carry trader goes short the base currency if the base currency has an overnight interest rate that is lower than the counter currency. It is possible though that due to the bid/ask spread quoted by traders making markets in tomnext swaps, long and short carry will be negative. Therefore we allow the carry trader to abstain from trading completely in such circumstances.

4. Experiment Design

In this section, we establish the design of the experiment, beginning with a description of the data we use.

4.1 The Data

We obtain our experiment data from Refinitiv. We extract daily sampled data for 36 of the major cash fx pairs, for which there are available tomorrow / next forward points and outrights. These fx pairs are listed in table 1.

	spot rate (long name)	short name	ric	tn ric	pip value	start date	end date	n days
0	Australian Dollar/US Dollar FX	AUDUSD	AUD=	AUDTN=	1.00E-04	2006-12-14	2021-10-01	5405
1	Euro/Australian Dollar FX Cross Rate	EURAUD	EURAUD=	EURAUDTN=	1.00E-04	2009-09-23	2021-10-01	4391
2	Euro/Swiss Franc FX Cross Rate	EURCHF	EURCHF=	EURCHFTN=	1.00E-04	2009-04-15	2021-10-01	4552
3	Euro/Czech Koruna FX Cross Rate	EURCZK	EURCZK=	EURCZKTN=	1.00E-03	2006-09-20	2021-10-01	5490
4	Euro/Danish Krone FX Cross Rate	EURDKK	EURDKK=	EURDKKTN=	1.00E-04	2007-04-04	2021-10-01	5294
5	Euro/UK Pound Sterling FX Cross Rate	EURGBP	EURGBP=	EURGBPTN=	1.00E-04	2005-08-17	2021-10-01	5889
6	Euro/Hungarian Forint FX Cross Rate	EURHUF	EURHUF=	EURHUFTN=	1.00E-02	2006-09-11	2021-10-01	5499
7	Euro/Japanese Yen FX Cross Rate	EURJPY	EURJPY=	EURJPYTN=	1.00E-02	2008-02-06	2021-10-01	4986
8	Euro/Norwegian Krone FX Cross Rate	EURNOK	EURNOK=	EURNOKTN=	1.00E-04	2007-04-04	2021-10-01	5294
9	Euro/Polish Zloty FX Cross Rate	EURPLN	EURPLN=	EURPLNTN=	1.00E-04	2008-02-11	2021-10-01	4981
10	Euro/Swedish Krona FX Cross Rate	EURSEK	EURSEK=	EURSEKTN=	1.00E-04	2007-04-04	2021-10-01	5294
11	Euro/US Dollar FX	EURUSD	EUR=	EURTN=	1.00E-04	2005-08-17	2021-10-01	5889
12	UK Pound Sterling/US Dollar FX	GBPUSD	GBP=	GBPTN=	1.00E-04	2005-08-17	2021-10-01	5889
13	New Zealand Dollar/US Dollar FX	NZDUSD	NZD=	NZDTN=	1.00E-04	2006-12-14	2021-10-01	5405
14	US Dollar/Canadian Dollar FX	USDCAD	CAD=	CADTN=	1.00E-04	2005-08-17	2021-10-01	5889
15	US Dollar/Swiss Franc FX	USDCHF	CHF=	CHFTN=	1.00E-04	2006-12-14	2021-10-01	5405
16	US Dollar/Chinese Yuan Offshore FX	USDCNH	CNH=	CNHTN=	1.00E-04	2010-12-07	2021-10-01	3951
17	US Dollar/Czech Koruna FX	USDCZK	CZK=	CZKTN=	1.00E-03	2007-03-06	2021-10-01	5323
18	US Dollar/Danish Krone FX	USDDKK	DKK=	DKKTN=	1.00E-04	2007-03-06	2021-10-01	5323
19	US Dollar/Hong Kong Dollar FX	USDHKD	HKD=	HKDTN=	1.00E-04	1991-10-08	2021-10-01	10951
20	US Dollar/Hungarian Forint FX	USDHUF	HUF=	HUFTN=	1.00E-02	2007-03-06	2021-10-01	5323
21	US Dollar/Indonesian Rupiah FX	USDIDR	IDR=	IDRTN=	1.00E-02	1995-03-17	2021-10-01	9695
22	US Dollar/Israeli Shekel FX	USDILS	ILS=	ILSTN=	1.00E-04	2006-06-05	2021-10-01	5597
23	US Dollar/Indian Rupee FX	USDINR	INR=	INRTN=	1.00E-04	1999-05-27	2021-10-01	8163
24	US Dollar/Japanese Yen FX	USDJPY	JPY=	JPYTN=	1.00E-02	2005-08-17	2021-10-01	5889
25	US Dollar/Korean Won FX	USDKRW	KRW=	KRWTN=	1.00E-02	2004-10-18	2021-10-01	6192
26	US Dollar/Mexican Peso FX	USDMXN	MXN=	MXNTN=	1.00E-04	2007-01-19	2021-10-01	5369
27	US Dollar/Norwegian Krone FX	USDNOK	NOK=	NOKTN=	1.00E-04	2006-12-14	2021-10-01	5405
28	US Dollar/Polish Zloty FX	USDPLN	PLN=	PLNTN=	1.00E-04	2007-03-06	2021-10-01	5323
29	US Dollar/Russian Rouble FX	USDRUB	RUB=	RUBTN=	1.00E-04	2007-03-06	2021-10-01	5323
30	US Dollar/Swedish Krona FX	USDSEK	SEK=	SEKTN=	1.00E-04	2006-12-14	2021-10-01	5405
31	US Dollar/Singapore Dollar FX	USDSGD	SGD=	SGDTN=	1.00E-04	1991-10-08	2021-10-01	10951
32	US Dollar/Thai Baht FX	USDTHB	THB=	THBTN=	1.00E-03	1995-03-17	2021-10-01	9695
33	US Dollar/Turkish Lira FX	USDTRY	TRY=	TRYTN=	1.00E-04	2007-03-06	2021-10-01	5323
34	US Dollar/Taiwan Dollar FX	USDUSD	TWD=	TWDTN=	1.00E-03	2007-06-14	2021-10-01	5223
35	US Dollar/South African Rand FX	USDZAR	ZAR=	ZARTN=	1.00E-04	2006-09-14	2021-10-01	5496

