

XGBoost Learning of Dynamic Wager Placement for In-Play Betting on an Agent-Based Model of a Sports Betting Exchange

Chawin Terawong and Dave Cliff^a

Department of Computer Science, University of Bristol, Bristol BS8 1UB, U.K.

Keywords: Agent-Based Models, Sports Betting Exchanges, In-Play Betting, Dynamic Wager Placement, Machine Learning, XGBoost.

Abstract: We present first results from the use of XGBoost, highly effective machine learning (ML) method, within the *Bristol Betting Exchange* (BBE), an open-source agent-based model (ABM) designed to simulate a contemporary sports-betting exchange with in-play betting during track-racing events such as horse races. We use the BBE ABM and its array of minimally-simple bettor-agents as a synthetic data generator which feeds into our XGBoost ML system, with the intention that XGBoost discovers profitable dynamic betting strategies by learning from the more profitable bets made by the BBE bettor-agents. After this XGBoost training, which results in one or more decision trees, a bettor-agent with a betting strategy determined by the XGBoost-learned decision tree(s) is added to the BBE ABM and made to bet on a sequence of races under various conditions and betting-market scenarios, with profitability serving as the primary metric of comparison and evaluation. Our initial findings presented here show that XGBoost trained in this way can indeed learn profitable betting strategies, and can generalise to learn strategies that outperform each of the set of strategies used for creation of the training data. To foster further research and enhancements, the complete version of our extended BBE, including the XGBoost integration, has been made freely available as an open-source release on GitHub.


1 INTRODUCTION

Like many other long-standing aspects of human culture, despite its five-thousand-year history, gambling activity and opportunities were transformed by the rise of the World-Wide-Web in the dot-com boom of the late 1990s. One particular technology innovation from that time subsequently proved to be a seismic shift within the gambling industry: this was the arrival of commercial web-based *betting exchanges*.

In much the same way that financial markets such as stock exchanges offer platforms where potential buyers and potential sellers of a stock can interact to buy and sell shares, with buyers and sellers indicating their intended prices in *bid* and *ask* orders, which are then matched to compatible counterparties by the exchange's internal mechanisms, so betting exchanges are platforms where potential *backers* and potential *layers* can interact and be matched by the exchange, to find one or more people to take the other side of a bet. In the terminology of betting markets, a backer is someone who places a *back* bet, i.e. a bet which

will be paid if a specific event-outcome does occur; and a layer is someone who places a *lay* bet, i.e. a bet that's paid if the specific event-outcome does *not* occur. The revolutionary aspect of betting exchanges is that they operate as platform businesses: the exchange does not take a position as either a layer or a backer, it simply serves to match customers who want to back at a particular odds with other customers who want to lay at those same odds, and the exchange makes its money by taking a small fee from each customer's winnings. In contrast, traditional bookmakers (or "bookies") are the counterparty to each customer's bet, and lose money if they miscalculate their odds.

The first notably commercially successful sports betting exchange was created by British company BetFair (see www.betfair.com), a start-up which grew with explosive pace after its founding in 2000, and by 2006 was valued at £1.5billion. In 2016 Betfair merged with another gambling company, Paddy Power, in a deal worth £5bn, and the Betfair-branded component of the merged company (now known as Flutter Entertainment PLC) remains the world's largest online betting exchange to this day; at the time of writing this paper in late 2023, Flutter's market

^a  <https://orcid.org/0000-0003-3822-9364>

capitalization is £22.5billion. For further discussion of BetFair, see e.g. (Davies et al., 2005; Houghton, 2006; Cameron, 2009).

Creating an online exchange for matching layer and backer bettors was not the only innovation that BetFair introduced. They also led in the development of *in-play* betting, which allowed bettors to continue to place back and lay bets after a sports event had started, and to continue betting as the event progressed, until some pre-specified cut-off time or situation occurred, or the event finished. This is in contrast to conventional human-operated bookmakers, who ceased to take any further bets once the event of interest was underway: because Betfair’s betting exchange system was entirely automated, it could process large numbers of bets while an event is underway, operating in real time with flows of information that would overwhelm a human bookie.

Just as most stock-exchanges publish real-time summary data of all the bids and asks currently seeking a counterparty, often showing the quantity available to be bought or sold at each potential price for a particular stock, so a betting exchange publishes real-time summary data for any one event E showing all the currently unmatched backs and lays, the odds (or “price”) for each of them, and the amount of money available to be wagered at each price – in the terminology of betting exchanges, this collection of data is the “market” for event E .

During in-play betting, the prices in the market can shift rapidly, and while some types of events such as tennis matches might last for hours, allowing for hours of in-play betting to endure for a single match, for other types of event such as horse-racing the event may only last a few minutes. The exploratory work that we describe in this paper is motivated by the hypothesis that it may be possible to use machine learning (ML) methods to process the rapidly-changing data on a betting-exchange market for short-duration events such as horse races, and for the ML system to thereby produce novel profitable betting strategies.

For the rest of this paper, without loss of generality, we’ll limit our descriptions to talking only of betting on horse races because this is a very widely known form of sport on which much money is wagered, because the duration of most horse races is only a few minutes, and also because it is reasonably easy to create an appropriately realistic agent-based model (ABM), a simulation of a horse race, where each agent in the model represents a horse/rider combination, and where during the race each agent has a particular position on the track, is travelling at a specific velocity, and may or may not be blocked or otherwise influenced by other horses in the race. Ex-

actly such a simulation of a horse race was introduced by Cliff (Cliff, 2021), as one component of the *Bristol Betting Exchange* (BBE), an agent-based simulator not only of horse races, but also of a betting exchange, and also of a population of bettor-agents who each form their own private opinion of the outcome of a race, and then place back or lay bets accordingly. Various implementations of BBE have been reported previously by (Cliff et al., 2021) and at ICAART2022 by (Guzelyte and Cliff, 2022), but to the best of our knowledge ours is the first study to explore use of XGBoost (Chen and Guestrin, 2016) on in-play betting data to develop profitable betting strategies.

The bettor-agents in the BBE ABM each form their own private opinion on the outcome of a race on the basis of their own internal logic, i.e. their own individual betting strategy, and the original specification of BBE in (Cliff, 2021) included a number of minimally simple strategies, described in more detail in Section 2.3 below, and the BBE ABM usually operates with a bettor population having a heterogeneous mix of such betting strategies. As the dynamics of a simulated race unfold, so the bettor-agents react to changes in the competitors’s pace and relative positions by making and/or cancelling in-play bets, altering the market for that race. The BBE ABM records every change in the market over the duration of a race, along with the rank-order positions of the competitors at the time of each market change (i.e., which competitor is in first place, which is in second, and so on): this we refer to as a *race record*.

In the work described here, we typically run 1000 race simulations, gathering a race record from each. The set of race records then go through an automated analysis process to identify the actions of the most profitable bettors in each race. That is, for each race, we look to see which bettors made the most profit from in-play betting on that race, and we then work backward in time to see what actions those bettors took during the race, and what the state of the market and the state of the race was at the time of each such action. This then forms the training and/or test data for XGBoost: for any one item of such data, the input to XGBoost is the state of the market and the state of the race, and the desired output is the action that the bettor took.

To accomplish this, we modified the existing source-code of the most recent version of BBE, which is the multi-threaded BBE integrated with Opinion Dynamics Platform used in Guzelyte’s research (Guzelyte, 2021b; Guzelyte and Cliff, 2022), hosted on Guzelyte’s GitHub (Guzelyte, 2021a), to incorporate the XGBoost betting agent. After the integration of XGBoost, we conduct experiments to train and

evaluate the XGBoost bettor-agent’s performance in different market scenarios.

A surprising result we present here is that although XGBoost is trained on the profitable betting activity of a population of minimally simple betting strategies, the betting strategy that it then learns can outperform even the best of those simple strategies. That is, XGBoost generalises over the training data sufficiently well that the profitability of XGBoost-enabled bettor-agents can eventually be better than those of the non-learning bettors whose behaviors formed the training data for XGBoost.

Our work described here is exploratory: we use the BBE agent-based model (ABM) as a synthetic data generator to create the training data needed for XGBoost, and then we take the XGBoost-trained bettor agent and test it in the BBE ABM. We are doing this in an attempt to answer the research question of whether the multidimensional time series of data from the in-play betting market for horse races can in principle be fed into a machine learning system such as XGBoost and result in a learned profitable automated betting system. What we develop here is a proof-of-concept, and as we show in Section 4 our results thus far do show some promise, but we strongly caution against any readers of this paper actually gambling with real money on the basis of the system we describe here: a lot of further development work and much more extensive testing would be required before we would ever want to deploy this system live-betting with our own money.

Section 2 gives further background information, and then our experiment design is described in Section 3. Section 4 presents our results, followed by discussion of future work in Section 5.

2 BACKGROUND

2.1 BBE Race Simulator

Guzelyte’s research (Guzelyte, 2021b; Guzelyte and Cliff, 2022), relies heavily on Keen’s thesis (Keen, 2021) and Cliff’s original paper (Cliff, 2021) to create the Bristol Betting Exchange (BBE) race-simulator. BBE isn’t designed to perfectly mimic a real track horse-race, but rather to generate convincing data resembling real race dynamics: the changes in pace and rank-order position that occur within real races.

In every race simulation, competitors are selected from a pool and placed on a one-dimensional track of a fixed length. Each competitor’s position on the track at a given time is represented as a real-valued distance. The race begins at $t = 0$ and con-

cludes when the last competitor crosses the finish line. The progress of each competitor is calculated using a discrete-time stochastic process which provides for modelling of individual differences in pace (e.g., some might start the race at a fast pace but subsequently slow down; others might instead hold back in the early stages of the race and then speed up at the end) and for inter-competitor interactions such as a competitor being blocked and hence slowed by another competitor immediately in front, or a competitor being “hurried” by another competitor closing on to its rear. Full details of the race simulator were first published in (Cliff, 2021) to which the reader is referred for more information

2.2 BBE Betting Exchange

The Bristol Betting Exchange (BBE) uses a matching-engine that tracks all of the bets that are placed. The details of the bets of each bettor (time of bet, amount of bet, etc) is recorded for every race. If a bet hasn’t been matched with a counterparty, it can be cancelled, but once it’s matched, it can’t be. When it comes to matching bets, older bets are prioritized if the odds are the same. To create the market for a race, for each competitor BBE collects all back and lay bets at the same odds, calculating the total money bet. After a race finishes, the money from losing bets is gathered and given to the winners. BBE earns by taking a small commission from the winnings. BBE’s matching-engine is designed to implement exactly the same processes as are used in real-world betting exchanges, so in this sense it can be argued that the BBE’s exchange module is not just a simulation, but an actual instance of a betting exchange.