Table 1: fx pairs

One of the challenges that the models will face in the experiment, is that these daily data show the last known top of book spot and outright prices at the end of the trading day, which is 5pm EST. The bid / ask spread for these prices are at their widest statistically at this time. Therefore, the execution cost will be higher, negative carry will be larger and positive carry will be smaller, than a trader who is able to trade at a more liquid time of day, such as 2pm GMT. If we try use intraday data, say data sampled minutely, then Refinitiv restricts us to 41 trading days, which isn't a very large sample size at all. Plot 3 illustrates the challenge succinctly. It shows bid/ask intraday spreads for both spot and outright (tomnext) for GBPUSD as quoted on Refinitiv in the past two calendar months. The data are sampled minutely. The global maximum bid / ask spread peak occurs exactly when the daily data are sampled.

4.2 Performance Evaluation

We have a little over 11 years of daily data to use in our experiment. From these data, we construct daily returns for each of the 36 currency pairs. We reserve the first third as a training set, feeding in the daily returns into the radial basis networks of section 3.2 for each target currency pair. We also use the training set to start learning the recurrent reinforcement learners' weights shown in equation 4. The carry trader baseline does not require any model fitting. In the test set, we continue to fit the radial basis function networks, momentum baseline and recurrent reinforcement learner, sequentially online.

Figure 3: Refinitiv GBPUSD bid/ask spreads by minute of day

We gather the test set pnls, net of execution and funding cost and finally aggregate them into a daily set of portfolio pnls.