2.3 BBE Betting Agents

BBE has a variety of bettor-agents each utilizing a unique approach. These bettors’ strategies range from the wholly irrational *Zero Intelligence* (ZI) strategy, where choice of bets is made entirely at random; to wholly rational, where the best available information guides their decisions. The most rational strategy in BBE at present is the *Rational Predictor* (RP), which makes its race outcome predictions based on a series of simulated “dry-runs” of the race: at the start of the race, an RP bettor runs n independent and identically distributed (IID) private (i.e., known only to that RP bettor) simulations of the entire race from time $t = 0$ to whatever time the last horse crosses the finish-line, using the same race simulator engine as is used to implement the actual ‘public’ race that all BBE bettors are betting on; then at various times $t = t_i$ during the

race, the RP bettor will run another fresh set of n IID simulations of the current race, running the simulation forwards from the *current* positions of the horses at time t_i forward to the end of the race, and then may then make a fresh in-play bet if the most frequent winning horse in those n simulations is different from whatever horse it had previously bet on. In (Cliff, 2021), the behavior of an RP agent was defined as being determined primarily by the hyperparameter n , how many IID simulations it runs each time it reevaluates the odds, and so these agents were denoted as RP(n) bettors. Other authors who have worked with BBE since the publication of (Cliff, 2021) have found it useful to also be explicit about the time interval between an RP(n) bettor’s successive sets of n IID simulations: this can be captured by two additional hyperparameters, $\Delta_{r:\min}$ and $\Delta_{r:\max}$, such that the wait (in seconds) until the next set of n IID simulations is conducted by an RP bettor is given by a fresh draw from a uniform distribution between $\Delta_{r:\min}$ and $\Delta_{r:\max}$ (denoted by $\mathcal{U}[\Delta_{r:\min}, \Delta_{r:\max}]$) as that bettor concludes its current set of IID simulations. For this reason, an RP bettor is fully denoted by RP($n, \Delta_{r:\min}, \Delta_{r:\max}$).

The computational costs of simulating any one RP($n, \Delta_{r:\min}, \Delta_{r:\max}$) bettor agent over the duration of an entire race manifestly rises sharply as n increases and/or as the expected value $E(\mathcal{U}[\Delta_{r:\min}, \Delta_{r:\max}])$ falls. Authors such as (Keen, 2021; Guzelyte, 2021b) have concentrated on using the relatively low-computational-cost instance of RP(1,10,15), which — because this type of bettor is in receipt of privileged information — has come to be referred to as the *Privileged* bettor strategy. In the work presented here, we follow their convention and also use Privileged bettors as our form of RP agent.

There are several other types of BBE bettor strategies. The *Linear Extrapolator* (LinEx) employs a strategy of estimating competitor speed and predicting the outcome based on linear extrapolation. The *Leader Wins* (LW) bettor operates on the assumption that the leading competitor will maintain their position and win the race. The *Underdog* strategy (UD) supports the second-placed competitor as long as they are within a certain threshold of the lead. The *Back The Favourite* (BTF) bettor, on the other hand, aligns their predictions with the market’s favourite.

The *Representative Bettor* (RB) is a unique agent designed to mimic real-world human betting behaviours. It factors in betting preferences such as an inclination towards certain stake amounts, often seen in human bettors who prefer multiples of 2, 5, or 10. This bettor also exhibits the well-known *favorite-longshot bias*, reflecting the tendencies of human bettors to bet disproportionately on the favourite or the

longshot, regardless of the actual odds.

The presence of these various bettors, each with different parameters, within BBE gives rise to an engaging and complex dynamic in the in-play betting market. For full details and implementation notes on each of these bettor strategies, see (Cliff, 2021; Keen, 2021; Guzelyte and Cliff, 2022).

2.4 XGBoost Model Training

XGBoost, an abbreviation for eXtreme Gradient Boosting, introduced by (Chen and Guestrin, 2016), is an ML technique celebrated for its speed, precision, and flexibility. The algorithm operates on the gradient boosting framework, sequentially crafting decision trees that progressively enhance prediction accuracy. XGBoost has proven its effectiveness by frequently being used in winning solutions to international data science competitions, particularly on Kaggle, a competitive data science platform (Adebayo, 2020). Moreover, its real-world use extends to different fields like predictive modelling and recommendation systems, demonstrating its versatility. With its combination of computational power and predictive capability, XGBoost continues to drive advancements in machine learning.

Gradient boosting is a technique utilized in machine learning for both regression and classification problems. It operates by iteratively combining a series of simple predictive models, typically decision trees. Each subsequent model is designed to rectify the residual errors of its predecessor, thus enhancing accuracy incrementally (Friedman, 2001).

The technique gets its name from Gradient Descent, an optimization method used to minimize the chosen loss function. Every new model reduces the loss by moving in the direction of the steepest descent. It does this by incorporating a new tree that minimizes the loss most effectively, which is the essence of gradient boosting (Friedman, 2001). Unlike other boosting algorithms such as AdaBoost (Freund and Schapire, 1997), gradient boosting identifies the weaknesses of weak learners via gradients in the loss function, while AdaBoost does so by examining high-weight data points.

XGBoost (Chen and Guestrin, 2016) is designed to be highly scalable and parallelizable, making it suitable for handling large-scale datasets. The algorithm effectively manages sparse data and missing values without additional input. Additionally, XGBoost includes a regularization term into its objective function. This regularization term helps control the model complexity and avoid overfitting, a common problem found with other gradient boosting algorithm

(Friedman, 2001).

Space limitations prevent us from giving here full details of the XGBoost algorithm, for which the reader is instead referred to (Chen and Guestrin, 2016). For the purposes of this paper, it is sufficient to treat XGBoost as a “black-box” learning method that produces a set of decision trees that can be used as a betting strategy. We used the XGBoost Python library (Chen, 2023) which provides a flexible interface with the scikit-learn API (Pedregosa et al., 2011).

XGBoost’s performance can be further improved with a technique called parameter tuning. The algorithm’s parameters and hyperparameters are divided into three types: general parameters, booster parameters, and learning task parameters (Chen, 2023). General parameters control the overall function, booster parameters influence individual boosters, and learning task parameters oversee the optimization process. Example of hyperparameters include: maximum tree depth (`max_depth`); step size shrinkage (`learning_rate`), number of trees (`n_estimators`); minimum loss reduction (`gamma`); minimum sum of instance weight (`min_child_weight`); and the subsample ratio of training instances (`subsample`). Through the adjustment and tuning of these parameters, XGBoost can be adjusted to efficiently address a broad range of machine learning tasks.

2.5 Tuning and Cross Validation

Hyperparameter Tuning Using Grid Search. In machine learning, hyperparameters are important configurations that pre-determine the algorithm’s training process which influences the model’s final performance on prediction accuracy (Nyuytiybiy, 2020). In particular, within the scope of this research involving the XGBoost algorithm, key hyperparameters such as the learning rate and the maximum depth of the decision trees can significantly affect the model performance. The optimization of these hyperparameters is core to achieving the highest model prediction ability. A widely recognized technique to find the best set of hyperparameters is called ‘Hyperparameter Tuning’. Among the variety of methods available for this purpose, ‘Grid Search’ emerges as particularly robust. Grid Search conducts a methodical exploration of all potential combinations of hyperparameter values within a predefined boundary (list). This is for finding the combination of hyperparameters that offers optimal model performance (Malato, 2021).

Cross Validation (K-Fold Cross-Validation). K-fold cross-validation is a technique to determine the performance of a machine learning model. It involves partitioning the dataset into k equally-sized subsam-

ples. Each iteration uses one subsample for validation and the remaining $k - 1$ for training. The model undergoes k evaluations, with each subsample serving once as the validation set. The outcomes from the k tests are averaged to obtain a consolidated performance metric. This approach ensures every data point contributes to both training and validation, preventing overfitting of the model (Scikit-Learn, 2023a).

GridSearchCV. The GridSearchCV module from the scikit-learn library (Scikit-Learn, 2023b) integrates grid search and cross-validation, offering a streamlined mechanism for hyperparameter tuning. To use it, one provides a model (like XGBoost), a parameter grid defining the hyperparameter value combinations to test, a scoring method, and a number of k-fold cross-validations. GridSearchCV then systematically explores these combinations using cross-validation, where the dataset is partitioned into subsets and each subset is iteratively used for validation. Leveraging GridSearchCV with the XGBoost algorithm not only saves effort but also ensures the identification of optimal hyperparameters, enhancing model robustness and performance on unseen data.

3 EXPERIMENT DESIGN

3.1 Overview

The primary objective of this experiment is to integrate the XGBoost betting agent into the existing suite of agents in the BBE system. The agent will leverage a trained XGBoost model to make informed decisions on whether to ‘Back’ or ‘Lay’ a bet, predicated on the input data.

Figure 1 provides the high-level design of this experiment. The first step is to gather the data from BBE by running 1,000 race sessions. These race records are then pre-processed, narrowing down the data to the top 20 percent of the most significant transactions (Back or Lay actions). This refined dataset was then used for training with the XGBoost Python library. By tuning the Hyperparameters and using Cross-validation, the goal was to ensure the model could perform well in many scenarios.

Once a trained XGBoost model is ready, it is added as a new betting agent into the BBE system. To further test its performance, various race scenarios were run, each scenario with 100 races. This approach ensured we had enough data to assess how the new agent compared to the existing ones. Lastly, the collected data is used for statistical hypothesis tests to validate if the new XGBoost agent is more profitable.

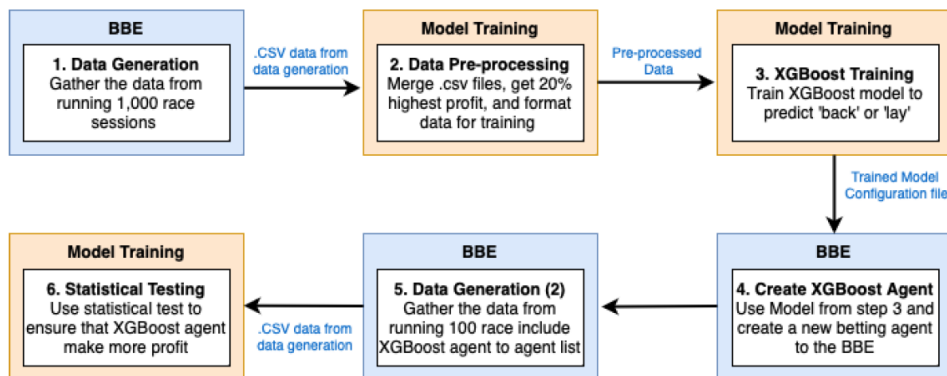


Figure 1: High-level overview of the experiment and the data flow of the system.