5. Experiment Results

The carry baseline performs poorly, reflecting the low interest rate differential environment since the 2008 financial crisis. The momentum trader achieves the highest return with an annual compound net return of 11.7% and an information ratio of 0.4. The recurrent reinforcement learner achieves an annual compound net return of 9.3%, with an information ratio of 0.52. Its information ratio is driven higher by virtue of the fact that its standard deviation of daily portfolio returns, is two-thirds of that of the momentum trader. Table 2 shows a summary of net pnl returns statistics by strategy. Table 3 shows the funding or carry in returns space for each strategy. We can see that the carry baseline does indeed capture positive carry, although this return is not enough to offset the execution cost and the pnl associated with holding risk, which evidently moves in a trend-following way, largely opposite to the funding pnl. This is expected. When currencies depreciate considerably, central banks invariably increase overnight rates, in an attempt to make their currency more attractive and stem the tide of depreciation. The Turkish Lira and Russian Ruble are two cases in point. Finally, we see evidence in table 3 that the recurrent reinforcement learner captures more carry relative to the momentum trader. This too is expected, as the funding pnl makes its way into equation 3 and is propagated through the derivative of the utility function with respect to the model weights, as per equation 5.

	drl	mom	carry
count	1888.0	1888.0	1888.0
mean	0.001	0.001	-0.002
std	0.032	0.048	0.052
min	-0.141	-0.202	-0.344
25%	-0.019	-0.028	-0.028
50%	-0.0	-0.002	0.0
75%	0.019	0.028	0.026
max	0.245	0.423	0.2
sum	1.953	2.296	-4.328
ir	0.518	0.403	-0.701

Table 2: portfolio net pnl returns by strategy

	drl	mom	carry
count	1888.0	1888.0	1888.0
mean	-0.0003	-0.0005	0.00048
std	0.00019	0.00031	0.00036
min	-0.00395	-0.00576	7e-05
25%	-0.00035	-0.00059	0.00029
50%	-0.00024	-0.0004	0.00035
75%	-0.00019	-0.00032	0.00051
max	0.00153	0.00072	0.00518
sum	-0.56226	-0.94769	0.90655

Table 3: portfolio funding pnl returns by strategy

Figure 4: cumulative daily returns

Figure 5: distribution of daily returns

6. Discussion

Both baselines make decisions using incomplete information. The momentum trader focuses on learning the fx trends, but ignores the execution cost and the carry / funding cost. The carry trader tries to earn funding, but ignores both execution cost and both the depreciation and appreciation of the underlying currency pair. The recurrent reinforcement learner however, is optimising the desired position as a function of market moves, the cost of realising a position and the funding cost associated with overnight position rolls. To demonstrate that the recurrent reinforcement learner is indeed learning from these reward inputs, we compare the realised positions of a USDRUB trader where, in the former case, transaction costs and carry are removed (figure 6a) and in the latter case, transaction costs and carry are included (figure 6b). We see that without cost, the recurrent reinforcement learner realises broadly a long position (buying USD and selling RUB), as the Ruble depreciates over time. In contrast, even though the Ruble is depreciating, the differential overnight interest rate differential is roughly 6% and the recurrent reinforcement learner learns a short position (selling USD and buying RUB), capturing this positive carry. The positive carry is not enough to offset the rapid depreciation of the Ruble.

How significant are these results? Grinold and Kahn (2019) show a table of empirical information ratios. The results are for US data over the five-year period from January 2003 through December 2007. Empirical studies included 338 equity mutual funds, 1,679 equity long-only institutional funds, 56 equity long-short institutional funds and 537 fixed-income mutual funds. Although now a bit dated, the results indicate that our recurrent reinforcement learner and momentum baseline, who are trading statistically at the worst time of day in the fx market, achieve an information ratio at the 75'th percentile of information ratios achieved empirically by various passive and active fund managers within fixed income and equities.

asset class	equity	equity	equity	fixed income	
percentile	mutual funds	long	long short	institutional	average
90	1.04	0.77	1.17	0.96	0.99
75	0.64	0.42	0.57	0.50	0.53
50	0.20	0.02	0.25	0.01	0.12
25	-0.21	-0.38	-0.22	-0.45	-0.32
10	-0.62	-0.77	-0.58	-0.90	-0.72