3.2 Setup for Model Training

Here we outline the specific scenario used for the data gathering process for XGBoost model training. Complete further details are available in (Terawong, 2023). We ran 1000 races, each of the same fixed distance so that the duration of each race was approximately the same, and all races had 5 competitors. The population of bettors in the ABM was made up from 10 each of Guzelyte’s (Guzelyte, 2021b; Guzelyte and Cliff, 2022) “opinionated” versions of the ZI, LW, BTF, LinEx, and Underdog strategies plus 5 of Guzelyte’s Opinionated-Privileged strategy; and then 55 of the original un-opinionated versions of these strategies, again split 10/10/10/10/10/5, for a grand total of 110 bettors.

3.3 Setup for Profit Validation

After implementing the XGBoost agent into the BBE system, data (including the data generated by the XGBoost agent) was gathered to evaluate whether the XGBoost agent generates more profit than the other agents. Two scenarios are created to experiment with this new XGBoost agent.

- **Scenario 1.** The number of simulations is reduced from 1,000 to 100, as this is only for profit testing and not for model training. A total of 5 XGBoost-trained bettor agents were added to the population.
- **Scenario 2.** Everything remains the same as in Scenario 1, except the number of agents is set to 5 for every agent type. This is for testing how the XGBoost agent performs when the environment changes.

3.4 XGBoost Parameters

We used the scikit-learn XGBoost API instead of the native XGBoost API. The native API of XGBoost provides a highly flexible and efficient way to train models, making it suitable for those experiment that prioritize performance and more refined configuration. On the other hand, the scikit-learn XGBoost API is a wrapper around this native API that integrates seamlessly with the widely used scikit-learn Python Library. This compatibility is the primary reason for selecting it in this research, mainly due to its seamless connection with GridSearchCV. This tool aids in hyperparameter tuning, an important aspect of model optimization. Moreover, the scikit-learn XGBoost API offers a user-friendly interface that reduces some complexities while still retaining robustness and enough flexibility.

In the model training process, specific choices shaped its direction. One crucial decision was the selection of the model’s objective function. This function dictates what the model aims to achieve during the learning process. In this work reported here the objective chosen was *binary:logistic*. This objective means the model is made to perform binary classification, determining an output as one of two distinct classes. The term logistic refers to the logistic function, mapping any input into a value between 0 and 1, making it suitable for probability estimation in binary decisions. Given the context of our betting decisions being binary (Back or Lay), the *binary:logistic* objective was a proper selection (Chen, 2023).

For evaluating the effectiveness of the model, logistic loss (“logloss”) was chosen as the metric (Chen, 2023). In machine learning, the choice of evaluation metric is crucial as it directly influences how the model’s performance is estimated and how it learns during training. The logloss is a measure for binary classification that quantifies the accuracy of a classifier by penalizing false classifications. It estimates the




Figure 2: Box-plot of the influence of XGBoost hyperparameter `eta` on the mean test score.




Figure 4: Box-plot of the influence of XGBoost hyperparameter `colsample_bytree` on the mean test score.




Figure 3: Box-plot of the influence of XGBoost hyperparameter `max_depth` on the mean test score.




Figure 5: Illustration of the influence of ‘gamma’ and ‘subsample’ on the mean test score, visualized using box plots. Left: the effects of `subsample`. Right: the effects of `gamma`.

probabilities associated with the accuracy of predictions. The closer the predicted probability is to the actual class, the lower the log loss.

Choice of hyperparameter values is crucial in shaping model performance. Hyperparameter tuning played a crucial role in the training, and we used GridSearchCV. The parameters tuned in our work are as follows: the learning rate `eta` which regulates step-size during boosting; `max_depth`, the maximum depth of the decision tree; `subsample`, the fraction of training data to be used in each boosting round; `colsample_bytree`, the fraction of features to be used for constructing each tree; `gamma`, the minimum loss reduction required to make a further partition of a leaf node in the decision tree, and `n_estimators`, which indicates the number of boosting rounds or trees to be constructed. This sequence is followed because `n_estimators` influences training time. Identifying the best hyperparameter set initially reduces the training time on finding the appropriate `n_estimators`.

For further details of the design and implementation of this extended version of BBE, see (Terawong, 2023).

4 RESULTS

4.1 Evaluation of XGBoost Training

4.1.1 Evaluation of Hyperparameters

The metric for evaluating hyperparameter performance in the experiment is the ‘accuracy’ score of the classification. Figures 2 to 5 present results from a 5-fold cross-validation combined with a hyperparameter tuning using grid search on the training dataset for the XGBoost model. Key insights are summarized as follows:

- Hyperparameters Impacting Model Performance.** As shown in Figures 2 to 4, the hyperparameters `eta`, `max_depth`, and `colsample_bytree` significantly influence model performance. An increase in the values of these hyperparameters generally correlates with an improved accuracy. Among them, `eta` demonstrates a pronounced effect, whereas `colsample_bytree` exhibits a more subtle impact.
- Hyperparameters with Minimal Impact.** Variations in hyperparameters like `subsample` and `gamma` seem to have little to no effect on the model’s performance according to Figure 5.




Figure 6: Illustration of the learning process of the model. It can be observed that the ‘LogLoss’ gradually decreased when the number of boosting round increases.

Through the process of cross-validation and hyperparameter tuning on the training data using the XGBoost machine learning algorithm, an optimal model was derived with a specific set of hyperparameters: `colsample_bytree=1.0`; `eta=0.3`; `gamma=0`; `max_depth=6`, and `subsample=1.0`.

4.1.2 Evaluation of ‘n estimators’ Parameters

After identifying the optimal set of hyperparameters, the next parameter under evaluation was n estimators. This parameter indicates the number of boosting rounds or trees that should be constructed. While in the code n estimators was initially set to 1,000, implying the intention to execute 1,000 boosting rounds, the inclusion of `early_stopping_rounds=10` ensures that training halts if the validation metric doesn’t demonstrate any enhancement over 10 successive boosting rounds. The combined usage of n estimators and early stopping rounds aids in mitigating both underfitting and overfitting of the model.

As depicted in Figure 6, the ‘LogLoss’ consistently decreases as the number of boosting rounds increases. This suggests that the model continues to learn and improve. The dashed line represents the early stopping point, beyond which the model no longer exhibits significant improvement. The model halted at the 452nd boosting round, as indicated by the early stopping mechanism. With these refinements, an optimal model has been obtained.

4.1.3 Evaluation of Optimal XGBoost Model

The SciKitLearn XGBoost training function that we use in this work provides insights into how different features contribute to the model’s predictions using the XGBoost F-score, which denotes the frequency with which a feature is used to split the data across all trees: this tell us the relative importance of each feature. The three most important features in our XGBoost bettor were:

Table 1: Confusion Matrix from the XGBoost training.

		Predicted Labels	
		lay	back
True Labels	lay	72521	1197
	back	9541	6730

1. Distance (F-score: 8680). As the most influential feature, ‘distance’ plays a central role in the model’s decision-making, indicating its significance in predicting the patterns the model identifies.
2. Time (F-score: 6642). The ‘time’ feature emerges as the second most influential feature, though it lags behind ‘distance’ by over 2000 points. Nevertheless, its considerable F-score denotes its relevance in the model’s predictions.
3. Rank (F-score: 1276). With a considerably lower F-score compared to the other two features, ‘rank’ seems to have a more marginal impact on the model’s decision process.

The confusion matrix of Table 1 summarizes the performance of the classification model. The following interpretations can be drawn:

- True Positive (6730). These instances correctly identify a ‘back’ bet. This means the model correctly predicted 6,730 instances where one should back a bet.
- True Negative (72521). These instances represent cases where the model correctly predicted a ‘lay’ bet. In this context, it means the model accurately identified 72,521 instances where one should not back the bet.
- False Positive (1197). These instances represent errors in prediction. The model mistakenly identified 1,197 bets as ‘back’ when they should have been ‘lay’.
- False Negative (9541). This count represents instances where the model wrongly classified bets as ‘lay’ when they should have been ‘back’. These could be seen as signifying areas where the model could be improved for better accuracy since the number is high.

According to this matrix, the model shows high accuracy in predicting when to place a lay bet because the number of true negatives is high. However, the number of false negatives indicates an area of potential improvement for the model, highlighting the potential to better recognise instances when the bettor should place a back bet.

The classification report, summarized in Table 2, provides in-depth metrics to assess the XGBoost

model's performance on the betting decision. Here, Class 0 corresponds to *lay* and Class 1 to *back*:

Precision (Class 0: 0.88, Class 1: 0.85). Precision measures how accurate the model's positive predictions are. For Class 0, a precision of 0.88 means that out of all predicted lays, 88% were correct. For Class 1, 85% of the model's predictions were correct.

Recall (Class 0: 0.98, Class 1: 0.41). Recall assesses the model's ability to detect all actual positives. For Class 0, the model identified 98% correctly. However, for Class 1, it was only 41%, showing room for improvement here.

F1 Score (Class 0: 0.93, Class 1: 0.56). F1-Score balances precision and recall. While it shows a high score for Class 0 of 93%, it has a moderate score for Class 1 of 56%.

Accuracy (0.88). The model correctly predicted the outcome for 88% of the bets. Referring to Table 2, it's clear that the model performs well for Class 0 because precision and recall are high. For Class 1, while precision remains high, recall drops, indicating challenges in detecting this class. However, the model's overall prediction accuracy is still high at 88%.

Table 2: Classification Report of the best model trained by XGBoost: Class 0 is lays; Class 1 is backs; *Acc.* is Accuracy; *MA* is Macro Average; and *WA* is Weighted Average.

Class	Precision	Recall	F1	Support
0	0.88	0.98	0.93	73718
1	0.85	0.41	0.56	16271
Acc.			0.88	89989
MA	0.87	0.70	0.74	89989
WA	0.88	0.88	0.86	89989

4.2 Hypothesis Testing

4.2.1 Scenario 1

Simulation Setup. The sessions were set at 100 rounds, incorporating various agents in predefined quantities. Specifically, 10 agents each of types ZI, LW, Underdog, BTF, and LinEx were introduced, along with 5 each of XGBoost and Privilege agents. This specific configuration was reminiscent of the one used during the model training phase, albeit here with the inclusion of the XGBoost agent.

In Figures 7 through 12, a consistent layout is used: the leftmost plot shows the average profit time-series comparison between XGBoost and its counterpart agent; the central plot exhibits the Kernel Density Estimation (KDE) contrasting XGBoost and the corresponding agent; the rightmost plot shows a box plot of this comparative data.