Table 4: empirical information ratios, source: Blackrock

In terms of future work, one might consider a multi-layer perceptron version of our recurrent reinforcement learner. One might also consider an echo state network (Yildiz et al., 2012) version of the model. In addition, one might be able to improve the results further by applying a portfolio overlay. The utility function of equation 2, is readily treated as a portfolio problem

$$v_t = \mathbf{h}^T \boldsymbol{\mu}_t - \frac{\lambda}{2} \mathbf{h}^T \boldsymbol{\Sigma}_t \mathbf{h},$$

where the optimal, unconstrained portfolio weights are obtained by differentiating the portfolio utility with respect to the weight vector

$$\mathbf{h}^* = \frac{1}{\lambda} \boldsymbol{\Sigma}_t^{-1} \boldsymbol{\mu}_t.$$

Another approach is to treat portfolio selection as a reinforcement learning policy gradient problem, where the policy of picking actions, or this case portfolio constituents, is estimated via function approximation techniques.

(a) USDRUB trading without execution or funding cost

(b) USDRUB trading with execution and funding cost

7. Conclusion

We conduct a detailed experiment on major cash fx pairs, accurately accounting for transaction and funding costs. These sources of profit and loss, including the price trends that occur in the currency markets, are made available to our recurrent reinforcement learner via a quadratic utility (Sharpe, 2007), which learns to target a position directly. We improve upon the earlier work of Moody and Wu (1997), Moody and Saffell (1998), Moody et al. (1998) and Gold (2003), by casting the problem of learning to target a risk position in an online learning context. This online learning occurs sequentially in time, but also via transfer learning (Yang et al., 2020). This transfer learning takes the form of radial basis function hidden processing units, whose means, covariances and overall size are determined by an unsupervised learning procedure for finite Gaussian mixture models (Figueiredo and Jain, 2002). The intrinsic nature of the feature space is learnt and made available to the recurrent reinforcement learner and baseline momentum trader, an exponentially weighted recursive least squares model. The recurrent reinforcement learning trader achieves an annualised portfolio information ratio of 0.52 with compound return of 9.3%, net of execution and funding cost, over a 7 year test set. This is despite forcing the model to trade at the close of the trading day 5pm EST, when trading costs are statistically the most expensive. The momentum baseline trader achieves similar results. Thus, we cannot demonstrate in our experiment that this direct reinforcement learning approach works better than the traditional approach of minimising mean square prediction error, in the way that the momentum baseline trading model does. We put this down to the low interest rate differential environment that has been observed since the 2008 financial crisis, and the clear impact of momentum (trending) in currency returns. The recurrent reinforcement learner does however maintain an important advantage in that the model’s weights can be adapted to reflect the different sources of profit and loss variation, including returns momentum, transaction costs and funding costs. This is demonstrated visually in plots 6a and 6b, where a USDRUB trading agent learns to target different positions that reflect trading in the absence or presence of cost.

References

- R. Almgren and N. Chriss. Optimal execution of portfolio transactions. *Journal of Risk*, 3:5–40, 2001.
- J. Baz, N. M. Granger, C. R. Harvey, N. Le Roux, and S. Rattray. Dissecting investment strategies in the cross section and time series. *SSRN Electronic Journal*, 2015. ISSN 1556-5068. doi: 10.2139/ssrn.2695101.
- Y. Bengio. Using a financial training criterion rather than a prediction criterion. *International Journal of Neural Systems*, 08, 8 1997. ISSN 0129-0657.
- G. Borrageiro, N. Firoozye, and P. Barucca. Online learning with radial basis function networks, 2021.
- L. Bottou. Large-scale machine learning with stochastic gradient descent. In *Proceedings of the 19th International Conference on Computational Statistics (COMPSTAT’2010)*, pages 177–187, Paris, France, August 2010. Springer.
- A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39:1–38, 1977. ISSN 00359246.