Statistical Examination. Visual inspection of the Kernel Density Estimation (KDE) plots of Figures 7 through 12, indicated that the data distributions were non-Normal, and in each case when we applied the Shapiro-Wilks test, the test outcome confirmed non-Normality. This prompted us to use the Wilcoxon-Mann-Whitney U-Test to with null hypothesis that there's no difference in the profit averages between the XGBoost agent and other agents, and alternate hypothesis that the XGBoost agent is more profitable. In all cases the null hypothesis is roundly rejected.

4.2.2 Scenario 2

Simulation Setup. Retaining the simulation sessions at 100 rounds, a different composition of agents was employed: 5 agents each for Random, Leader Wins, Underdog, Back Favourite, Linex, XGBoost, and Privilege.

Statistical Examination. Similar to Scenario 1, the data's non-normal distribution was confirmed in each case by the Shapiro-Wilk test and this is confirmed by visual inspection of the Kernel Density Estimation (KDE) plots of Figures 13 through 18. Consequently, the Wilcoxon-Mann-Whitney U test was used, and in each case the results from the U-test led to the rejection of the null hypothesis (the largest p value, for Privileged/XGBoost, was 0.0017), further emphasizing the XGBoost agent's performance.

Furthermore, from examination of the plots of Figures 13 through 18 it's also evident that the XGBoost betting agent consistently outperforms its peers, the same as in Scenario 1. The line graph highlights XGBoost's superior performance, with its values often trending higher. Similarly, the box plot emphasizes its strong placement, often residing in the upper range of outcomes.

In conclusion, for both scenarios, the XGBoost betting agent demonstrably outperformed its peers in terms of profit generation. Given these consistent results across different scenarios, it's obvious that the XGBoost agent, as modelled and implemented, offers a notable advantage in the context of this simulation.

5 FUTURE WORK

Numerous opportunities and potential areas remain for further investigation, including:

Variability of Data Collection. BBE, being a complex system, offers several scenarios and parameters setting that can influence outcomes. Each race's length, the number of competitors, and the mixture of participating agents all contribute uniquely to the fi-




Figure 7: Profit generated from XGBoost compare with Random Betting Agent for Scenario 1.




Figure 8: Profit generated from XGBoost compare with Leader Win Agent for Scenario 1.




Figure 9: Profit generated from XGBoost compare with Underdog Agent for Scenario 1.




Figure 10: Profit generated from XGBoost compare with Back Favorite Agent for Scenario 1.




Figure 11: Profit generated from XGBoost compare with Linex Agent for Scenario 1.




Figure 12: Profit generated from XGBoost compare with Privileged Agent for Scenario 1.




Figure 13: Profit generated from XGBoost compare with Random Betting Agent for Scenario 2.




Figure 14: Profit generated from XGBoost compare with Leader Win Agent for Scenario 2.




Figure 15: Profit generated from XGBoost compare with Underdog Agent for Scenario 2.




Figure 16: Profit generated from XGBoost compare with Back Favorite Agent for Scenario 2.




Figure 17: Profit generated from XGBoost compare with Linex Agent for Scenario 2.




Figure 18: Profit generated from XGBoost compare with Privilege Agent for Scenario 2.

nal dataset. Currently, the data extraction from BBE has been largely uniform. However, by introducing more randomness or systematically varying these parameters, it's possible to simulate a broader spectrum of race scenarios. Gathering data from these diverse conditions would likely provide a dataset with richer contextual information.

Feature Engineering. The model currently relies on four primary features is both a strength, for simplicity, and a limitation, for depth of insight. While distance, rank, time, and stake are crucial, there exist other features that might further refine the model's understanding. For instance, the rate of change of rank over time, interactions between distance and stake, or even cyclic patterns in betting behaviour could be potential features. Incorporating such sophisticated features could refine the model's decision boundaries and offer more precise predictions.

Model Optimization and Evaluation Metrics. The choice of the [binary:logistic] objective function has been pivotal for the model's current design, aiming for binary classification. However, XGBoost offers a large number of objective functions and evaluation metrics tailored for different kinds of predictive tasks. By experimenting with other objectives, such as 'multi:softmax' for multiclass problems or 'reg:squarederror' for regression tasks, new insights or even potential performance improvements could be achieved. This could also lead to the development of betting agents that can predict more than just binary outcomes, potentially increasing the versatility of the agent in different betting scenarios.

Expanded Testing Scenarios. Here we have used only two testing scenarios. However, the dynamic nature of the betting domain suggests the potential benefit of a more comprehensive evaluation, encompassing a broader spectrum of conditions and parameters. The performance and limitations of the XGBoost betting agent, in comparison to other available agents in the system, could be further illuminated under an array of diversified scenarios.

For instance, the impact of varying the number of competitors in a race could provide insights into the agent's robustness across different competitive landscapes. Similarly, the distance of races can influence outcome predictability, with certain agents potentially excelling in short sprints while others might have an edge in longer, more strategic races.

Furthermore, exploring races with different odds ranges can unveil how well the XGBoost betting agent navigates between high-risk, high-reward situations versus more conservative betting scenarios. Extending the testing to these more diverse scenarios would offer a richer, more comprehensive view of the XG-

Boost betting agent's capabilities, strengths, and potential areas for improvement

Online Learning and Feedback Loop Integration.

The field of online machine learning, where models learn on the go, adapting to new data as it arrives, offers a chance to improve the static nature of the current implementation. Instead of periodic manual data extraction and retraining, an integrated feedback loop would allow the XGBoost agent to continuously refine its strategies during every BBE session.