- M. A. T Figueiredo and A. K Jain. Unsupervised learning of finite mixture models. *IEEE transactions on pattern analysis and machine intelligence*, 24:381–396, 2002. ISSN 0162-8828. doi: 10.1109/34.990138.
- C. Gold. Fx trading via recurrent reinforcement learning. In *IEEE*. IEEE, 2003.
- I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016.
- R.C. Grinold and R.N. Kahn. *Advances in Active Portfolio Management: New Developments in Quantitative Investing*. McGraw-Hill Education, 2019. ISBN 9781260453720.
- S. Haykin. *Kalman filtering and neural networks*. Wiley, 2001. ISBN 1-280-36756-3.
- S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- D. P. Kingma and J. Ba. Adam: A method for stochastic optimization, 2017.
- S. Luo, X. Lin, and Z. Zheng. A novel cnn-ddpg based ai-trader: Performance and roles in business operations. *Transportation research. Part E, Logistics and transportation review*, 131:68–79, 2019. ISSN 1366-5545.
- R. C. Merton. Option pricing when underlying stock returns are discontinuous. *Journal of financial economics*, 3:125–144, 1976.
- V. Mnih, Kavukcuoglu K, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller. Playing atari with deep reinforcement learning, 2013.
- V. Mnih, A. Puigdomènech Badia, M. Mirza, A. Graves, T. P. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu. Asynchronous methods for deep reinforcement learning, 2016.
- J. Moody and M. Saffell. Reinforcement learning for trading. In *Advances in Neural Information Processing Systems*, 1998.
- J. Moody and Lizhong Wu. Optimization of trading systems and portfolios. In *IEEE*, 1997.
- J. Moody, L. Wu, Y. Liao, and M. Saffell. Performance functions and reinforcement learning for trading systems and portfolios. *Journal of Forecasting*, 17:441–470, 1998.
- K. P. Murphy. *Machine learning a probabilistic perspective*. Adaptive Computation and Machine Learning. MIT Press, 2012.
- W. F. Sharpe. Mutual fund performance. *The Journal of Business*, 39, 1 1966.
- W. F. Sharpe. Expected utility asset allocation. *Financial Analysts Journal*, 63(5):18–30, 2007. ISSN 0015198X.
- D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis. Mastering the game of go with deep neural networks and tree search. *Nature*, 529, 1 2016. ISSN 0028-0836. doi: 10.1038/nature16961.
- M. Sugiyama. *Statistical Reinforcement Learning*. CRC Press, 1st edition edition, 2015. ISBN 9781439856901.

- R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. A Bradford Book, Cambridge, MA, USA, 2018.
- A. Tamar, Y. Chow, M. Ghavamzadeh, and S. Mannor. Sequential decision making with coherent risk. *IEEE transactions on automatic control*, 62(7):3323–3338, 2017. ISSN 0018-9286.
- R. S. Tsay and R. Chen. *Nonlinear time series analysis*. Wiley, 2019.
- C. S. Wallace and D. L. Dowe. Minimum Message Length and Kolmogorov Complexity. *The Computer Journal*, 42(4):270–283, 01 1999. ISSN 0010-4620. doi: 10.1093/comjnl/42.4.270.
- C. J. Watkins. *Learning from Delayed Rewards*. PhD thesis, King’s College, Oxford, 1989.
- R. J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8, 5 1992a. ISSN 0885-6125. doi: 10.1007/BF00992696.
- R. J. Williams and D. Zipser. A learning algorithm for continually running fully recurrent neural networks. *Neural computation*, 1:270–280, 1989.
- R.J. Williams. Training recurrent networks using the extended kalman filter. In *[Proceedings 1992] IJCNN International Joint Conference on Neural Networks*, volume 4, pages 241–246 vol.4, 1992b. doi: 10.1109/IJCNN.1992.227335.
- Q. Yang, Y. Zhang, W. Dai, and S. J. Pan. *Transfer Learning*. Cambridge University Press, 2020.
- Z. Ye, W. Deng, S. Zhou, Y. Xu, and J. Guan. Optimal trade execution based on deep deterministic policy gradient. In *Database Systems for Advanced Applications*, volume 12112 of *Lecture Notes in Computer Science*, pages 638–654. Springer International Publishing, 2020. ISBN 3030594092.
- I. B. Yildiz, H. Jaeger, and S. J. Kiebel. Re-visiting the echo state property. *Neural networks*, 35: 1–9, 2012. ISSN 0893-6080. doi: 10.1016/j.neunet.2012.07.005.
- Z. Zhang, S. Zohren, and S. Roberts. Deep reinforcement learning for trading. In *ArXiv*, 2019.