6 CONCLUSION

The primary contribution of this paper is the introduction of XGBoost learning to the bettor-agents in the BBE agent-based model, offering the opportunity to use BBE as a synthetic data generator and for XGBoost to then learn profitable betting strategies from the data provided from BBE. Comparing the XGBoost-learned betting strategy with the performance of the minimally simple strategies pre-coded into BBE demonstrates that XGBoost does indeed offer a distinct advantage in adaptively learning in-play betting strategies which are more profitable than any of the strategies that were used to create the training data. This serves as a proof-of-concept and in future work we intend to explore application of the methods described here to automatically learn betting strategies that could be profitable if deployed in betting on real-world races.

REFERENCES

- Adebayo, S. (2020). How the Kaggle winners algorithm XGBoost works. <https://dataaspirant.com/xgboost-algorithm>.
- Cameron, C. (2009). *You Bet: The Betfair Story; How Two Men Changed the World of Gambling*. Harper Collins.
- Chen, T. (2023). XGBoost Documentation. <https://xgboost.readthedocs.io/en/stable/index.html>.
- Chen, T. and Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD2016*, pages 785–794.
- Cliff, D. (2021). *BBE: Simulating the Microstructural Dynamics of an In-Play Betting Exchange via Agent-Based Modelling*. SSRN 3845698.
- Cliff, D., Hawkins, J., Keen, J., and Lau-Soto, R. (2021). Implementing the BBE agent-based model of a sports-betting exchange. In Affenzeller, M., Bruzzone, A., Longo, F., and Petrillo, A., editors, *Proceedings of the 33rd European Modelling and Simulation Symposium (EMSS2021)*, pages 230–240.

- Davies, M., Pitt, L., Shapiro, D., and Watson, R. (2005). Betfair.com: Five technology forces revolutionise worldwide wagering. *European Management Journal*, 23(5):533–541.
- Freund, Y. and Schapire, R. (1997). A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139.
- Friedman, J. (2001). Greedy function approximation: A gradient boosting machine. *The Annals of Statistics*, 29(5):1189–1232.
- Guzelyte, R. (2021a). BBE.OD: Threaded Bristol Betting Exchange with Opinion Dynamics. <https://github.com/Guzelyte/TBBE.OD>.
- Guzelyte, R. (2021b). Exploring opinion dynamics of agent-based bettors in an in-play betting exchange. Master’s thesis, Department of Engineering Mathematics, University of Bristol.
- Guzelyte, R. and Cliff, D. (2022). Narrative economics of the racetrack: An agent-based model of opinion dynamics in in-play betting on a sports betting exchange. In Rocha, A.-P., Steels, L., and van den Herik, J., editors, *Proceedings of the 14th International Conference on Agents and Artificial Intelligence (ICAART2022)*, volume 1, pages 225–236. Scitepress.
- Houghton, J. (2006). *Winning on Betfair for Dummies*. Wiley.
- Keen, J. (2021). Discovering transferable and profitable algorithmic betting strategies within the simulated microcosm of a contemporary betting exchange. Master’s thesis, University of Bristol, Department of Computer Science; SSRN 3879677.
- Malato, G. (2021). Hyperparameter Tuning, Grid Search and Random Search. <https://www.yourdata-teacher.com/2021/05/19/hyperparameter-tuning-grid-search-and-random-search/>.
- Nyuytiybiy, K. (2020). Parameters and hyperparameters in machine learning and deep learning. <https://towardsdatascience.com/parameters-and-hyperparameters-aa609601a9ac>.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Scikit-Learn, (2023a). Cross-validation: evaluating estimator performance. https://scikit-learn.org/stable/modules/cross_validation.html.
- Scikit-Learn, (2023b). SKLearn model selection: GridSearchCV. https://scikit-learn.org/stable/modules/generated/sklearn.model_selection.GridSearchCV.html.
- Terawong, C. (2023). An XGBoost Agent Based Model of In-Play Betting on a Sports Betting Exchange. Master’s thesis, Department of Computer Science, University of Bristol, UK.

APPENDIX: GITHUB REPOS

The integration of the XGBoost machine learning algorithm into BBE has been separated into two GitHub repositories, both of which are freely available as open-source Python code from: <https://github.com/ChawinT/>

Synthetic Data Generator

GitHub repo: `XGBoost_TBBE/tree/main`

1. Data Collection. Enhancements were made to the Bristol Betting Exchange (BBE) to facilitate an efficient data acquisition process.
2. XGBoost Betting Agent. A new component, named the XGBoost betting agent, was introduced within the `betting_agent.py` file. This serves as a blueprint for embedding machine learning capabilities into BBE.
3. Model Configuration. The `model.json` file encapsulates the trained XGBoost model utilized by the agent for bet predictions.

This repository lays the foundation for data collection and demonstrates a practical blueprint for integrating machine learning models into the BBE system.

Model Training & Validation

GitHub repo: `XGBoost_ModelTraining`

1. Model Training. Comprehensive training of the XGBoost model has been conducted, supplemented with optimization techniques and visualization tools.
2. Statistical Hypothesis Testing: Dedicated sections have been allocated for rigorous statistical hypothesis testing to validate the reliability of the results.

Together, these repositories form a comprehensive suite for introducing and harnessing the power of machine learning, specifically XGBoost, in the realm of betting on the BBE platform.